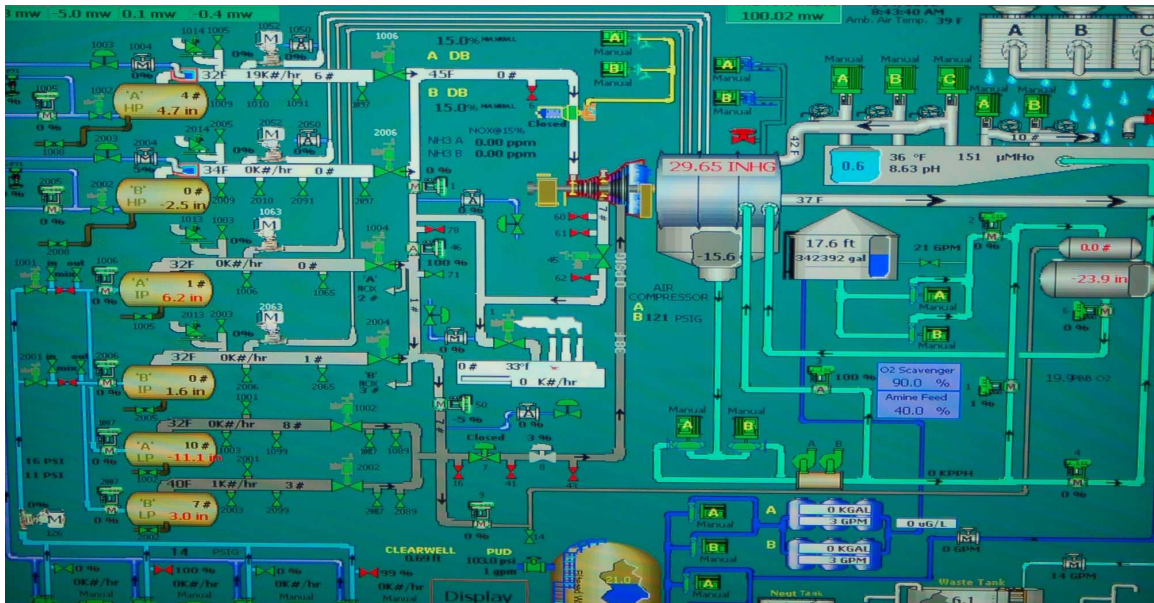


LESSONS IN INDUSTRIAL INSTRUMENTATION



© 2008-2018 by TONY R. KUPHALDT – UNDER THE TERMS AND CONDITIONS OF THE CREATIVE COMMONS ATTRIBUTION 4.0 INTERNATIONAL PUBLIC LICENSE

VERSION 2.28 (STABLE) – RELEASED 1 APRIL 2018

Lessons In Industrial Instrumentation © 2008-2018 by Tony R. Kuphaldt

This book is a copyrighted work, but licensed under the Creative Commons Attribution 4.0 International Public License. To view a copy of this license, turn to Appendix F, or visit <http://creativecommons.org/licenses/by/4.0/> or send a letter to Creative Commons, 171 Second Street, Suite 300, San Francisco, California, 94105, USA. The terms and conditions of this license allow for free copying, distribution, and/or modification of all licensed works by the general public.

Revision history¹

- Version 0.1 – July 2008 to September 2008 (initial development)
- Version 0.2 – released September 29, 2008 for Fall quarter student use (SafeCreative registration code 0810111072182)
- Version 0.4 – released January 12, 2009 for Winter quarter student use (SafeCreative registration code 0901122394919)
- Version 0.6 – released April 21, 2009 for public use (SafeCreative registration code 0904213106101)
- Version 0.8 – released September 8, 2009 for public use (SafeCreative registration code 0909094400980)
- Version 1.0 – released September 28, 2009 for public use (SafeCreative registration code 0909284601913)
- Version 1.2 – released January 12, 2010 for public use (SafeCreative registration code 1001125303242)
- Version 1.4 – released April 11, 2010 for public use (SafeCreative registration code 1004125976416)
- Version 1.6 – released May 4, 2010 for public use (SafeCreative registration code 1005046194859)
- Version 1.8 – released June 1, 2010 for public use (SafeCreative registration code 1006016477958)
- Version 1.10 – released June 28, 2010 for public use (SafeCreative registration code 1006286691573)
- Version 1.12 – released September 20, 2010 for public use (SafeCreative registration code 1009217397803)

¹Version numbers ending in odd digits are developmental (e.g. 0.7, 1.23, 4.5), with only the latest revision made accessible to the public. Version numbers ending in even digits (e.g. 0.6, 1.0, 2.14) are considered “public-release” and will be archived. Version numbers beginning with zero (e.g. 0.1, 0.2, etc.) represent early editions that were substantially incomplete.

- Version 1.14 – released January 11, 2011 for public use (SafeCreative registration code 1101118240391)
- Version 1.16 – released April 11, 2011 for public use (SafeCreative registration code 1104118950581)
- Version 1.18 – released July 4, 2011 for public use (SafeCreative registration code 1107069618739)
- Version 1.20 – released September 19, 2011 for public use (SafeCreative registration code 1109190093232)
- Version 1.22 – released January 8, 2012 for public use (SafeCreative registration code 1201090880285)
- Version 1.24 – released April 8, 2012 for public use (SafeCreative registration code 1204091446627)
- Version 1.26 – released June 30, 2012 for public use
- Version 1.28 – released September 18, 2012 for public use
- Version 1.30 – released January 5, 2013 for public use
- Version 1.32 – released April 1, 2013 for public use
- Version 1.34 – released July 1, 2013 for public use
- Version 1.36 – released September 23, 2013 for public use
- Version 2.00 – released January 12, 2014 for public use
- Version 2.02 – released April 14, 2014 for public use
- Version 2.04 – released August 8, 2014 for public use
- Version 2.06 – released September 22, 2014 for public use
- Version 2.08 – released January 12, 2015 for public use
- Version 2.10 – released April 7, 2015 for public use
- Version 2.12 – released June 29, 2015 for public use
- Version 2.14 – released September 21, 2015 for public use
- Version 2.16 – released January 1, 2016 for public use
- Version 2.18 – released April 11, 2016 for public use
- Version 2.20 – released September 16, 2016 for public use
- Version 2.22 – released April 8, 2017 for public use
- Version 2.24 – released July 4, 2017 for public use
- Version 2.26 – released September 25, 2017 for public use

Contents

Preface	1
1 Calculus	5
1.1 Introduction to calculus	6
1.2 The concept of differentiation	9
1.3 The concept of integration	14
1.4 How derivatives and integrals relate to one another	23
1.5 Symbolic versus numerical calculus	27
1.6 Numerical differentiation	31
1.7 Numerical integration	41
2 Physics	55
2.1 Terms and Definitions	56
2.2 Metric prefixes	57
2.3 Areas and volumes	57
2.3.1 Common geometric shapes	58
2.4 Unit conversions and physical constants	60
2.4.1 Unity fractions	61
2.4.2 Conversion formulae for temperature	64
2.4.3 Conversion factors for distance	65
2.4.4 Conversion factors for volume	65
2.4.5 Conversion factors for velocity	65
2.4.6 Conversion factors for mass	65
2.4.7 Conversion factors for force	65
2.4.8 Conversion factors for area	65
2.4.9 Conversion factors for pressure (either all gauge or all absolute)	66
2.4.10 Conversion factors for pressure (absolute pressure units only)	66
2.4.11 Conversion factors for energy or work	66
2.4.12 Conversion factors for power	66
2.4.13 Terrestrial constants	67
2.4.14 Properties of water	67
2.4.15 Miscellaneous physical constants	68
2.4.16 Weight densities of common materials	69
2.5 Dimensional analysis	71

2.6	The International System of Units	72
2.7	Conservation Laws	73
2.8	Classical mechanics	73
2.8.1	Newton's Laws of Motion	74
2.8.2	Work, energy, and power	75
2.8.3	Mechanical springs	91
2.8.4	Rotational motion	93
2.9	Simple machines	99
2.9.1	Levers	99
2.9.2	Pulleys	101
2.9.3	Inclined planes	104
2.9.4	Gears	106
2.9.5	Belt drives	117
2.9.6	Chain drives	120
2.10	Elementary thermodynamics	121
2.10.1	Heat versus Temperature	122
2.10.2	Temperature	123
2.10.3	Heat	125
2.10.4	Heat transfer	126
2.10.5	Specific heat and enthalpy	137
2.10.6	Phase changes	144
2.10.7	Phase diagrams and critical points	152
2.10.8	Saturated steam table	155
2.10.9	Thermodynamic degrees of freedom	158
2.10.10	Applications of phase changes	159
2.11	Fluid mechanics	170
2.11.1	Pressure	170
2.11.2	Pascal's Principle and hydrostatic pressure	175
2.11.3	Fluid density expressions	180
2.11.4	Manometers	182
2.11.5	Systems of pressure measurement	186
2.11.6	Negative pressure	189
2.11.7	Buoyancy	191
2.11.8	Gas Laws	197
2.11.9	Fluid viscosity	199
2.11.10	Reynolds number	201
2.11.11	Law of Continuity	205
2.11.12	Viscous flow	207
2.11.13	Bernoulli's equation	208
2.11.14	Torricelli's equation	218
2.11.15	Flow through a venturi tube	219

3	Chemistry	225
3.1	Terms and Definitions	228
3.2	Atomic theory and chemical symbols	230
3.3	Periodic table of the elements	236
3.4	Electronic structure	241
3.5	Spectroscopy	249
3.5.1	Emission spectroscopy	253
3.5.2	Absorption spectroscopy	256
3.6	Formulae for common chemical compounds	258
3.7	Molecular quantities	262
3.8	Stoichiometry	265
3.8.1	Balancing chemical equations by trial-and-error	266
3.8.2	Balancing chemical equations using algebra	268
3.8.3	Stoichiometric ratios	271
3.9	Energy in chemical reactions	273
3.9.1	Heats of reaction and activation energy	274
3.9.2	Heats of formation and Hess's Law	278
3.10	Periodic table of the ions	281
3.11	Ions in liquid solutions	282
3.12	pH	284
4	DC electricity	291
4.1	Electrical voltage	292
4.2	Electrical current	298
4.2.1	Electron versus conventional flow	301
4.3	Electrical sources and loads	306
4.4	Electrical power	312
4.5	Electrical resistance and Ohm's Law	313
4.6	Series versus parallel circuits	316
4.7	Kirchhoff's Laws	320
4.8	Circuit fault analysis	324
4.9	Bridge circuits	327
4.9.1	Component measurement	328
4.9.2	Sensor signal conditioning	330
4.10	Null-balance voltage measurement	336
4.11	Electromagnetism	340
4.12	Capacitors	346
4.13	Inductors	350
5	AC electricity	355
5.1	RMS quantities	356
5.2	Resistance, Reactance, and Impedance	360
5.3	Series and parallel circuits	361
5.4	Transformers	361
5.4.1	Basic principles	362
5.4.2	Loading effects	365

5.4.3	Step ratios	367
5.4.4	Transformer impedance	369
5.5	Phasors	371
5.5.1	Circles, sine waves, and cosine waves	372
5.5.2	Phasor expressions of phase shifts	377
5.5.3	Phasor expressions of impedance	386
5.5.4	Phasor arithmetic	390
5.5.5	Phasors and circuit measurements	394
5.6	The s variable	402
5.6.1	Meaning of the s variable	403
5.6.2	Impedance expressed using the s variable	406
5.7	Transfer function analysis	411
5.7.1	Example: LR low-pass filter circuit	412
5.7.2	Example: RC high-pass filter circuit	418
5.7.3	Example: LC “tank” circuit	423
5.7.4	Example: RLC band-pass filter circuit	428
5.7.5	Summary of transfer function analysis	437
5.8	Polyphase AC power	439
5.8.1	Delta and Wye configurations	447
5.8.2	Power in three-phase circuits	451
5.8.3	Grounded three-phase circuits	454
5.8.4	Symmetrical components	457
5.9	Phasor analysis of transformer circuits	466
5.9.1	Phasors in single-phase transformer circuits	468
5.9.2	Phasors in three-phase transformer circuits	474
5.10	Transmission lines	475
5.10.1	Open-ended transmission lines	477
5.10.2	Shorted transmission lines	479
5.10.3	Properly terminated transmission lines	481
5.10.4	Discontinuities	483
5.10.5	Velocity factor	483
5.10.6	Cable losses	484
5.11	Antennas	485
5.11.1	Maxwell and Hertz	489
5.11.2	Antenna size	491
5.11.3	Antenna orientation and directionality	492
6	Introduction to industrial instrumentation	497
6.1	Example: boiler water level control system	500
6.2	Example: wastewater disinfection	505
6.3	Example: chemical reactor temperature control	507
6.4	Other types of instruments	510
6.4.1	Indicators	511
6.4.2	Recorders	514
6.4.3	Process switches and alarms	520
6.5	Summary	529

6.6	Review of fundamental principles	530
7	Instrumentation documents	531
7.1	Process Flow Diagrams	533
7.2	Process and Instrument Diagrams	535
7.3	Loop diagrams	537
7.4	Functional diagrams	540
7.5	Instrument and process equipment symbols	543
7.5.1	Line types	544
7.5.2	Process/Instrument line connections	545
7.5.3	Instrument bubbles	545
7.5.4	Process valve types	546
7.5.5	Valve actuator types	547
7.5.6	Valve failure mode	548
7.5.7	Liquid level measurement devices	549
7.5.8	Flow measurement devices (flowing left-to-right)	550
7.5.9	Process equipment	552
7.5.10	Functional diagram symbols	553
7.5.11	Single-line electrical diagram symbols	554
7.5.12	Fluid power diagram symbols	556
7.6	Instrument identification tags	558
8	Instrument connections	563
8.1	Pipe and pipe fittings	563
8.1.1	Flanged pipe fittings	564
8.1.2	Tapered thread pipe fittings	571
8.1.3	Parallel thread pipe fittings	574
8.1.4	Sanitary pipe fittings	575
8.2	Tube and tube fittings	579
8.2.1	Compression tube fittings	580
8.2.2	Common tube fitting types and names	585
8.2.3	Bending instrument tubing	588
8.2.4	Special tubing tools	590
8.3	Electrical signal and control wiring	592
8.3.1	Connections and wire terminations	593
8.3.2	DIN rail	602
8.3.3	Cable routing	605
8.3.4	Signal coupling and cable separation	613
8.3.5	Electric field (capacitive) de-coupling	617
8.3.6	Magnetic field (inductive) de-coupling	623
8.3.7	High-frequency signal cables	626
8.4	Fiber optics	627
8.4.1	Fiber optic data communication	628
8.4.2	Fiber optic sensing applications	631
8.4.3	Fiber optic cable construction	636
8.4.4	Multi-mode and single-mode optical fibers	639

8.4.5	Fiber optic cable connectors, routing, and safety	641
8.4.6	Fiber optic cable testing	644
8.5	Review of fundamental principles	650
9	Discrete process measurement	653
9.1	“Normal” status of a switch	654
9.2	Hand switches	659
9.3	Limit switches	661
9.4	Proximity switches	663
9.5	Pressure switches	668
9.6	Level switches	673
9.6.1	Float-type level switches	674
9.6.2	Tuning fork level switches	676
9.6.3	Paddle-wheel level switches	677
9.6.4	Ultrasonic level switches	678
9.6.5	Capacitive level switches	679
9.6.6	Conductive level switches	680
9.7	Temperature switches	682
9.8	Flow switches	686
9.9	Review of fundamental principles	688
10	Discrete control elements	689
10.1	On/off valves	690
10.2	Fluid power systems	692
10.3	Solenoid valves	702
10.3.1	2-way solenoid valves	703
10.3.2	3-way solenoid valves	706
10.3.3	4-way solenoid valves	710
10.3.4	Normal energization states	715
10.4	On/off electric motor control circuits	718
10.4.1	AC induction motors	719
10.4.2	Motor contactors	728
10.4.3	Motor protection	730
10.4.4	Motor control circuit wiring	738
10.5	Review of fundamental principles	745
11	Relay control systems	747
11.1	Control relays	749
11.2	Relay circuits	753
11.3	Interposing relays	760
11.4	Review of fundamental principles	763

12 Programmable Logic Controllers	765
12.1 PLC examples	766
12.2 Input/Output (I/O) capabilities	775
12.2.1 Discrete I/O	777
12.2.2 Analog I/O	784
12.2.3 Network I/O	786
12.3 Logic programming	787
12.3.1 Relating I/O status to virtual elements	788
12.3.2 Memory maps and I/O addressing	796
12.4 Ladder Diagram (LD) programming	802
12.4.1 Contacts and coils	804
12.4.2 Counters	822
12.4.3 Timers	827
12.4.4 Data comparison instructions	832
12.4.5 Math instructions	835
12.4.6 Sequencers	838
12.5 Structured Text (ST) programming	848
12.6 Instruction List (IL) programming	848
12.7 Function Block Diagram (FBD) programming	848
12.8 Sequential Function Chart (SFC) programming	848
12.9 Human-Machine Interfaces	849
12.10 How to teach yourself PLC programming	855
12.11 Review of fundamental principles	858
13 Analog electronic instrumentation	861
13.1 4 to 20 mA analog current signals	861
13.2 Relating 4 to 20 mA signals to instrument variables	865
13.2.1 Example calculation: controller output to valve	868
13.2.2 Example calculation: flow transmitter	869
13.2.3 Example calculation: temperature transmitter	871
13.2.4 Example calculation: pH transmitter	874
13.2.5 Example calculation: reverse-acting I/P transducer signal	876
13.2.6 Example calculation: PLC analog input scaling	878
13.2.7 Graphical interpretation of signal ranges	881
13.2.8 Thinking in terms of per unit quantities	883
13.3 Controller output current loops	886
13.4 4-wire (“self-powered”) transmitter current loops	889
13.5 2-wire (“loop-powered”) transmitter current loops	891
13.6 4-wire “passive” versus “active” output transmitters	893
13.7 Troubleshooting current loops	894
13.7.1 Using a standard milliammeter to measure loop current	896
13.7.2 Using a clamp-on milliammeter to measure loop current	898
13.7.3 Using “test” diodes to measure loop current	899
13.7.4 Using shunt resistors to measure loop current	901
13.7.5 Troubleshooting current loops with voltage measurements	902
13.7.6 Using loop calibrators	906

13.7.7	NAMUR signal levels	913
13.8	Review of fundamental principles	914
14	Pneumatic instrumentation	915
14.1	Pneumatic sensing elements	921
14.2	Self-balancing pneumatic instrument principles	923
14.3	Pilot valves and pneumatic amplifying relays	928
14.4	Analogy to opamp circuits	938
14.5	Analysis of practical pneumatic instruments	950
14.5.1	Foxboro model 13A differential pressure transmitter	951
14.5.2	Foxboro model E69 “I/P” electro-pneumatic transducer	956
14.5.3	Fisher model 546 “I/P” electro-pneumatic transducer	961
14.5.4	Fisher-Rosemount model 846 “I/P” electro-pneumatic transducer	966
14.6	Proper care and feeding of pneumatic instruments	969
14.7	Advantages and disadvantages of pneumatic instruments	970
14.8	Review of fundamental principles	971
15	Digital data acquisition and networks	973
15.1	Digital representation of numerical data	977
15.1.1	Integer number formats	978
15.1.2	Fixed-point number formats	981
15.1.3	Floating-point number formats	982
15.1.4	Example of industrial number formats	984
15.2	Digital representation of text	987
15.2.1	Morse and Baudot codes	988
15.2.2	EBCDIC and ASCII	989
15.2.3	Unicode	991
15.3	Analog-digital conversion	991
15.3.1	Converter resolution	992
15.3.2	Converter sampling rate and aliasing	996
15.4	Analog signal conditioning and referencing	999
15.4.1	Instrumentation amplifiers	1000
15.4.2	Analog input references and connections	1009
15.5	Digital data communication theory	1022
15.5.1	Serial communication principles	1024
15.5.2	Physical encoding of bits	1027
15.5.3	Communication speed	1030
15.5.4	Data frames	1032
15.5.5	Channel arbitration	1040
15.5.6	The OSI Reference Model	1046
15.6	EIA/TIA-232, 422, and 485 networks	1049
15.6.1	EIA/TIA-232	1050
15.6.2	EIA/TIA-422 and EIA/TIA-485	1054
15.7	Ethernet networks	1061
15.7.1	Repeaters (hubs)	1062
15.7.2	Ethernet cabling	1065

15.7.3	Switching hubs	1069
15.8	Internet Protocol (IP)	1071
15.8.1	IP addresses	1073
15.8.2	Subnetworks and subnet masks	1077
15.8.3	Routing tables	1081
15.8.4	IP version 6	1082
15.8.5	ARP	1083
15.8.6	DNS	1084
15.8.7	Command-line diagnostic utilities	1084
15.9	Transmission Control Protocol (TCP) and User Datagram Protocol (UDP)	1088
15.10	The HART digital/analog hybrid standard	1091
15.10.1	Basic concept of HART	1092
15.10.2	HART physical layer	1097
15.10.3	HART multidrop mode	1104
15.10.4	HART multi-variable transmitters and burst mode	1105
15.11	Modbus	1106
15.11.1	Modbus overview	1107
15.11.2	Modbus data frames	1111
15.11.3	Modbus function codes and addresses	1113
15.11.4	Modbus relative addressing	1115
15.11.5	Modbus function command formats	1118
15.11.6	Floating-point values in Modbus	1126
15.12	Review of fundamental principles	1129
16	FOUNDATION Fieldbus instrumentation	1133
16.1	FF design philosophy	1134
16.2	H1 FF Physical layer	1140
16.2.1	Segment topology	1141
16.2.2	Coupling devices	1145
16.2.3	Electrical parameters	1148
16.2.4	Cable types	1150
16.2.5	Segment design	1153
16.3	H1 FF Data Link layer	1155
16.3.1	Device addressing	1157
16.3.2	Communication management	1158
16.3.3	Device capability	1167
16.4	FF function blocks	1167
16.4.1	Analog function blocks versus digital function blocks	1168
16.4.2	Function block location	1169
16.4.3	Standard function blocks	1174
16.4.4	Device-specific function blocks	1176
16.4.5	FF signal status	1177
16.4.6	Function block modes	1179
16.5	H1 FF device configuration and commissioning	1180
16.5.1	Configuration files	1181
16.5.2	Device commissioning	1183

16.5.3	Calibration and ranging	1192
16.6	H1 FF segment troubleshooting	1200
16.6.1	Cable resistance	1200
16.6.2	Signal strength	1201
16.6.3	Electrical noise	1201
16.6.4	Using an oscilloscope on H1 segments	1202
16.6.5	Message re-transmissions	1204
16.7	Review of fundamental principles	1204
17	Wireless instrumentation	1207
17.1	Radio systems	1207
17.1.1	Antennas	1208
17.1.2	Decibels	1210
17.1.3	Antenna radiation patterns	1217
17.1.4	Antenna gain calculations	1220
17.1.5	Effective radiated power	1222
17.1.6	RF link budget	1224
17.1.7	Fresnel zones	1230
17.2	WirelessHART	1232
17.2.1	Introduction to <i>WirelessHART</i>	1233
17.2.2	<i>WirelessHART</i> network protocol	1236
17.2.3	<i>WirelessHART</i> network gateway device	1240
17.2.4	<i>WirelessHART</i> device commissioning and configuration	1245
17.3	Review of fundamental principles	1249
18	Instrument calibration	1251
18.1	Calibration versus re-ranging	1251
18.2	Zero and span adjustments (analog instruments)	1252
18.3	Calibration errors and testing	1254
18.3.1	Typical calibration errors	1255
18.3.2	As-found and as-left documentation	1260
18.3.3	Up-tests and Down-tests	1261
18.3.4	Automated calibration	1262
18.4	Damping adjustments	1265
18.5	LRV and URV settings, digital trim (digital transmitters)	1269
18.6	An analogy for calibration versus ranging	1275
18.7	Calibration procedures	1276
18.7.1	Linear instruments	1278
18.7.2	Nonlinear instruments	1279
18.7.3	Discrete instruments	1280
18.8	Instrument turndown	1281
18.9	NIST traceability	1282
18.10	Practical calibration standards	1285
18.10.1	Electrical standards	1287
18.10.2	Temperature standards	1289
18.10.3	Pressure standards	1296

18.10.4	Flow standards	1301
18.10.5	Analytical standards	1302
18.11	Review of fundamental principles	1306
19	Continuous pressure measurement	1307
19.1	Manometers	1308
19.2	Mechanical pressure elements	1314
19.3	Electrical pressure elements	1320
19.3.1	Piezoresistive (strain gauge) sensors	1321
19.3.2	Differential capacitance sensors	1326
19.3.3	Resonant element sensors	1334
19.3.4	Mechanical adaptations	1337
19.4	Force-balance pressure transmitters	1338
19.5	Differential pressure transmitters	1341
19.5.1	DP transmitter construction and behavior	1342
19.5.2	DP transmitter applications	1349
19.5.3	Inferential measurement applications	1355
19.6	Pressure sensor accessories	1358
19.6.1	Valve manifolds	1359
19.6.2	Bleed (vent) fittings	1365
19.6.3	Pressure pulsation damping	1368
19.6.4	Remote and chemical seals	1371
19.6.5	Filled impulse lines	1381
19.6.6	Purged impulse lines	1383
19.6.7	Heat-traced impulse lines	1387
19.6.8	Self-purged impulse lines	1390
19.6.9	Water traps and pigtail siphons	1393
19.6.10	Mounting brackets	1395
19.6.11	Heated enclosures	1396
19.7	Process/instrument suitability	1398
19.8	Review of fundamental principles	1400
20	Continuous level measurement	1403
20.1	Level gauges (sightglasses)	1403
20.1.1	Basic concepts of sightglasses	1404
20.1.2	Interface problems	1407
20.1.3	Temperature problems	1409
20.2	Float	1410
20.3	Hydrostatic pressure	1416
20.3.1	Pressure of a fluid column	1417
20.3.2	Bubbler systems	1422
20.3.3	Transmitter suppression and elevation	1426
20.3.4	Compensated leg systems	1430
20.3.5	Tank expert systems	1438
20.3.6	Hydrostatic interface level measurement	1443
20.4	Displacement	1449

20.4.1	Buoyant-force instruments	1450
20.4.2	Torque tubes	1455
20.4.3	Displacement interface level measurement	1462
20.5	Echo	1466
20.5.1	Ultrasonic level measurement	1468
20.5.2	Radar level measurement	1473
20.5.3	Laser level measurement	1485
20.5.4	Magnetostrictive level measurement	1486
20.6	Weight	1490
20.7	Capacitive	1495
20.8	Radiation	1497
20.9	Level sensor accessories	1502
20.10	Review of fundamental principles	1506
21	Continuous temperature measurement	1509
21.1	Bi-metal temperature sensors	1511
21.2	Filled-bulb temperature sensors	1513
21.3	Thermistors and Resistance Temperature Detectors (RTDs)	1517
21.3.1	Temperature coefficient of resistance (α)	1518
21.3.2	Two-wire RTD circuits	1520
21.3.3	Four-wire RTD circuits	1521
21.3.4	Three-wire RTD circuits	1522
21.3.5	Proper RTD sensor connections	1525
21.3.6	Self-heating error	1528
21.4	Thermocouples	1529
21.4.1	Dissimilar metal junctions	1530
21.4.2	Thermocouple types	1532
21.4.3	Connector and tip styles	1533
21.4.4	Manually interpreting thermocouple voltages	1537
21.4.5	Reference junction compensation	1539
21.4.6	Law of Intermediate Metals	1543
21.4.7	Software compensation	1547
21.4.8	Extension wire	1549
21.4.9	Side-effects of reference junction compensation	1554
21.4.10	Burnout detection	1560
21.5	Non-contact temperature sensors	1561
21.5.1	Concentrating pyrometers	1562
21.5.2	Distance considerations	1565
21.5.3	Emissivity	1570
21.5.4	Thermal imaging	1570
21.6	Temperature sensor accessories	1574
21.7	Process/instrument suitability	1578
21.8	Review of fundamental principles	1579

22 Continuous fluid flow measurement	1581
22.1 Pressure-based flowmeters	1582
22.1.1 Venturi tubes and basic principles	1588
22.1.2 Volumetric flow calculations	1594
22.1.3 Mass flow calculations	1597
22.1.4 Square-root characterization	1600
22.1.5 Orifice plates	1609
22.1.6 Other differential producers	1625
22.1.7 Proper installation	1633
22.1.8 High-accuracy flow measurement	1638
22.1.9 Equation summary	1645
22.2 Laminar flowmeters	1648
22.3 Variable-area flowmeters	1649
22.3.1 Rotameters	1650
22.3.2 Weirs and flumes	1653
22.4 Velocity-based flowmeters	1661
22.4.1 Turbine flowmeters	1662
22.4.2 Vortex flowmeters	1671
22.4.3 Magnetic flowmeters	1675
22.4.4 Ultrasonic flowmeters	1687
22.4.5 Optical flowmeters	1695
22.5 Positive displacement flowmeters	1699
22.6 Standardized volumetric flow	1702
22.7 True mass flowmeters	1708
22.7.1 Coriolis flowmeters	1711
22.7.2 Thermal flowmeters	1726
22.8 Weighfeeders	1730
22.9 Change-of-quantity flow measurement	1733
22.10 Insertion flowmeters	1736
22.11 Process/instrument suitability	1742
22.12 Review of fundamental principles	1743
23 Continuous analytical measurement	1749
23.1 Conductivity measurement	1751
23.1.1 Dissociation and ionization in aqueous solutions	1752
23.1.2 Two-electrode conductivity probes	1753
23.1.3 Four-electrode conductivity probes	1756
23.1.4 Electrodeless conductivity probes	1759
23.2 pH measurement	1761
23.2.1 Colorimetric pH measurement	1761
23.2.2 Potentiometric pH measurement	1763
23.3 Chromatography	1784
23.3.1 Manual chromatography methods	1785
23.3.2 Automated chromatographs	1786
23.3.3 Species identification	1788
23.3.4 Chromatograph detectors	1789

23.3.5	Measuring species concentration	1791
23.3.6	Industrial applications of chromatographs	1794
23.3.7	Chromatograph sample valves	1797
23.3.8	Improving chromatograph analysis time	1800
23.4	Introduction to optical analyses	1805
23.5	Dispersive spectroscopy	1812
23.6	Non-dispersive Luft detector spectroscopy	1815
23.6.1	Single-beam analyzer	1818
23.6.2	Dual-beam analyzer	1819
23.6.3	Luft detectors	1822
23.6.4	Filter cells	1826
23.7	Gas Filter Correlation (GFC) spectroscopy	1829
23.8	Laser spectroscopy	1834
23.9	Fluorescence	1838
23.10	Chemiluminescence	1847
23.11	Analyzer sample systems	1851
23.12	Safety gas analyzers	1856
23.12.1	Oxygen gas	1860
23.12.2	Lower explosive limit (LEL)	1861
23.12.3	Hydrogen sulfide gas	1862
23.12.4	Carbon monoxide gas	1863
23.12.5	Chlorine gas	1864
23.13	Review of fundamental principles	1866
24	Machine vibration measurement	1871
24.1	Vibration physics	1871
24.1.1	Sinusoidal vibrations	1872
24.1.2	Non-sinusoidal vibrations	1877
24.2	Vibration sensors	1883
24.3	Monitoring hardware	1888
24.4	Mechanical vibration switches	1891
24.5	Review of fundamental principles	1892
25	Electric power measurement and control	1895
25.1	Introduction to power system automation	1896
25.2	Electrical power grids	1908
25.3	Interconnected generators	1912
25.4	Single-line electrical diagrams	1916
25.5	Circuit breakers and disconnects	1922
25.5.1	Low-voltage circuit breakers	1924
25.5.2	Medium-voltage circuit breakers	1927
25.5.3	High-voltage circuit breakers	1934
25.5.4	Reclosers	1942
25.6	Electrical sensors	1944
25.6.1	Potential transformers	1945
25.6.2	Current transformers	1949

25.6.3	Transformer polarity	1956
25.6.4	Instrument transformer safety	1963
25.6.5	Instrument transformer test switches	1965
25.6.6	Instrument transformer burden and accuracy	1973
25.7	Introduction to protective relaying	1985
25.8	ANSI/IEEE function number codes	1992
25.9	Instantaneous and time-overcurrent (50/51) protection	1996
25.10	Differential (87) current protection	2004
25.11	Directional overcurrent (67) protection	2023
25.12	Distance (21) protection	2025
25.12.1	Zone overreach and underreach	2026
25.12.2	Line impedance characteristics	2029
25.12.3	Using impedance diagrams to characterize faults	2032
25.12.4	Distance relay characteristics	2038
25.13	Auxiliary and lockout (86) relays	2049
25.14	Review of fundamental principles	2052
26	Signal characterization	2055
26.1	Flow measurement from differential pressure	2056
26.2	Flow measurement in open channels	2063
26.3	Material volume measurement	2065
26.4	Radiative temperature measurement	2073
26.5	Analytical measurements	2074
26.6	Review of fundamental principles	2077
27	Control valves	2079
27.1	Sliding-stem valves	2080
27.1.1	Globe valves	2081
27.1.2	Gate valves	2090
27.1.3	Diaphragm valves	2090
27.2	Rotary-stem valves	2092
27.2.1	Ball valves	2093
27.2.2	Butterfly valves	2093
27.2.3	Disk valves	2094
27.3	Dampers and louvres	2095
27.4	Valve packing	2098
27.5	Valve seat leakage	2107
27.6	Control valve actuators	2108
27.6.1	Pneumatic actuators	2109
27.6.2	Hydraulic actuators	2117
27.6.3	Self-operated valves	2118
27.6.4	Electric actuators	2122
27.6.5	Hand (manual) actuators	2126
27.7	Valve failure mode	2127
27.7.1	Direct/reverse actions	2128
27.7.2	Available failure modes	2130

27.7.3	Selecting the proper failure mode	2131
27.8	Actuator bench-set	2133
27.9	Pneumatic actuator response	2138
27.10	Valve positioners	2142
27.10.1	Force-balance pneumatic positioners	2147
27.10.2	Motion-balance pneumatic positioners	2151
27.10.3	Electronic positioners	2154
27.11	Split-ranging	2159
27.11.1	Complementary valve sequencing	2160
27.11.2	Exclusive valve sequencing	2163
27.11.3	Progressive valve sequencing	2165
27.11.4	Valve sequencing implementations	2168
27.12	Control valve sizing	2175
27.12.1	Physics of energy dissipation in a turbulent fluid stream	2176
27.12.2	Importance of proper valve sizing	2181
27.12.3	Gas valve sizing	2186
27.12.4	Relative flow capacity	2187
27.13	Control valve characterization	2188
27.13.1	Inherent versus installed characteristics	2189
27.13.2	Control valve performance with constant pressure	2191
27.13.3	Control valve performance with varying pressure	2194
27.13.4	Characterized valve trim	2197
27.14	Control valve problems	2203
27.14.1	Mechanical friction	2204
27.14.2	Flashing	2209
27.14.3	Cavitation	2213
27.14.4	Choked flow	2220
27.14.5	Valve noise	2222
27.14.6	Erosion	2224
27.14.7	Chemical attack	2229
27.15	Review of fundamental principles	2230
28	Variable-speed motor controls	2235
28.1	DC motor speed control	2237
28.2	AC motor speed control	2245
28.3	AC motor braking	2250
28.3.1	DC injection braking	2251
28.3.2	Dynamic braking	2252
28.3.3	Regenerative braking	2254
28.3.4	Plugging	2256
28.4	Motor drive features	2257
28.5	Use of line reactors	2258
28.6	Metering pumps	2262
28.7	Review of fundamental principles	2264

29 Closed-loop control	2265
29.1 Basic feedback control principles	2266
29.2 Diagnosing feedback control problems	2273
29.3 On/off control	2275
29.4 Proportional-only control	2277
29.5 Proportional-only offset	2284
29.6 Integral (reset) control	2290
29.7 Derivative (rate) control	2295
29.8 Summary of PID control terms	2297
29.8.1 Proportional control mode (P)	2297
29.8.2 Integral control mode (I)	2298
29.8.3 Derivative control mode (D)	2299
29.9 P, I, and D responses graphed	2299
29.9.1 Responses to a single step-change	2300
29.9.2 Responses to a momentary step-and-return	2301
29.9.3 Responses to two momentary steps-and-returns	2303
29.9.4 Responses to a ramp-and-hold	2304
29.9.5 Responses to an up-and-down ramp	2305
29.9.6 Responses to a multi-slope ramp	2306
29.9.7 Responses to a multiple ramps and steps	2307
29.9.8 Responses to a sine wavelet	2308
29.9.9 Note to students regarding quantitative graphing	2310
29.10 Different PID equations	2314
29.10.1 Parallel PID equation	2315
29.10.2 Ideal PID equation	2316
29.10.3 Series PID equation	2317
29.11 Pneumatic PID controllers	2318
29.11.1 Proportional control action	2319
29.11.2 Automatic and manual modes	2323
29.11.3 Derivative control action	2325
29.11.4 Integral control action	2326
29.11.5 Fisher MultiTrol	2329
29.11.6 Foxboro model 43AP	2332
29.11.7 Foxboro model 130	2334
29.11.8 External reset (integral) feedback	2337
29.12 Analog electronic PID controllers	2339
29.12.1 Proportional control action	2340
29.12.2 Derivative and integral control actions	2342
29.12.3 Full-PID circuit design	2346
29.12.4 Single-loop analog controllers	2349
29.12.5 Multi-loop analog control systems	2352
29.13 Digital PID controllers	2355
29.13.1 Stand-alone digital controllers	2356
29.13.2 Direct digital control (DDC)	2361
29.13.3 SCADA and telemetry systems	2368
29.13.4 Distributed Control Systems (DCS)	2373

29.13.5	Fieldbus control	2379
29.14	Practical PID controller features	2382
29.14.1	Manual and automatic modes	2383
29.14.2	Output and setpoint tracking	2384
29.14.3	Alarm capabilities	2386
29.14.4	Output and setpoint limiting	2386
29.14.5	Security	2387
29.15	Digital PID algorithms	2388
29.15.1	Introduction to pseudocode	2388
29.15.2	Position versus velocity algorithms	2394
29.16	Note to students	2399
29.16.1	Proportional-only control action	2400
29.16.2	Integral-only control action	2401
29.16.3	Proportional plus integral control action	2402
29.16.4	Proportional plus derivative control action	2403
29.16.5	Full PID control action	2404
29.17	Review of fundamental principles	2405
30	Process dynamics and PID controller tuning	2407
30.1	Process characteristics	2408
30.1.1	Self-regulating processes	2409
30.1.2	Integrating processes	2412
30.1.3	Runaway processes	2419
30.1.4	Steady-state process gain	2424
30.1.5	Lag time	2429
30.1.6	Multiple lags (orders)	2435
30.1.7	Dead time	2442
30.1.8	Hysteresis	2447
30.2	Before you tune	2451
30.2.1	Identifying operational needs	2452
30.2.2	Identifying process and system hazards	2454
30.2.3	Identifying the problem(s)	2455
30.2.4	Final precautions	2456
30.3	Quantitative PID tuning procedures	2457
30.3.1	Ziegler-Nichols closed-loop (“Ultimate Gain”)	2458
30.3.2	Ziegler-Nichols open-loop	2463
30.4	Heuristic PID tuning procedures	2466
30.4.1	Features of P, I, and D actions	2468
30.4.2	Tuning recommendations based on process dynamics	2469
30.4.3	Recognizing an over-tuned controller by phase shift	2470
30.4.4	Recognizing a “porpoising” controller	2474
30.5	Tuning techniques compared	2475
30.5.1	Tuning a “generic” process	2476
30.5.2	Tuning a liquid level process	2481
30.5.3	Tuning a temperature process	2485
30.6	Note to students	2489

30.6.1 Electrically simulating a process	2490
30.6.2 Building a “Desktop Process” unit	2491
30.6.3 Simulating a process by computer	2495
30.7 Review of fundamental principles	2496
31 Basic process control strategies	2499
31.1 Supervisory control	2500
31.2 Cascade control	2502
31.3 Ratio control	2510
31.4 Relation control	2518
31.5 Feedforward control	2520
31.5.1 Load Compensation	2521
31.5.2 Proportioning feedforward action	2532
31.6 Feedforward with dynamic compensation	2539
31.6.1 Dead time compensation	2540
31.6.2 Lag time compensation	2546
31.6.3 Lead/Lag and dead time function blocks	2552
31.7 Limit, Selector, and Override controls	2563
31.7.1 Limit controls	2566
31.7.2 Selector controls	2571
31.7.3 Override controls	2580
31.8 Techniques for analyzing control strategies	2585
31.8.1 Explicitly denoting controller actions	2585
31.8.2 Determining the design purpose of override controls	2593
31.9 Review of fundamental principles	2595
32 Process safety and instrumentation	2597
32.1 Classified areas and electrical safety measures	2598
32.1.1 Classified area taxonomy	2599
32.1.2 Explosive limits	2601
32.1.3 Protective measures	2604
32.2 Concepts of probability	2610
32.2.1 Mathematical probability	2612
32.2.2 Laws of probability	2614
32.2.3 Applying probability laws to real systems	2625
32.3 Practical measures of reliability	2628
32.3.1 Failure rate and MTBF	2628
32.3.2 The “bathtub” curve	2634
32.3.3 Reliability	2636
32.3.4 Probability of failure on demand (PFD)	2639
32.4 High-reliability systems	2640
32.4.1 Design and selection for reliability	2641
32.4.2 Preventive maintenance	2642
32.4.3 Component de-rating	2644
32.4.4 Redundant components	2645
32.4.5 Proof tests and self-diagnostics	2650

32.5	Overpressure protection devices	2655
32.5.1	Rupture disks	2657
32.5.2	Direct-actuated safety and relief valves	2658
32.5.3	Pilot-operated safety and relief valves	2668
32.6	Safety Instrumented Functions and Systems	2669
32.6.1	SIS sensors	2674
32.6.2	SIS controllers (logic solvers)	2680
32.6.3	SIS final control elements	2682
32.6.4	Safety Integrity Levels	2685
32.6.5	SIS example: burner management systems	2687
32.6.6	SIS example: water treatment oxygen purge system	2694
32.6.7	SIS example: nuclear reactor scram controls	2698
32.7	Review of fundamental principles	2702
33	Instrumentation cyber-security	2705
33.1	Stuxnet	2706
33.1.1	A primer on uranium enrichment	2707
33.1.2	Gas centrifuge vulnerabilities	2713
33.1.3	The Natanz uranium enrichment facility	2715
33.1.4	Stuxnet version 0.5	2716
33.1.5	Stuxnet version 1.x	2717
33.2	Motives	2717
33.2.1	Technical challenge	2718
33.2.2	Profit	2718
33.2.3	Espionage	2719
33.2.4	Sabotage	2720
33.2.5	Terrorism	2721
33.3	Lexicon of cyber-security terms	2722
33.4	Fortifying strategies	2725
33.4.1	Policy-based strategies	2725
33.4.2	Design-based strategies	2733
33.5	Review of fundamental principles	2739
34	Problem-solving and diagnostic strategies	2743
34.1	Learn principles, not procedures	2744
34.2	Active reading	2744
34.2.1	Don't limit yourself to one text!	2745
34.2.2	Marking versus outlining a text	2746
34.3	General problem-solving techniques	2751
34.3.1	Identifying and classifying all "known" conditions	2751
34.3.2	Re-cast the problem in a different format	2751
34.3.3	Working backwards from a known solution	2752
34.3.4	Using thought experiments	2755
34.3.5	Explicitly annotating your thoughts	2756
34.4	Mathematical problem-solving techniques	2760
34.4.1	Manipulating algebraic equations	2761

34.4.2	Linking formulae to solve mathematical problems	2764
34.4.3	Double-checking calculations	2778
34.4.4	Maintaining relevance of intermediate calculations	2783
34.5	Problem-solving by simplification	2785
34.5.1	Limiting cases	2786
34.6	Scientific system diagnosis	2791
34.6.1	Scientific method	2792
34.6.2	Occam’s Razor	2792
34.6.3	Diagnosing intermittent problems	2793
34.6.4	Strategy: tracing data paths	2794
34.7	Common diagnostic mistakes	2796
34.7.1	Failing to gather data	2797
34.7.2	Failing to use relevant documentation	2798
34.7.3	Fixating on the first hypothesis	2799
34.7.4	Failing to build and test a new system in stages	2801
34.8	Helpful “tricks” using a digital multimeter (DMM)	2803
34.8.1	Recording unattended measurements	2803
34.8.2	Avoiding “phantom” voltage readings	2804
34.8.3	Non-contact AC voltage detection	2807
34.8.4	Detecting AC power harmonics	2808
34.8.5	Identifying noise in DC signal paths	2809
34.8.6	Generating test voltages	2810
34.8.7	Using the meter as a temporary jumper	2811
A	Flip-book animations	2813
A.1	Polyphase light bulbs animated	2813
A.2	Polyphase induction motor animated	2840
A.3	Rotating phasor animated	2865
A.4	Differentiation and integration animated	2906
A.5	Guided-wave radar level measurement	3043
A.6	Basic chromatograph operation	3096
B	<i>Doctor Strangeflow</i>, or how I learned to relax and love Reynolds numbers	3157
C	Disassembly of a sliding-stem control valve	3167
D	How to use this book – some advice for teachers	3177
D.1	Teaching technical theory	3178
D.1.1	The problem with lecture	3180
D.1.2	A more accurate model of learning	3182
D.1.3	The ultimate goal of education	3183
D.2	Teaching technical practices (labwork)	3186
D.3	Teaching diagnostic principles and practices	3194
D.3.1	Deductive diagnostic exercises	3195
D.3.2	Inductive diagnostic exercises	3203
D.4	Practical topic coverage	3210

D.5 Principles, not procedures	3211
D.6 Assessing student learning	3213
D.7 Summary	3215
E Contributors	3217
E.1 Error corrections	3218
E.2 New content	3222
F Creative Commons Attribution License	3223
F.1 A simple explanation of your rights	3224
F.2 Legal code	3225

Preface

I did not want to write this book . . . honestly.

My first book project began in 1998, titled *Lessons In Electric Circuits*, and I didn't call "quit" until six volumes and five years later. Even then it was not complete, but being an open-source project it gained traction on the Internet to the point where other people took over its development and it grew fine without me. The impetus for writing this first tome was a general dissatisfaction with available electronics textbooks. Plenty of textbooks exist to describe things, but few really *explain* things well for students, and the field of electronics is no exception. I wanted my book(s) to be different, and so they were. No one told me how time-consuming it was going to be to write them, though!

The next few years' worth of my spare time went to developing a set of question-and-answer worksheets designed to teach electronics theory in a Socratic, active-engagement style. This project proved quite successful in my professional life as an instructor of electronics. In the summer of 2006, my job changed from teaching electronics to teaching industrial instrumentation, and I decided to continue the Socratic mode of instruction with another set of question-and-answer worksheets.

However, the field of industrial instrumentation is not as well-represented as general electronics, and thus the array of available textbooks is not as vast. I began to re-discover the drudgery of trying to teach with inadequate texts as source material. The basis of my active teaching style was that students would spend time researching the material on their own, then engage in Socratic-style discussion with me on the subject matter when they arrived for class. This teaching technique functions in direct proportion to the quality and quantity of the research sources at the students' disposal. Despite much searching, I was unable to find a textbook adequately addressing my students' learning needs. Many textbooks I found were written in a shallow, "math-phobic" style well below the level I intended to teach to. Some reference books I found contained great information, but were often written for degreed engineers with lots of Laplace transforms and other mathematical techniques well above the level I intended to teach to. Few on either side of the spectrum actually made an effort to explain certain concepts students generally struggle to understand. I needed a text giving good, practical information and theoretical coverage at the same time.

In a futile effort to provide my students with enough information to study outside of class, I scoured the Internet for free tutorials written by others. While some manufacturer's tutorials were nearly perfect for my needs, others were just as shallow as the textbooks I had found, and/or were little more than sales brochures. I found myself starting to write my own tutorials on specific topics to "plug the gaps," but then another problem arose: it became troublesome for students to navigate through dozens of tutorials in an effort to find the information they needed in their studies. What my students really needed was a *book*, not a smorgasbord of tutorials.

So here I am again, writing another textbook. This time around I have the advantage of wisdom gained from the first textbook project. For this project, I will *not*:

- . . . attempt to maintain a parallel book in HTML markup (for direct viewing on the Internet). I had to go to the trouble of inventing my own quasi-XML markup language last time in an effort to generate multiple format versions of the book from the same source code. Instead, this time I will use stock L^AT_EX as the source code format and regular Adobe PDF format for the final output, which anyone may read thanks to its ubiquity. If anyone else desires the book in a different format, I will gladly let them deal with issues of source code translation. Not that this should be a terrible problem for anyone technically competent in markup languages, as L^AT_EX source is rather easy to work with.
- . . . use a GNU GPL-style copyleft license. Instead, I will use the Creative Commons Attribution-only license, which is far more permissive for anyone wishing to incorporate my work into derivative works. My interest is maximum flexibility for those who may adapt my material to their own needs, not the imposition of certain philosophical ideals.
- . . . start from a conceptual state of “ground zero.” I will assume the reader has certain familiarity with electronics and mathematics, which I will build on. If a reader finds they need to learn more about electronics, they should go read *Lessons In Electric Circuits*.
- . . . avoid using calculus to help explain certain concepts. Not all my readers will understand these parts, and so I will be sure to explain what I can without using calculus. However, I want to give my more mathematically adept students an opportunity to see the power of calculus applied to instrumentation where appropriate. By occasionally applying calculus and explaining my steps, I also hope this text will serve as a practical guide for students who might wish to learn calculus, so they can see its utility and function in a context that interests them.

There do exist many fine references on the subject of industrial instrumentation. I only wish I could condense their best parts into a single volume for my students. Being able to do so would certainly save me from having to write my own! Listed here are some of the best books I can recommend for those wishing to explore instrumentation outside of my own presentation:

- *Instrument Engineers’ Handbook* series (Volumes I, II, and III), edited by Béla Lipták – by far my favorite modern references on the subject. Unfortunately, there is a fair amount of material within that lies well beyond my students’ grasp (Laplace transforms, etc.), and the volumes are incredibly bulky and expensive (nearly 2000 pages, and at a cost of nearly \$200.00, *apiece!*). These texts also lack some of the basic content my students do need, and I don’t have the heart to tell them to buy yet *another* textbook to fill the gaps.
- *Handbook of Instrumentation and Controls*, by Howard P. Kallen. Perhaps the best-written textbook on general instrumentation I have ever encountered. Too bad it is both long out of print – my copy dates 1961 – and technologically dated. Like most American textbooks written during the years immediately following Sputnik, it is a masterpiece of practical content and conceptual clarity. I consider books like this useful for their presentations of “first principles,” which of course are timeless.
- *Industrial Instrumentation Fundamentals*, by Austin E. Fribance. Another great post-Sputnik textbook – my copy dates 1962.

- *Instrumentation for Process Measurement and Control*, by Norman A. Anderson. An inspiring effort by someone who knows the art of teaching as well as the craft of instrumentation. Too bad the content doesn't seem to have been updated since 1980.
- *Applied Instrumentation in the Process Industries* (Volume I), edited by William G. Andrew. A very clear and fairly comprehensive overview of industrial instrumentation. Sadly, this fine book is out of print, and much of the material is dated (second edition written in 1979).
- Practically anything written by Francis Greg Shinsky.

Whether or not I achieve my goal of writing a better textbook is a judgment left for others to make. One decided advantage my book will have over all the others is its *openness*. If you don't like anything you see in these pages, you have the right to modify it to your liking! Delete content, add content, modify content – it's all fair game thanks to the Creative Commons licensing. My only conditions are declared in the license: you must give me credit for my original authorship, include a copyright notice, and also include the full text of the Creative Commons license informing readers of their rights (see Section 3(a)(1) of the License for more detail). What you do with it beyond that is wholly up to you². This way, perhaps I can spare someone else from having to write their own textbook from scratch!

²This includes selling copies of it, either electronic or print. Of course, you must include the Creative Commons license as part of the text you sell, which means anyone will be able to tell it is an open text and can probably figure out how to download an electronic copy off the Internet for free. The only way you're going to make significant money selling this text is to add your own value to it, either in the form of expansions or bundled product (e.g. simulation software, learning exercises, etc.), which of course is perfectly fair – you must profit from your *own* labors. All my work does for you is give you a starting point.

Chapter 1

Calculus

Mathematics is the investigation of an artificial world: a universe populated by abstract entities and rigid rules governing those entities. Mathematicians devoted to the study and advancement of pure mathematics have an extremely well-developed respect for these rules, for the integrity of this artificial world depends on them. In order to preserve the integrity of their artificial world, their collective work must be *rigorous*, never allowing for sloppy handling of the rules or allowing intuitive leaps to be left unproven.

However, many of the tools and techniques developed by mathematicians for their artificial world happen to be extremely useful for understanding the real world in which we live and work, and therein lies a problem. In applying mathematical rules to the study of real-world phenomena, we often take a far more pragmatic approach than any mathematician would feel comfortable with.

The tension between pure mathematicians and those who apply math to real-world problems is not unlike the tension between linguists and those who use language in everyday life. All human languages have rules (though none as rigid as in mathematics!), and linguists are the guardians of those rules, but the vast majority of human beings play fast and loose with the rules as they use language to describe and understand the world around them. Whether or not this “sloppy” adherence to rules is good depends on which camp you are in. To the purist, it is offensive; to the pragmatist, it is convenient.

I like to tell my students that mathematics is very much like a language. The more you understand mathematics, the larger “vocabulary” you will possess to describe principles and phenomena you encounter in the world around you. Proficiency in mathematics also empowers you to grasp relationships between different things, which is a powerful tool in learning new concepts.

This book is not written for (or by!) mathematicians. Rather, it is written for people wishing to make sense of industrial process measurement and control. This chapter of the book is devoted to a very pragmatic coverage of certain mathematical concepts, for the express purpose of applying these concepts to real-world systems.

Mathematicians, cover your eyes for the rest of this chapter!

1.1 Introduction to calculus

Few areas of mathematics are as powerfully useful in describing and analyzing the physical world as calculus: the mathematical study of *changes*. Calculus also happens to be tremendously confusing to most students first encountering it. A great deal of this confusion stems from mathematicians' insistence on rigor¹ and denial of intuition.

Look around you right now. Do you see any mathematicians? If not, good – you can proceed in safety. If so, find another location to begin reading the rest of this chapter. I will frequently appeal to practical example and intuition in describing the basic principles of single-variable calculus, for the purpose of expanding your mathematical “vocabulary” to be able to describe and better understand phenomena of change related to instrumentation.

Silvanus P. Thompson, in his wonderful book *Calculus Made Simple* originally published in 1910, began his text with a short chapter entitled, “To Deliver You From The Preliminary Terrors².” I will follow his lead by similarly introducing you to some of the notations frequently used in calculus, along with very simple (though not mathematically rigorous) definitions.

When we wish to speak of a change in some variable's value (let's say x), it is common to precede the variable with the capital Greek letter “delta” as such:


$$\Delta x = \text{“Change in } x\text{”}$$

An alternative interpretation of the “delta” symbol (Δ) is to think of it as denoting a *difference* between two values of the same variable. Thus, Δx could be taken to mean “the difference between two values of x ”. The cause of this difference is not important right now: it may be the difference between the value of x at one point in time versus another point in time, it may be the difference between the value of x at one point in space versus another point in space, or it may simply be the difference between values of x as it relates to some other variable (e.g. y) in a mathematical function. If we have some variable such as x that is known to change value relative to some other variable (e.g. time, space, y), it is nice to be able to express that change using precise mathematical symbols, and this is what the “delta” symbol does for us.

¹In mathematics, the term *rigor* refers to a meticulous attention to detail and insistence that each and every step within a chain of mathematical reasoning be thoroughly justified by deductive logic, not intuition or analogy.

²The book's subtitle happens to be, *Being a very-simplest introduction to those beautiful methods of reckoning which are generally called by the terrifying names of the differential calculus and the integral calculus*. Not only did Thompson recognize the anti-pragmatic tone with which calculus is too often taught, but he also infused no small amount of humor in his work.

For example, if the temperature of a furnace (T) increases over time, we might wish to describe that change in temperature as ΔT :




$$\Delta T = T_{10:32} - T_{9:45}$$

$$\Delta T = 1276\text{ °F} - 1255\text{ °F} = 21\text{ °F}$$

The value of ΔT is nothing more than the difference (subtraction) between the recent temperature and the older temperature. A rising temperature over time thus yields a positive value for ΔT , while a falling temperature over time yields a negative value for ΔT .

We could also describe differences between the temperature of two *locations* (rather than a difference of temperature between two *times*) by the notation ΔT , such as this example of heat transfer through a heat-conducting wall where one side of the wall is hotter than the other:



Once again, ΔT is calculated by subtracting one temperature from another. Here, the sign (positive or negative) of ΔT denotes the *direction* of heat flow through the thickness of the wall.

One of the major concerns of calculus is changes or differences between variable values lying *very*

close to each other. In the context of a heating furnace, this could mean increases in temperature over miniscule time periods. In the context of heat flowing through a wall, this could mean differences in temperature sampled between points within the wall immediately adjacent each other. If our desire is to express the change in a variable between neighboring points along a continuum rather than over some discrete period, we may use a different notation than the capital Greek letter delta (Δ); instead, we use a lower-case Roman letter d (or in some cases, the lower-case Greek letter delta: δ).

Thus, a change in furnace temperature from one instant in time to the next could be expressed as dT (or δT), and likewise a difference in temperature between two adjacent positions within the heat-conducting wall could also be expressed as dT (or δT). Just as with the “delta” (Δ) symbol, the changes references by either d or δ may occur over a variety of different domains.

We even have a unique name for this concept of extremely small differences: whereas ΔT is called a *difference* in temperature, dT is called a *differential* of temperature. The concept of a differential may seem redundant to you right now, but they are actually quite powerful for describing *continuous changes*, especially when one differential is related to another differential by ratio (something we call a *derivative*).

Another major concern in calculus is how quantities accumulate, especially how differential quantities add up to form a larger whole. A furnace’s temperature rise since start-up (ΔT_{total}), for example, could be expressed as the accumulation (sum) of many temperature differences (ΔT) measured periodically. The total furnace temperature rise calculated from a sampling of temperature once every minute from 9:45 to 10:32 AM could be written as:

$$\Delta T_{total} = \Delta T_{9:45} + \Delta T_{9:46} + \cdots + \Delta T_{10:32} = \text{Total temperature rise over time, from 9:45 to 10:32}$$

A more sophisticated expression of this series uses the capital Greek letter sigma (meaning “sum of” in mathematics) with notations specifying which temperature differences to sum:

$$\Delta T_{total} = \sum_{n=9:45}^{10:32} \Delta T_n = \text{Total temperature rise over time, from 9:45 to 10:32}$$

However, if our furnace temperature monitor scans at an infinite pace, measuring temperature *differentials* (dT) and summing them in rapid succession, we may express the same accumulated temperature rise as an *infinite* sum of *infinitesimal* (infinitely small) changes, rather than as a finite sum of temperature changes measured once every minute. Just as we introduced a unique mathematical symbol to represent differentials (d) over a continuum instead of differences (Δ) over discrete periods, we will introduce a unique mathematical symbol to represent the summation of differentials (\int) instead of the summation of differences (\sum):


$$\Delta T_{total} = \int_{9:45}^{10:32} dT = \text{Total temperature rise over time, from 9:45 to 10:32}$$

This summation of infinitesimal quantities is called *integration*, and the elongated “S” symbol (\int) is the *integral* symbol.

These are the two major ideas and notations of calculus: *differentials* (tiny changes represented by d or δ) and *integrals* (accumulations represented by \int). Now that wasn’t so frightening, was it?

1.2 The concept of differentiation

Suppose we wished to measure the rate of propane gas flow through a hose to a torch:



Flowmeters appropriate for measuring low flow rates of any gas are typically very expensive, making it impractical to directly measure the flow rate of propane fuel gas consumed by this torch at any given moment. We could, however, *indirectly* measure the flow rate of propane by placing the tank on a scale where its mass (m) could be monitored over time. By taking measurements of mass over short time periods (Δt), we could calculate the corresponding differences in mass (Δm), then calculate the ratio of mass lost over time to calculate average mass flow rate (\overline{W}):

$$\overline{W} = \frac{\Delta m}{\Delta t} = \text{Average mass flow rate}$$

Where,

\overline{W} = Average mass flow rate within each time period (kilograms per minute)


Δm = Measured mass difference over time period (kilograms)

Δt = Time period of mass measurement sample (minutes)

Note that flow rate is a ratio (quotient) of mass change over time change. The units used to express flow even reflect this process of division: kilograms *per* minute.

$$\overline{W} = \frac{[\text{kg}]}{[\text{min}]} = \text{Average mass flow rate} = \left[\frac{\text{kg}}{\text{min}} \right]$$

Graphed as a function over time, the tank's mass will be seen to decrease as time elapses. Each dot represents a mass and time measurement coordinate pair (e.g. 20 kilograms at 7:38, 18.6 kilograms at 7:51, etc.):




We should recall from basic geometry that the slope of a line or line segment is defined as its *rise* (vertical height) divided by its *run* (horizontal width). Thus, the average mass flow rate calculated within each time period may be represented as the pitch (slope) of the line segments connecting dots, since mass flow rate is defined as a change in mass per (divided by) change in time.

Periods of high propane flow (large flame from the torch) show up on the graph as steeply-pitched line segments. Periods of no propane flow reveal themselves as flat portions on the graph (no rise or fall over time).

If the determination of average flow rates between significant gaps in time is good enough for our application, we need not do anything more. However, if we wish to detect mass flow rate at any particular *instant* in time, we need to perform the same measurements of mass loss, time elapse, and division of the two at an infinitely fast rate.

Supposing such a thing were possible, what we would end up with is a smooth graph showing mass consumed over time. Instead of a few line segments roughly approximating a curve, we would have an *infinite* number of infinitely short line segments connected together to form a seamless curve. The flow rate at any particular point in time would be the ratio of the mass and time differentials (the slope of the infinitesimal line segment) at that point:



$$W = \frac{dm}{dt} = \text{Instantaneous mass flow rate}$$

Where,

W = Instantaneous mass flow rate at a given time (kilograms per minute)

dm = Mass differential at a single point in time (kilograms)

dt = Time differential at a single point in time (minutes)

Flow is calculated just the same as before: a quotient of mass and time differences, except here the differences are infinitesimal in magnitude. The unit of flow measurement reflects this process of division, just as before, with mass flow rate expressed in units of kilograms *per* minute. Also, just as before, the rate of flow is graphically represented by the *slope* of the graph: steeply-sloped points on the graph represent moments of high flow rate, while shallow-sloped points on the graph represent moments of low flow rate.

Such a ratio of differential quantities is called a *derivative* in calculus³. Derivatives – especially time-based derivatives such as flow rate – find many applications in instrumentation as well as the general sciences. Some of the most common time-based derivative functions include the relationships between *position* (x), *velocity* (v), and *acceleration* (a).

Velocity (v) is the rate at which an object changes position over time. Since position is typically denoted by the variable x and time by the variable t , the derivative of position with respect to time may be written as such:

$$v = \frac{dx}{dt} \quad [\text{meters/second}] = \frac{[\text{meters}]}{[\text{seconds}]}$$

The metric units of measurement⁴ for velocity (meters per second, miles per hour, etc.) betray this process of division: a differential of position (meters) divided by a differential of time (second).

Acceleration (a) is the rate at which an object changes velocity over time. Thus, we may express acceleration as the time-derivative of velocity, just as velocity was expressed as the time-derivative of position:

$$a = \frac{dv}{dt} \quad [\text{meters/second}^2] = \frac{[\text{meters/second}]}{[\text{seconds}]}$$

We may even express acceleration as a function of position (x), since it is the rate of change of the rate of change in position over time. This is known as a *second derivative*, since it is applying the process of “differentiation” twice:


$$a = \frac{d}{dt} \left(\frac{dx}{dt} \right) = \frac{d^2x}{dt^2} \quad [\text{meters/second}^2] = \frac{[\text{meters}]}{[\text{seconds}^2]}$$

As with velocity, the units of measurement for acceleration (meters per second squared, or alternatively meters per second per second) suggest a compounded quotient.

³Isaac Newton referred to derivatives as *fluxions*, and in Silvanus Thompson’s day they were known as *differential coefficients*.

⁴British units of measurement for velocity indicate this same process of division: the number of feet traveled in a time period of seconds yields a velocity in feet per second. There is nothing unique about metric units in this regard.

It is also possible to express rates of change between different variables not involving time. A common example in the engineering realm is the concept of *gain*, generally defined as the ratio of output change to input change. An electronic amplifier, for example, with an input signal of 2 volts (peak-to-peak) and an output signal of 8.6 volts (peak-to-peak), would be said to have a gain of 4.3, since the change in output measured in peak-to-peak volts is 4.3 times larger than the corresponding change in input voltage:




This gain may be expressed as a quotient of differences ($\frac{\Delta V_{out}}{\Delta V_{in}}$), or it may be expressed as a derivative instead:

$$\text{Gain} = \frac{dV_{out}}{dV_{in}}$$

If the amplifier's behavior is perfectly linear, there will be no difference between gain calculated using differences and gain calculated using differentials (the derivative), since the average slope of a straight line is the same as the instantaneous slope at any point along that line. If, however, the amplifier does not behave in a perfectly linear fashion, gain calculated from large changes in voltage ($\frac{\Delta V_{out}}{\Delta V_{in}}$) will not be the same as gain calculated from infinitesimal changes at different points along the amplifier's operating voltage range.

1.3 The concept of integration

Suppose we wished to measure the consumption of propane over time for a large propane storage tank supplying a building with heating fuel, because the tank lacked a level indicator to show how much fuel was left at any given time. The flow rate is sufficiently large, and the task sufficiently important, to justify the installation of a mass flowmeter⁵, which registers flow rate at an indicator inside the building:



By measuring true mass flow rate, it should be possible to indirectly measure how much propane has been used at any time following the most recent filling of the tank. For example, if the mass flow rate of propane out of the tank happened to be a constant 5 kilograms per hour for 30 hours straight, it would be a simple matter of multiplication to calculate the consumed mass:

$$\left(\frac{5 \text{ kg}}{\text{hr}}\right) \left(\frac{30 \text{ hrs}}{1}\right) = 150 \text{ kg of propane consumed}$$

Expressing this mathematically as a function of differences in mass and differences in time, we may write the following equation:

$$\Delta m = \bar{W} \Delta t$$

Where,

\bar{W} = Average mass flow rate within the time period (kilograms per hour)

Δm = Mass difference over time period (kilograms)

Δt = Time period of flow measurement sample (hours)

It is easy to see how this is just a variation of the quotient-of-differences equation used previously in this chapter to define mass flow rate:


$$\bar{W} = \frac{\Delta m}{\Delta t} = \text{Average mass flow rate}$$

Inferring mass flow rate from changes in mass over time periods is a process of *division*. Inferring changes in mass from flow rate over time periods is a process of *multiplication*. The units of measurement used to express each of the variables makes this quite clear.

⁵Most likely a thermal mass flowmeter or a Coriolis flowmeter.


As we learned previously, the process of differentiation is really just a matter of determining the *slope* of a graph. A graph of propane fuel mass (m) plotted over discrete points in time (t) has a slope corresponding to mass flow rate ($W = \frac{\Delta m}{\Delta t}$). Here, we are attempting to do the opposite: the data reported by the sensing instrument is propane mass flow rate (W), and our goal is to determine total mass lost (Δm) as the propane is consumed from the storage tank over a period of time (Δt). This operation is fundamentally distinct from differentiation, which means its graphical interpretation will not be the same. Instead of calculating the slope of the graph, we will have to do something else.

Using the previous example of the propane flowmeter sensing a constant mass flow rate (W) of 5 kilograms of propane per hour for 30 hours (for a total consumption of 150 kilograms), we may plot a trend graph showing flow rate (vertical) as a function of time (horizontal). We know the consumed propane quantity is the simple product (multiplication) of constant flow rate and time, which relates to the geometric *area* enclosed by the graph, since the area of any rectangle is height times width:




To summarize: the height of this graph represents the rate at which propane exits the storage tank, the width of the graph represents the length of time propane has been consumed from the storage tank, and the geometric area enclosed by these two boundaries represents the total mass of propane consumed during that time.

The task of inferring lost mass over time becomes more complicated if the propane flow rate is not constant over time. Consider the following graph, showing periods of increased and decreased flow rate due to different gas-fired appliances turning on and off inside the building:



Here, the propane gas flow rate does not stay constant throughout the entire time interval covered by the graph. This graph is obviously more challenging to analyze than the previous example where the propane flow rate was constant. From that previous example, though, we have learned that the geometric area enclosed by the boundaries of the graph's height (flow rate) and width (time duration) has physical meaning, representing the total quantity of propane passed through the flowmeter. Despite the fact that the graph's area is now more complex to calculate, the basic principle remains the same as before: the enclosed area represents the amount of propane consumed.

In order to accurately calculate the amount of propane mass consumed by the building over time, we must treat each period of constant flow as its own propane quantity, calculating the mass lost during each period, then summing those mass differences to arrive at a total mass loss for the entire time interval covered by the graph. Since we know the difference (loss) in mass over a time period is equal to the average flow rate for that period multiplied by the period's duration ($\Delta m = W \Delta t$), we may calculate each period's mass as an *area* underneath the graph line, each rectangular area being equal to height (W) times width (Δt):



Each rectangular area underneath the flow line on the graph ($W_n \Delta t_n$) represents a quantity of propane gas consumed during that time period. To find the total amount of propane consumed in the time represented by the entire graph, we must sum these mass quantities together:


$$\Delta m = (W_1 \Delta t_1) + (W_2 \Delta t_2) + (W_3 \Delta t_3) + (W_4 \Delta t_4) + (W_5 \Delta t_5) + (W_6 \Delta t_6) + (W_7 \Delta t_7) + (W_8 \Delta t_8)$$

A “shorthand” notation for this sum uses the capital Greek letter sigma to represent a series of repeated products (multiplication) of mass flow and time periods for the eight rectangular areas enclosed by the graph:

$$\Delta m = \sum_{n=1}^8 W_n \Delta t_n$$


While $W_n \Delta t_n$ represents the area of just one of the rectangular periods, $\sum_{n=1}^8 W_n \Delta t_n$ represents the total combined areas, which in this example represents the total mass of propane consumed over the eight time periods shown on the graph.

The task of inferring total propane mass consumed over time becomes even more complicated if the flow does not vary in stair-step fashion as it did in the previous example. Suppose the building were equipped with *throttling* gas appliances instead of on/off gas appliances, thus creating a continuously variable flow rate demand over time. A typical flow rate graph might look something like this:

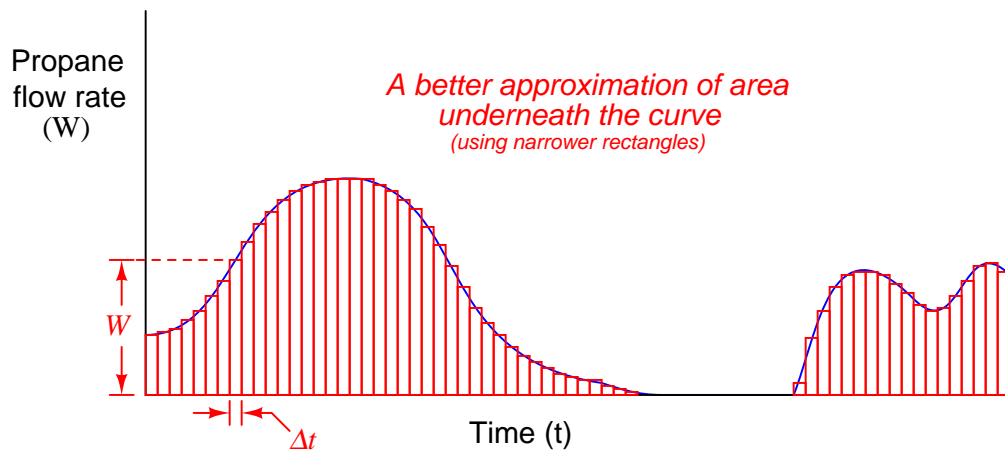


The physics of gas flow and gas mass over time has not changed: total propane mass consumed over time will still be the area enclosed beneath the flow curve. The only difference between this example and the two previous examples is the complexity of actually calculating that enclosed area.


We can, however, *approximate* the area underneath this curve by overlaying a series of rectangles, the area of each rectangle being height (W) times width (Δt):



It should be intuitively obvious that this strategy of approximating the area underneath a curve improves with the number of rectangles used. Each rectangle still has an area $W\Delta t$, but since the Δt periods are shorter, it is easier to fit the rectangles to the curve of the graph. The summation of a series of rectangular areas intended to approximate the area enclosed by a graphed function is commonly referred to as a *Riemann Sum* in honor of the mathematician Bernhard Riemann:



Taking this idea to its ultimate realization, we could imagine a super-computer sampling mass flow rates at an infinite speed, then calculating the rectangular area covered by each flow rate (W) times each infinitesimal increment of time (dt). With time increments of negligible width, the “approximation” of area underneath the graph found by the sum of all these rectangles would be perfect – indeed, it would not be an approximation at all, but rather an exact match:



If we represent infinitesimal time increments using the notation “ dt ” as opposed to the notation “ Δt ” used to represent discrete time periods, we must also use different notation to represent the mathematical sum of those quantities. Thus, we will replace the “sigma” symbol (\sum) used for summation and replace it with the integral symbol (\int), which means a *continuous* summation of infinitesimal quantities:

$$\Delta m = \sum_{n=0}^x W \Delta t_n \quad \text{Summing discrete quantities of } W \Delta t$$

$$\Delta m = \int_0^x W dt \quad \text{Summing continuous quantities of } W dt$$

This last equation tells us the total change in mass (Δm) from time 0 to time x is equal to the continuous sum of mass quantities found by multiplying mass flow rate measurements (W) over corresponding increments of time (dt). We refer to this summation of infinitesimal quantities as *integration* in calculus. Graphically, the *integral* of a function is the geometric area enclosed by the function over a specified interval.

An important detail to note is that this process of integration (multiplying flow rates by infinitesimal time increments, then summing those products) only tells us how much propane mass was consumed – it does *not* tell us how much propane remains in the tank, which was the purpose of installing the mass flowmeter and performing all this math! The integral of mass flow and time ($\int W dt$) will always be a negative⁶ quantity in this example, because a flow of propane gas out of the tank represents a *loss* of propane mass within the tank. In order to calculate the amount of propane mass left in the tank, we would need to know the initial value of propane in the tank before any of it flowed to the building, then we would add this initial mass quantity (m_0) to the negative mass loss calculated by integration.

Thus, we would mathematically express the propane mass inside the tank at time x as such⁷:

$$m_x = \int_0^x W dt + m_0$$


This initial value must always be considered in problems of integration if we attempt to absolutely define some integral quantity. Otherwise, all the integral will yield is a relative quantity (how much something has *changed* over an interval).

⁶Although we will measure time, and differentials of time, as positive quantities, the mass flowmeter should be configured to show a negative flow rate (W) when propane flows from the tank to the building. This way, the *integrand* (the product “inside” the integration symbol; $W dt$) will be a negative quantity, and thus the integral over a positive time interval (from 0 to x) will likewise be a negative quantity.

⁷According to calculus convention, the differential dt represents the end of the integrand. It is safe to regard the long “S” symbol and the differential (dx , dt , etc.) as complementary *grouping symbols* declaring the beginning and end of the integrand. This tells us m_0 is not part of the integrand, but rather comes after it. Using parentheses to explicitly declare the boundaries of the integrand, we may re-write the expression as $m_x = (\int_0^x W dt) + m_0$

The problem of initial values is very easy to relate to common experience. Consider the *odometer* indication in an automobile. This is an example of an integral function, the distance traveled (x) being the time-integral⁸ of speed (or velocity, v):

$$\Delta x = \int v dt$$



Although the odometer does accumulate to larger and larger values as you drive the automobile, its indication does not necessarily tell me how many miles *you* have driven it. If, for example, you purchased the automobile with 32411.6 miles on the odometer, its current indication of 52704.8 miles means that you have driven it 20293.2 miles. The automobile's *total* distance traveled since manufacture is equal to the distance you have accumulated while driving it ($\int v dt$) *plus* the initial mileage accumulated at the time you took ownership of it (x_0):

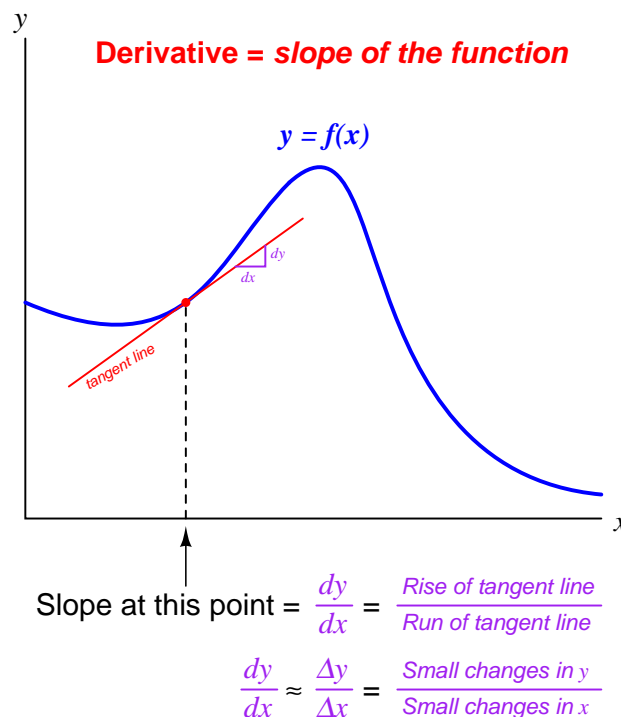
$$x_{total} = \int v dt + x_0$$

⁸Recall from the previous section (“The Concept of Differentiation”) that velocity could be defined as the time-derivative of position: $v = \frac{dx}{dt}$. All we have done here is algebraically solved for changes in x by first multiplying both sides of the equation by dt to arrive at $dx = v dt$. Next, we integrate both sides of the equation in order to “un-do” the differential (d) applied to x : $\int dx = \int v dt$. Since accumulations (\int) of any differential (dx) yields a discrete change for that variable, we may substitute Δx for $\int dx$ and get our final answer of $\Delta x = \int v dt$.

1.4 How derivatives and integrals relate to one another


First, let us review some of the properties of *differentials* and *derivatives*, referencing the expression and graph shown below:

- A *differential* is an infinitesimal increment of change (difference) in some continuously-changing variable, represented either by a lower-case Roman letter d or a lower-case Greek letter “delta” (δ). Such a change in time would be represented as dt ; a similar change in temperature as dT ; a similar change in the variable x as dx .
- A derivative is always a *quotient of differences*: a process of subtraction (to calculate the amount each variable changed) followed by division (to calculate the *rate* of one change to another change).
- The units of measurement for a derivative reflect this final process of division: one unit divided by some other unit (e.g. gallons *per* minute, feet *per* second).
- Geometrically, the derivative of a function is its graphical *slope* (its “rise over run”).
- When computing the value of a derivative, we must specify a single point along the function where the slope is to be calculated.
- The *tangent line* matching the slope at that point has a “rise over run” value equal to the derivative of the function at that point.



Next, let us review some of the properties of *integrals*, referencing the expression and graph shown below:

- An integral is always a *sum of products*: a process of multiplication (to calculate the product of two variables) followed by addition (to sum those quantities into a whole).
- The units of measurement for an integral reflect this initial process of multiplication: one unit times some other unit (e.g. kilowatt-hours, foot-pounds, volt-seconds).
- When computing the value of an integral, we must specify both the starting and ending points along the function defining the interval of integration (a and b).
- Geometrically, the integral of a function is the graphical *area* enclosed by the function and the interval boundaries.
- The area enclosed by the function may be thought of as an infinite sum of extremely narrow rectangles, each rectangle having a height equal to one variable (y) and a width equal to the differential of another variable (dx).



Just as division and multiplication are *inverse* mathematical functions (i.e. one “un-does” the other), differentiation and integration are also inverse mathematical functions. The two examples of propane gas flow and mass measurement highlighted in the previous sections illustrates this complementary relationship. We may use differentiation with respect to time to convert a mass measurement (m) into a mass flow measurement (W , or $\frac{dm}{dt}$). Conversely, we may use integration with respect to time to convert a mass flow measurement (W , or $\frac{dm}{dt}$) into a measurement of mass gained or lost (Δm).

Likewise, the common examples of position (x), velocity (v), and acceleration (a) used to illustrate the principle of differentiation are also related to one another by the process of integration. Reviewing the derivative relationships:

$$v = \frac{dx}{dt} \quad \text{Velocity is the derivative of position with respect to time}$$


$$a = \frac{dv}{dt} \quad \text{Acceleration is the derivative of velocity with respect to time}$$

Now, expressing position and velocity as *integrals* of velocity and acceleration, respectively⁹:

$$x = \int v dt \quad \text{Position is the integral of velocity with respect to time}$$

$$v = \int a dt \quad \text{Velocity is the integral of acceleration with respect to time}$$

Differentiation and integration may be thought of as processes transforming these quantities into one another. Note the transformation of units with each operation – differentiation always divides while integration always multiplies:



⁹To be perfectly accurate, we must also include initial values for position and velocity. In other words, $x = \int v dt + x_0$ and $v = \int a dt + v_0$

The inverse nature of these two calculus operations is codified in mathematics as the *Fundamental Theorem of Calculus*, shown here:

$$\frac{d}{dx} \left[\int_a^b f(x) dx \right] = f(x)$$

What this equation tells us is that the derivative of the integral of any continuous function is that original function. In other words, we can take any mathematical function of a variable that we know to be continuous over a certain range – shown here as $f(x)$, with the range of integration starting at a and ending at b – integrate that function over that range, then take the derivative of that result and end up with the original function. By analogy, we can take the *square-root* of any quantity, then *square* the result and end up with the original quantity, because these are inverse functions as well.

A feature of this book which may be helpful to your understanding of derivatives, integrals, and their relation to each other is found in an Appendix section (Appendix A.4 beginning on page 2906). In this section, a series of illustrations provides a simple form of animation you may “flip” through to view the filling and emptying of a water storage tank, with graphs showing stored volume (V) and volumetric flow rate (Q). Since flow rate is the time-derivative of volume ($Q = \frac{dV}{dt}$) and volume change is the time-integral of volumetric flow rate ($\Delta V = \int Q dt$), the animation demonstrates both concepts in action.

1.5 Symbolic versus numerical calculus

Calculus has a reputation for being difficult to learn, and with good reason. The traditional approach to teaching calculus is based on manipulating symbols (variables) in equations, learning how different types of mathematical functions become transformed by the calculus operations of differentiation and integration.

For example, suppose a first-semester calculus student were given the following function to differentiate. The function is expressed as y in terms of x :

$$y = \frac{3x^2 - 2x + 5}{x^2 - 8}$$

A calculus student would first apply two basic rules of symbolic differentiation (namely, the *Power Rule* and the *Quotient Rule*) followed by algebraic distribution and combination of like terms to arrive at the derivative of y with respect to x (written as $\frac{dy}{dx}$) in terms of x :

$$\begin{aligned} \frac{dy}{dx} &= \frac{(x^2 - 8)(6x - 2) - (3x^2 - 2x + 5)(2x)}{(x^2 - 8)^2} \\ \frac{dy}{dx} &= \frac{6x^3 - 2x^2 - 48x + 16 - (6x^3 - 4x^2 + 10x)}{x^4 - 16x^2 + 64} \\ \frac{dy}{dx} &= \frac{2x^2 - 58x + 16}{x^4 - 16x^2 + 64} \end{aligned}$$

The resulting derivative expresses the rate-of-change of y with respect to x of the original function for any value of x . In other words, anyone can now plug any arbitrary value of x they wish into the derivative equation, and the result ($\frac{dy}{dx}$) will tell them *how steep the slope is* of the original function at that same x value¹⁰.

Rules such as the Power Rule and even the Quotient Rule are not difficult to memorize, but they are far from intuitive. Although it is possible to formally prove each one of them from more fundamental principles of algebra, doing so is tedious, and so most students simply resign themselves to memorizing all the calculus rules of differentiation and integration. There are many such rules to memorize in symbolic calculus.

Symbolic integration is even more difficult to learn than symbolic differentiation. Most calculus textbooks reserve pages at the very end listing the general rules of differentiation and integration. Whereas a table of derivatives might occupy a single page in a calculus text, tables of integrals may fill five or more pages!

The next logical topic in the sequence of a calculus curriculum is *differential equations*. A “differential equation” is a function relating some variable to one or more of its own derivatives. To use the variables y and x , a differential equation would be one containing both y and at least one derivative of y ($\frac{dy}{dx}$, $\frac{d^2y}{dx^2}$, $\frac{d^3y}{dx^3}$, etc.). $\frac{dV}{dt} = -kV$ is an example of a simple differential equation. The various forms and solution techniques for different kinds of differential equations are numerous and complex.

¹⁰For instance, at $x = 1$, the original function tells us that y will be equal to $-\frac{6}{7}$. If we plug this same value of 1 into x of the derivative function, the result $\frac{dy}{dx} = -\frac{40}{49}$ tells us the original function $y = f(x)$ has a slope of $-\frac{40}{49}$ when $x = 1$.

It has been said that the laws of the universe are written in the language of calculus. This is immediately evident in the study of physics, but it is also true for chemistry, biology, astronomy, and other “hard sciences.” Areas of applied science including engineering (chemical, electrical, mechanical, and civil) as well as economics, statistics, and genetics would be impoverished if not for the many practical applications of symbolic calculus. To be able to express a function of real-life quantities as a set of symbols, then apply the rules of calculus to those symbols to transform them into functions relating rates of change and accumulations of those real-life quantities, is an incredibly powerful tool.

Two significant problems exist with symbolic calculus, however. The first problem with symbolic calculus is its complexity, which acts as a barrier to many people trying to learn it. It is quite common for students to drop out of calculus or to change their major of study in college because they find the subject so confusing and/or frustrating. This is a shame, not only because those students end up missing out on the experience of being able to see the world around them in a new way, but also because mastery of calculus is an absolute requirement of entry into many professions. One cannot become a licensed engineer in the United States, for example, without passing a series of calculus courses in an accredited university and demonstrating mastery of those calculus concepts on a challenging exam.

The second significant problem with symbolic calculus is its limitation to a certain class of mathematical functions. In order to be able to symbolically differentiate a function (e.g. $y = f(x)$) to determine its derivative ($\frac{dy}{dx}$), we must first have a function written in mathematical symbols to differentiate. This rather obvious fact becomes a barrier when the data we have from a real-life application defies symbolic expression. It is trivial for a first-semester calculus student to determine the derivative of the function $V = 2t^2 - 4t + 9$, but what if V and t only exist as recorded values in a table, or as a trend drawn by a process recorder? Without a mathematical formula showing V as a function of t , none of the rules learned in a calculus course for manipulating those symbols directly apply. The problem is even worse for differential equations, where a great many examples exist that have so far defied solution by the greatest mathematicians.

Such is the case when we apply calculus to recorded values of process variable, setpoint, and controller output in real-world automated processes. A trend showing a PV over time *never* comes complete with a formula showing you $PV = f(t)$. We must approach these practical applications from some perspective other than symbolic manipulation if we are to understand how calculus relates. Students of instrumentation face this problem when learning PID control: the most fundamental algorithm of feedback control, used in the vast majority of industrial processes to regulate process variables to their setpoint values.

An alternative approach to calculus exists which is easily understood by anyone with the ability to perform basic arithmetic (addition, subtraction, multiplication, and division) and sketching (drawing lines and points on a graph). Numerical calculus uses simple *arithmetic* to approximate derivatives and integrals on real-world data. The results are not as precise as with symbolic calculus, but the technique works on *any* data as well as most mathematical functions written in symbolic form. Furthermore, the simplicity of these techniques opens a door to those people who might otherwise be scared away by the mathematical rigor of symbolic calculus. Any way we can find to showcase the beauty and practicality of calculus principles to more people is a good thing!

Suppose we needed to calculate the derivative of some real-world function, such as the volume of liquid contained in a storage vessel. The derivative of volume (V) with respect to time (t) is *volumetric flow rate* ($\frac{dV}{dt}$), thus the time-derivative of the vessel's volume function at any specified point in time will be the net flow rate into (or out of) that vessel at that point in time.


To numerically determine the derivative of volume from raw data, we could follow these steps:

- Choose two values of volume both near the point in time we're interested in calculating flow rate.
- Subtract the two volume values: this will be ΔV .
- Subtract the two time values corresponding to those volume values: this will be Δt .
- Divide ΔV by Δt to approximate $\frac{dV}{dt}$ between those two points in time.

A slightly different approach to numerical differentiation follows these steps:

- Sketch a graph of the volume versus time data for this vessel (if this has not already been done for you by a trend recorder).
- Locate the point in time on this graph you are interested in, and sketch a tangent line to that point (a straight line having the same slope as the graphed data at that point).
- Estimate the rise-over-run slope of this tangent line to approximate $\frac{dV}{dt}$ at this point.

An illustration is a helpful reminder of what differentiation means for any graphed function: the *slope* of that function at a specified point:



Suppose we needed to calculate the integral of some real-world function, such as the flow rate of liquid through a pipe. The integral of volumetric flow (Q) with respect to time (t) is *total volume* (V), thus the time-integral of the flow rate over any specified time interval will be the total volume of liquid that passed by over that time.


To numerically determine the integral of flow from raw data, we could follow these steps:

- Identify the time interval over which we intend to calculate volume, and the duration of each measured data point within that interval.
- Multiply each measured value of flow by the duration of that measurement (the interval between that measurement and the next one) to obtain a volume over each duration.
- Repeat the last step for each and every flow data point up to the end of the interval we're interested in.
- Add all these volume values together – the result will be the approximate liquid volume passed through the pipe over the specified time interval.

A slightly different approach to numerical integration follows these steps:

- Sketch a graph of the flow versus time data for this pipe (if this has not already been done for you by a trend recorder).
- Mark the time interval over which we intend to calculate volume (two straight vertical lines on the graph).
- Use any geometrical means available to estimate the area bounded by the graph and the two vertical time markers – the result will be the approximate liquid volume passed through the pipe over the specified time interval.

An illustration is a helpful reminder of what integration means for any graphed function: the *area* enclosed by that function within a specified set of boundaries:



The next sections of this chapter delve into more specific details of numerical differentiation and integration, with realistic examples to illustrate.

1.6 Numerical differentiation

As we have seen, the concept of *differentiation* is finding the rate-of-change of one variable compared to another (related) variable. In this section, we will explore the practical application of this concept to real-world data, where actual numerical values of variables are used to calculate relative rates of change.

In industrial instrumentation, for example, we are often interested in knowing the rate of change of some process variable (pressure, level, temperature, flow, etc.) over time, and so we may use computers to calculate those rates of change, either after the fact (from recorded data) or in real time (as the data is being received by sensors and acquired by the computer). We may be similarly interested in calculating the rate at which one process variable changes with respect to another process variable, both of which measured and recorded as tables of data by instruments.

Numerical (data-based) differentiation is fundamentally a two-step arithmetic process. First, we must use *subtraction* to calculate the change in a variable between two different points. Actually, we perform this step twice to determine the change in *two variables* which we will later compare. Then, we must use *division* to calculate the ratio of the two variables' changes, one to the other (i.e. the “rise-over-run” steepness of the function's graph).

For example, let us consider the application of pressure measurement for a pipeline. One of the diagnostic indicators of a burst pipeline is that the measured pressure rapidly drops. It is not the existence of low pressure in and of itself that suggests a breach, but rather the *rate* at which the pressure falls that reveals a burst pipe. For this reason, pipeline control systems may be equipped with automatic shut-down systems triggered by rate-of-change pressure calculations.

The association of rapid pressure drop with pipeline ruptures is nothing new to pipeline operations. Here is an illustration taken from page 566 of volume 8 of *Cassier's Magazine* published in the year 1895, showing water pressure measurements taken by a paper strip chart recorder on a city water main line. The pressure drop created by a burst in that 36-inch pipe is clearly seen and commented on the recording:


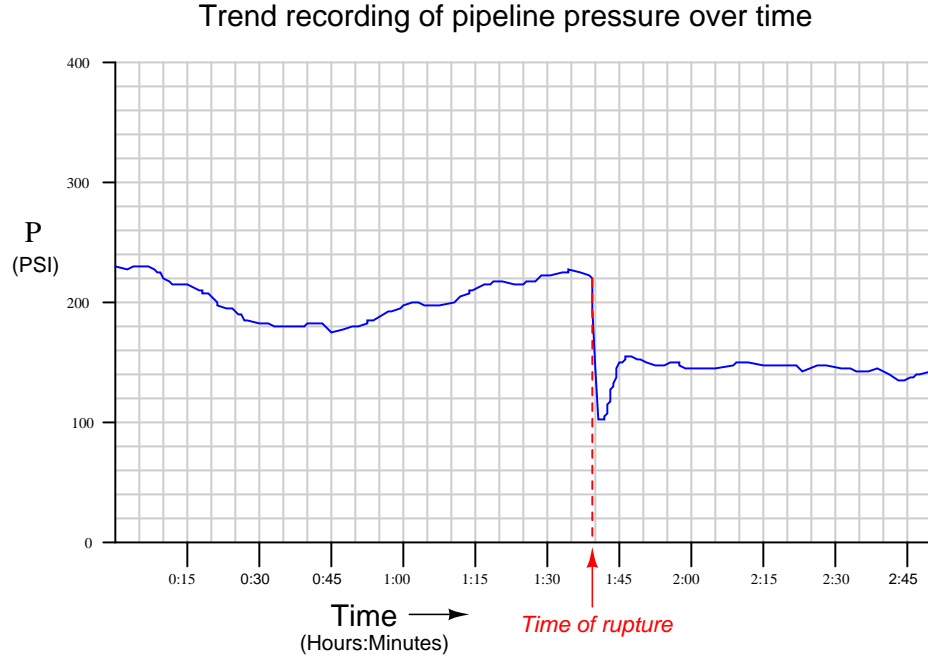


FIG. 5. CHART SHOWING THE BURSTING OF A WATER MAIN.

An example of a modern¹¹ pressure-trend recording during a pipeline rupture is shown here:




While it may seem surprising that pipeline pressure should recover after the low point immediately following the pipe's rupture, it is important to bear in mind that many pipelines are *pressure-controlled* processes. After a pipeline ruptures, the pumping equipment will attempt to compensate for the reduced pressure automatically, which is why the pressure jumps back up (although not to its previous level) after the initial drop.

This phenomenon helps explain why pressure rate-of-change is a more reliable diagnostic indicator of a ruptured pipe than pressure magnitude alone: any automatic rupture-detection scheme based on a simple comparison of pipeline pressure against a pre-set threshold may fail to reliably detect a rupture if the pressure-regulating equipment is able to quickly restore pipeline pressure following the rupture. A rate-of-change system, on the other hand, will still detect the rupture based on the sharp pressure decrease following the break, even if the pressure quickly recovers.

¹¹Unlike the recording shown from *Cassier's Magazine*, which runs chronologically from right to left, modern chart recordings all run from left to right.

A computer tasked with calculating the pressure's rate of change over time ($\frac{dP}{dt}$) would have to continuously sample the pressure value over short time periods, then calculate the quotient of pressure changes over time changes. Given a sample rate of once every 5 minutes, we see how the computer would tabulate the pressure data over time:



Pressure	Time
217.5 PSI	1 hour, 20 minutes
215.0 PSI	1 hour, 25 minutes
222.5 PSI	1 hour, 30 minutes
226.3 PSI	1 hour, 35 minutes
150.0 PSI	1 hour, 40 minutes
150.0 PSI	1 hour, 45 minutes
151.3 PSI	1 hour, 50 minutes
148.8 PSI	1 hour, 55 minutes
145.0 PSI	2 hours, 0 minutes
145.0 PSI	2 hours, 5 minutes

To calculate the rate of pressure change over time in each of these periods, the computer would subtract the two adjacent pressure values, subtract the two corresponding adjacent time values, and then divide those two differences to arrive at a figure in units of PSI per minute. Taking the first two data coordinates in the table as an example:

$$\frac{\Delta P}{\Delta t} = \frac{215.0 \text{ PSI} - 217.5 \text{ PSI}}{1:25 - 1:20} = \frac{-2.5 \text{ PSI}}{5 \text{ min}} = -0.5 \frac{\text{PSI}}{\text{min}}$$

The sample period where the computer would detect the pipeline rupture lies between 1:35 and 1:40. Calculating this rate of pressure change:

$$\frac{\Delta P}{\Delta t} = \frac{150.0 \text{ PSI} - 226.3 \text{ PSI}}{1:40 - 1:35} = \frac{-76.3 \text{ PSI}}{5 \text{ min}} = -15.26 \frac{\text{PSI}}{\text{min}}$$

Clearly, a pressure drop rate of -15.26 PSI per minute is far greater than a typical drop of -0.5 PSI per minute, thus signaling a pipeline rupture.


As you can see, the pipeline monitoring computer is not technically calculating *derivatives* ($\frac{dP}{dt}$), but rather *difference quotients* ($\frac{\Delta P}{\Delta t}$). Being a digital device, the best it can ever do is perform calculations at discrete points in real time. It is evident that calculating rates of change over 5-minute period misses a lot of detail¹². The actual rate of change at the steepest point of the pressure drop *far* exceeds -15.26 PSI per minute.

It is possible for us to calculate the instantaneous rate-of-change of pressure ($\frac{dP}{dt}$) at the moment of the rupture by examining the graph and sketching a straight line called a *tangent line* matching the slope where the graph is steepest. Our goal is to calculate the exact slope of that single (steepest) point on that graph, rather than an estimate of slope between two points as the computer did. In essence, the computer “drew” short line segments between pairs of points and calculated the slopes (rise-over-run) of those line segments. The slope of each line segment¹³ is a difference quotient: $\frac{\Delta P}{\Delta t}$. The slope of a tangent line matching the slope at a single point on the function graph, however, is a derivative: $\frac{dP}{dt}$.

¹²Not only does a 5-minute rate calculation period miss a lot of detail, but it also results in a time delay of (up to) 5 minutes detecting a pipeline rupture.

¹³The technical term for a line passing through a pair of points on a curve is called a *secant line*.

First we sketch a tangent line (by hand) matching the steepest portion of the pressure trend graph. Then, we calculate the slope of a tangent line by marking convenient points¹⁴ where the line intersects major division marks on the graph's graduated scale, then calculating rise over run:



$$\frac{dP}{dt} = \frac{150.0 \text{ PSI} - 340.0 \text{ PSI}}{1:40:00 - 1:37:30} = \frac{-190.0 \text{ PSI}}{2.5 \text{ min}} = -76.0 \frac{\text{PSI}}{\text{min}}$$

This distinction between calculating difference quotients ($\frac{\Delta P}{\Delta t}$) and calculating true derivative values ($\frac{dP}{dt}$) becomes less and less significant as the calculation period shortens. If the computer could sample and calculate at infinite speed, it *would* generate true derivative values instead of approximate derivative values.

¹⁴Please note that the pipeline pressure is *not* actually 340.0 PSI at a time of 1:37:30. This is simply a coordinate convenient to mark because it how it lines up with the divisions on the trend display. We choose coordinate points on the tangent line easy to visually discern, then calculate the tangent line's slope using those coordinates.

An algorithm applicable to calculating rates of change in a digital computer is shown here, using a notation called *pseudocode*¹⁵. For more information on pseudocode, refer to section 29.15.1 beginning on page 2388. Each line of text in this listing represents a command for the digital computer to follow, one by one, in order from top to bottom. The LOOP and ENDLOOP markers represent the boundaries of a program *loop*, where the same set of encapsulated commands are executed over and over again in cyclic fashion:

Pseudocode listing

```

LOOP
  SET x = analog_input_N      // Update x with the latest measured input
  SET t = system_time         // Sample the system clock

  SET delta_x = x - last_x    // Calculate change in x
  SET delta_t = t - last_t    // Calculate change in t (time)

  SET rate = (delta_x / delta_t) // Calculate ratio of changes

  SET last_x = x              // Update last_x value for next program cycle
  SET last_t = t              // Update last_t value for next program cycle
ENDLOOP

```


Each SET command tells the computer to assign a numerical value to the variable on the left-hand side of the equals sign (=), according to the value of the variable or expression on the right-hand side of the equals sign. Text following the double-dash marks (//) are *comments*, included only to help human readers interpret the code, not for the computer's benefit.

This computer program uses two variables to “remember” the values of the input (x) and time (t) from the previous scan, named `last_x` and `last_t`, respectively. These values are subtracted from the current values for x and t to yield differences (`delta_x` and `delta_t`, respectively), which are subsequently divided to yield a difference quotient. This quotient (`rate`) may be sampled in some other portion of the computer's program to trigger an alarm, a shutdown action, or simply display and/or record the rate value for a human operator's benefit.

The time period (Δt) for this program's difference quotient calculation is simply how often this algorithm “loops,” or repeats itself. For a modern digital microprocessor, this could be many thousands of times per second.

¹⁵“Pseudocode” is a name given to any imaginary computer language used for the purpose of illustrating some procedure or concept without having to make reference to any particular (real) computer programming language. I could have just as well shown you the same algorithm using BASIC, C, or Java code, but pseudocode does just as well without the burden of introducing unfamiliar syntax to the reader.

If a nearly-instantaneous calculation is required for a rate-of-change variable, we may turn to an older technology using analog electronic circuitry. Such a *differentiator* circuit uses the natural behavior of a capacitor to generate an output voltage proportional to the instantaneous rate-of-change of the input voltage:



$$V_{out} = -RC \frac{dV_{in}}{dt}$$

The negative feedback of the operational amplifier forms a *virtual ground* at the node where the capacitor, resistor, and inverting input connect. This means the capacitor “sees” the full input voltage (V_{in}) at all times. Current through a capacitor is a direct function of the voltage’s time-derivative:


$$I = C \frac{dV}{dt}$$

This current finds its way through the feedback resistor, developing a voltage drop that becomes the output signal (V_{out}). Thus, the output voltage of this analog differentiator circuit is directly proportional to the time-derivative of the input voltage (i.e. the input voltage’s rate-of-change).

It is indeed impressive that such a simple circuit, possessing far fewer components than a microprocessor, is actually able to do a *better* job at calculating the real-time derivative of a changing signal than modern digital technology. The only real limitations to this device are accuracy (tolerances of the components used) and the bandwidth of the operational amplifier.

It would be a mistake, though, to think that an analog differentiator circuit is better suited to industrial applications of rate calculation than a digital computer, even if it does a superior job differentiating live signals. A very good argument for favoring difference quotients over actual derivatives is the presence of *noise* in the measured signal. A true differentiator, calculating the actual time-derivative of a live signal, will pick up on *any* rise or fall of the signal over time, no matter how brief. This is a serious problem when differentiating real-world signals, because noise (small amounts of “jittering” in the signal caused by any number of phenomena) will be interpreted by a perfect differentiator as very large rates of change over time.

A close look at the previous pipeline pressure trend illustrates this problem. Note the areas circled (in red) on the graph, representing relatively small increases and decreases in signal occurring over very short periods of time:



Although each “step” in pressure at these circled locations is small in amplitude, each one occurs over an extremely brief time increment. Thus, each of these steps has a nearly *infinite* rate of change (i.e. a vertical slope). Any rate-of-change sensing system able to apply true differentiation to the pressure signal would falsely declare an alarm (high rate-of-change) condition every time it encountered one of these “steps” in the signal. This means that even under perfectly normal operating conditions the rate-detection system would periodically declare an alarm (or perhaps shut the pipeline down!) given the inevitable presence of small noise-induced¹⁶ “jitters” in the signal.


The best solution to this problem is to use a digital computer to calculate rates of change, setting the calculation period time slow enough that these small “jitters” will be averaged to very low values, yet fast enough that any serious pressure rate-of-change will be detected if it occurs. Back in the days when analog electronic circuits were the *only* practical option for calculating rates of signal change, the solution to this problem was to place a low-pass filter before the differentiator circuit to block such noise from ever reaching the differentiator.

Differentiation with respect to *time* has many applications, but there are other applications of differentiation in industrial measurement and control that are not time-based. For example, we may

¹⁶Another source of trouble for differentiation of live signals is when the signal originates from a digital sensor. Digital devices, by their very nature, break analog signals into a series of discrete amplitude steps. As a digital process transmitter encounters a steadily increasing or decreasing process variable, its output rises or falls in discrete “jumps” rather than continuously as a fully analog transmitter would. Now, each of these jumps is quite small, but since each one occurs almost instantly it still translates into an extremely large rate-of-change when detected by a differentiator sampling over small time increments or sampling continuously (as in the case of an analog differentiator circuit). This means the problem of false rates-of-change exists *even in perfectly noiseless systems*, when the detection device (and/or the information channel to the monitoring system) is digital rather than analog.


use differentiation to express the *sensitivity* of a non-linear device in terms of the rate-of-change of output over input.

One such application is the sensitivity of a mechanism called a *baffle/nozzle* assembly used in many pneumatic instruments to convert a small physical motion (x) into an air pressure signal (P). This very simple mechanism uses a flat piece of sheet metal (the baffle) to restrict air flow out of a small nozzle, causing a variable “backpressure” at the nozzle to develop as the baffle-to-nozzle clearance changes:



The graph expressing the relationship between P and x is clearly non-linear, having different slopes ($\frac{dP}{dx}$) at different points along its range. When used as part of the feedback mechanism for a self-balancing instrument, the purpose of the baffle/nozzle assembly is to detect baffle motion as sensitively as possible: that is, to generate the greatest change in pressure (ΔP) for the least change in motion (Δx). This means the designer of the pneumatic instrument should design it in such a way that the normal baffle/nozzle clearance gap rests at a point of maximum slope (maximum $\frac{dP}{dx}$) on the graph.

Sketching a tangent line near the point of maximum slope (maximum “steepness” on the graph) allows us to approximate the rate of change at that point:



Choosing convenient points¹⁷ on this tangent line aligning with major divisions on the graph’s scales, we find two coordinates we may use to calculate the derivative of the curve at its steepest point:

$$\frac{dP}{dx} = \frac{0 \text{ PSI} - 18 \text{ PSI}}{0.0025 \text{ inch} - 0.001 \text{ inch}} = \frac{-18 \text{ PSI}}{0.0015 \text{ inch}} = -12000 \text{ PSI per inch}$$

The phenomenally large value of -12000 PSI per inch is a *rate* of pressure change to clearance (baffle-nozzle gap) change. Do not mistakenly think that this value suggests the mechanism could ever develop a pressure of 12000 PSI – it is simply describing the extreme sensitivity of the mechanism in terms of PSI change per unit change of baffle motion. By analogy, just because an automobile travels at a *speed* of 70 miles per hour does not mean it must travel 70 miles in *distance*!

It should be clear from an examination of the graph that this high sensitivity extends approximately between the pressure values of 9 and 14 PSI. Outside of those pressure values, the graph’s slope begins to decrease. While still sensitive, the baffle/nozzle mechanism will not be as sensitive to baffle motion outside those pressure values as it is within.

¹⁷Once again, we are looking for points where the tangent line happens to intersect with major divisions on the graph’s scale. This makes it relatively easy to calculate the line’s slope, since the pressure and distance values for those coordinates are easy to read.

1.7 Numerical integration

As we have seen, the concept of *integration* is finding the accumulation of one variable multiplied by another (related) variable. In this section, we will explore the practical application of this concept to real-world data, where actual numerical values of variables are used to calculate accumulated sums.

In industrial instrumentation, for example, we are often interested in calculating the accumulation of some process fluid based on a measured flow *rate* of that fluid. The rate is, of course, expressed in either mass or volume units per unit time (e.g. gallons per minute), but the total accumulated quantity will be expressed plainly in either mass or volume units (e.g. gallons). We may use computers to calculate those accumulated quantities, either after the fact (from recorded data) or in real time.

Numerical (data-based) integration is fundamentally a two-step arithmetic process. First, we must use *multiplication* to calculate the product of a variable and a small increment of another variable (a change in the second variable between two different points). Then, we must use *addition* to calculate the accumulated sum of the products.

To illustrate, we will first focus on the integration of a flow measurement signal with respect to time. The flow rate of any fluid is always expressed in units of volume or mass *per unit time*. Common volumetric flow units are gallons *per minute*, liters *per second*, cubic feet *per day*, etc. Common mass flow units are pounds *per hour*, kilograms *per minute*, slugs *per second*, etc. If we desire to calculate the volume or mass of fluid passed through a pipe – representing fluid added to or removed from a system – over some interval of time, we may do so by integrating flow rate with respect to time:

$$\Delta V = \int_a^b Q dt$$

$$\Delta m = \int_a^b W dt$$

Where,

ΔV = Volume of fluid added or removed

Q = Volumetric flow rate of fluid

Δm = Mass of fluid added or removed

W = Mass flow rate of fluid

a = Starting point of integration interval

b = Ending point of integration interval

t = Time

As always, integration is fundamentally a matter of *multiplying* one variable by small increments of another variable. If a flow rate is integrated with respect to time, the result is that the unit for time becomes eliminated. Gallons per minute, for example, becomes gallons after integration; kilograms per second becomes kilograms; etc.

The elimination of time units is also evident if we re-write the integrands in the previous equations to show volumetric and mass flow rates (Q and W , respectively) as the rates of change they are ($Q = \frac{dV}{dt}$ and $W = \frac{dm}{dt}$):

$$\Delta V = \int_a^b \frac{dV}{dt} dt$$

$$\Delta m = \int_a^b \frac{dm}{dt} dt$$


It should be clear that the time differentials (dt) cancel in each integrand, leaving:

$$\Delta V = \int_a^b dV$$

$$\Delta m = \int_a^b dm$$

Since we know the integral symbol (\int) simply means the “continuous sum of” whatever follows it, we may conclude in each case that the continuous sum of infinitesimal increments of a variable is simply a larger change of that same variable. The continuous summation of dV is simply the total change in V over the interval beginning at time a and ending at time b ; likewise, the continuous summation of dm is simply the total change in m over the interval beginning at time a and ending at time b .

A flowmeter measuring the flow rate of a fluid outputs a signal representing either volume or mass units passing by per unit time. Integrating that signal with respect to time yields a value representing the total volume or mass passed through the pipe over a specific interval. A physical device designed to perform this task of integrating a signal with respect to time is called an *integrator* or a *totalizer*:



An example of a flow integrator, or flow totalizer, made for pneumatic instrument systems is the Foxboro model 14. A view of this instrument's front face shows an odometer-style display, in this particular case showing the total number of pounds (lb) of fluid passed through the pipe, with a multiplying factor of 10:



The fact that this instrument's display resembles the odometer of an automobile is no coincidence. Odometers are really just another form of mechanical integrator, “totalizing” the distance traveled by a vehicle. If the speedometer of a vehicle registers speed (v) in units of miles per hour, then the odometer will accumulate a distance (Δx) in units of miles, since distance (miles) is the time-integral of speed (miles per hour):

$$\Delta x = \int_a^b v \, dt \quad \dots \text{ or } \dots \quad \Delta x = \int_a^b \frac{dx}{dt} \, dt$$


$$[\text{miles}] = \int_a^b \left(\left[\frac{\text{miles}}{\text{hour}} \right] [\text{hours}] \right)$$

In this particular case, where the flowmeter measures pounds per hour, and the integrator registers accumulated mass in pounds, the integration of units is as follows:

$$\Delta m = \int_a^b W \, dt \quad \dots \text{ or } \dots \quad \Delta m = \int_a^b \frac{dm}{dt} \, dt$$

$$[\text{pounds}] = \int_a^b \left(\left[\frac{\text{pounds}}{\text{hour}} \right] [\text{hours}] \right)$$

The Foxboro model 14 used a turbine wheel driven by a jet of compressed air from a nozzle. The wheel's speed was made proportional to the process fluid flow rate sensed by a pneumatic DP transmitter. As process flow rate increased, the wheel spun faster. This spinning wheel drove a gear-reduction mechanism to slowly turn the odometer-style numerals, registering total fluid quantity passed through the flowmeter:



As pneumatic signal pressure (3-15 PSI) from a pneumatic flow transmitter entered the brass bellows of this instrument, it pressed down on a lever, forcing a baffle toward a nozzle. As nozzle backpressure rose, amplified air pressure spun the turbine wheel to drive the integrating “odometer” display. Mounted on the turbine wheel was a set of fly-weights, which under the influence of centrifugal force would press upward on the lever to re-establish a condition of force-balance to maintain a (relatively) constant baffle-nozzle gap. Thus, the force-balance mechanism worked to establish an accurate and repeatable relationship¹⁸ between instrument signal pressure and integration rate.

¹⁸The Foxboro model 14 totalizer's design was quite ingenious, since centrifugal force varies with the *square* of angular velocity. This had the effect of naturally performing the *square-root* characterization required of most pneumatic flow-measuring instruments due to the quadratic nature of most primary flow-sensing elements (e.g. orifice plate, venturi tubes, pitot tubes, etc.).

A very different style of integrator appears here, as part of the controller for a ball mill used to crush limestone into small pieces for the manufacturing of concrete. Limestone is fed into the ball mill on a device called a *weighfeeder*, which measures the mass of limestone as it passes over a conveyor belt. The controller maintains a limestone “flow rate” at a setpoint specified in tons per hour (mass flow of solid material). The red LED digital display shows the total number of tons passed through the mill:



The units involved in the integration of limestone “flow” into the ball mill are slightly different from the example shown with the Foxboro model 14 totalizer, but the concept is the same:

$$\Delta m = \int_a^b W dt$$

$$[\text{tons}] = \int_a^b \left(\left[\frac{\text{tons}}{\text{hour}} \right] [\text{hours}] \right)$$

As with all cases of numerical integration, an essential piece of information to know when “totalizing” any rate is the initial quantity at the start of the totalization interval. This is the *constant of integration* mentioned previously. For flow totalization, this constant would be the initial volume of fluid recorded at the starting time. For an automobile’s odometer, this constant is the initial “mileage” accumulated prior to driving on a trip¹⁹.

¹⁹Vehicles equipped with a *trip odometer* allow the driver to reset this integration constant to zero at will, thus

An algorithm applicable to integrating real signals with respect to time in a digital computer is shown here, once again using “pseudocode” as the computer language. Each line of text in this listing represents a command for the digital computer to follow, one by one, in order from top to bottom. The LOOP and ENDL00P markers represent the boundaries of a program *loop*, where the same set of encapsulated commands are executed over and over again in cyclic fashion:

Pseudocode listing

```

LOOP
  SET x = analog_input_N    // Update x with the latest measured input
  SET t = system_time      // Sample the system clock

  SET delta_t = t - last_t  // Calculate change in t (time)

  SET product = x * delta_t // Calculate product (integrand)
  SET total = total + product // Update the running total

  SET last_t = t           // Update last_t value for next program cycle
ENDLOOP

```

This computer program uses a variable to “remember” the value of time (t) from the previous scan, named `last_t`. This value is subtracted from the current value for t to yield a difference (`delta_t`), which is subsequently multiplied by the input value x to form a product. This product is then added to an accumulating total (named `total`), representing the integrated value. This “total” value may be sampled in some other portion of the computer’s program to trigger an alarm, a shutdown action, or simply display and/or record the totalized value for a human operator’s benefit.

The time period (Δt) for this program’s difference quotient calculation is simply how often this algorithm “loops,” or repeats itself. For a modern digital microprocessor, this could be upwards of many thousands of times per second. Unlike differentiation, where an excessive sampling rate may cause trouble by interpreting noise as extremely high rates of change, there is no danger of excessive sampling when performing numerical integration. The computer may integrate as fast as it can with no ill effect.

One of the fundamental characteristics of integration is that it *ignores* noise, which is a very good quality for industrial signal processing. Small “jittering” in the signal tends to be random, which means for every “up” spike of noise, one may expect a comparable “down” spike (or collection of “down” spikes having comparable weight) at some later time. Thus, noise tends to cancel itself out when integrated over time.

As with differentiation, applications exist for integration that are not time-based. One such application is the calculation of mechanical *work*, defined as the product of force and displacement (distance moved). In mechanical systems where there is no energy dissipated due to friction, work results in a change in the energy possessed by an object.

allowing the tracking of mileage for individual trips instead of over the life of the automobile.

For example, if we use a hoist to lift a mass weighing 700 pounds straight up against gravity a distance of 3 feet, we will have done 2100 foot-pounds of work. The work done on the mass increases its potential energy (ΔE) by 2100 foot-pounds:

$$\Delta E = Fx$$

Where,

ΔE = Change in potential energy resulting from work, in joules (metric) or foot-pounds (British)


F = Force doing the work, in newtons (metric) or pounds (British)

x = Displacement over which the work was done, in meters (metric) or feet (British)

We may also express this change in potential energy as an integral of force (F) multiplied by infinitesimal increments in displacement (dx) over some interval (from a to b), since we know integration is nothing more than a sophisticated way to multiply quantities:

$$\Delta E = \int_a^b F dx$$

Like any other integral, the energy change effected by lifting this mass a vertical distance may be represented graphically as the *area* enclosed by the graph. In this case, the area is very simple to calculate, being a simple rectangle (height times width):



Lifting the mass vertically constitutes a *positive* change in potential energy for this object, because each displacement differential (dx) is a positive quantity as we move from a height of 0 feet to a height of 3 feet:

$$2100 \text{ ft-lbs} = \int_{0 \text{ ft}}^{3 \text{ ft}} (700 \text{ lbs}) dx$$

A natural question to ask at this point is, *what would the resulting change in energy be if we lowered the mass from its height of 3 feet back down to 0 feet?* Doing so would cover the exact same distance (3 feet) while exerting the exact same amount of suspending force (700 lbs), and so we can safely conclude the work will have an absolute magnitude of 2100 ft-lbs. However, if we lower the mass, each displacement differential (dx) will be a negative quantity²⁰ as we move from a greater height to a lesser height. This makes the work – and the resulting energy change – a negative quantity as well:

$$-2100 \text{ ft-lbs} = \int_{3ft}^{0ft} (700 \text{ lbs}) dx$$

This means if we raise the mass to a height of 3 feet, then lower it back to its original starting height of 0 feet, the total change in potential energy will be zero:

$$0 \text{ ft-lbs} = \int_{0ft}^{3ft} (700 \text{ lbs}) dx + \int_{3ft}^{0ft} (700 \text{ lbs}) dx$$


This is true for any integral having an interval of zero (same starting and ending values), regardless of the integrand's value at any point in time:

$$0 \text{ ft-lbs} = \int_a^a F dx$$

²⁰As we lower the mass to ground level, height (x) goes from being a positive value to zero. This means each differential (infinitesimal change in value) for x will be negative, thus causing the integrand $F dx$ to have a negative value and thus causing the integrated total (work) to be negative as well.

The integration of force and displacement to calculate potential energy change really shows its utility when the force changes as a function of displacement. A classic example of this is the compression of a mechanical spring, described in section 2.8.3 beginning on page 91.


One practical example of this sort of calculation is the determination of energy stored in an archer's bow when drawn to a certain displacement. The so-called *force-draw curve* of a longbow is nearly ideal for a theoretical spring, with force increasing linearly as the string is drawn back by the archer. The force-draw curve for a compound bow²¹ is quite nonlinear, with a much lesser holding force required to maintain the bow at full draw:



The force required to draw a compound bow rises sharply during the first few inches of draw, peaks during the region where the archer's arms are ideally angled for maximum pulling strength, then "lets off" toward the end where the archer's drawing arm is weakest in the "holding" position. The result is a bow that requires substantial force to draw, but is relatively easy to hold in fully-drawn position.

²¹While a longbow is really nothing more than a long and flexible stick with a straight string drawn across it, a compound bow is a sophisticated machine with multiple passes of string and cam-shaped pulleys providing the nonlinear force-draw relationship.

While the compound bow may be easier to hold at full draw than the longbow, for any given holding force the compound bow stores *much* more energy than the longbow, owing to the far greater *area* (force-displacement integral) enclosed by the curve:




This is why a compound bow is so much more powerful than a longbow or a “recurve” bow with the same holding force: the energy represented by the greater area underneath the force-draw curve equates to greater energy imparted to the arrow when released, and therefore greater kinetic energy in the arrow during flight.

Like any other form of mechanical work, the energy invested into the bow by the archer is readily calculated and expressed in units of force \times displacement, typically newton-meters (joules) in metric units and foot-pounds in British units. This stands to reason, since we know integration is fundamentally a matter of *multiplying* quantities together, in this case force (pull) and displacement (draw).

To actually calculate the amount of energy stored in a fully-drawn bow, we could measure both force and displacement with sensors as the archer draws the bow, with a computer numerically integrating force over increments of draw in real time. Another method would be to simply graph force versus draw as we have done here, then use geometric methods²² to approximate the area underneath the curve.



²²One simple way to do this is to cover the entire integration area using nothing but rectangles and triangles, then measuring all the sketched shapes to totalize their areas.

A more sophisticated example of numerical integration used to calculate work is that of a heat engine, where a piston compresses an enclosed gas:



As the piston is pushed farther into the cylinder, the gas becomes compressed, exerting more force on the piston. This requires an ever-increasing application of force to continue the piston's motion. Unlike the example where a mass of constant weight was lifted against the pull of gravity, here the force is a dynamically changing variable instead of a constant. The graph shows this relationship between piston displacement and piston force.

If we push the piston into the cylinder, the force increases as the displacement decreases. The change in energy is described by the integral of force with respect to displacement, graphically equivalent to the area underneath the force curve:



$$\Delta E = \int_a^b F dx$$


If we slowly allow the piston to return to its original position (letting the pressure of the enclosed gas push it back out), the piston's force decreases as displacement increases. The force/displacement relationship is the same as before, the only difference being the direction of travel is opposite. This means the change in energy is happening over the same interval, in reverse direction (from b to a now instead of from a to b). Expressed as an integral:

$$\Delta E = \int_b^a F dx$$

As we have already learned, a reversal of direction means the sign of the integral will be opposite. If pushing the piston farther inside the cylinder represented work being done on the enclosed gas by the applied force, now the gas will be doing work on the source of the applied force as the piston returns to its extended position.


This means we will have done zero net work by pushing the piston into the cylinder and then letting it spring back out to its original position, just as we performed zero net work by lifting a mass 3 feet in the air and then letting it back down.

In order that this piston/cylinder mechanism might function as an engine, we must have some way of making the energy change greater in one direction than the other. This is done by *heating* the enclosed gas at the point of greatest compression. In a spark-ignition engine, the gas is actually a mixture of air and fuel, ignited by an electric spark. In a compression-ignition (diesel) engine, the gas is pure air, with fuel injected at the last moment to initiate combustion. The addition of heat (from combustion) will cause the gas pressure to rise, exerting more force on the piston than what it took to compress the gas when cold. This increased force will result in a *greater* energy change with the piston moving out of the cylinder than with the piston moving in:



Representing the work done by the hot gas as the area enclosed by the curve makes this clear: more mechanical energy is being released as the piston travels from b to a during the “power stroke” than the amount of energy invested in compressing the gas as the piston traveled from a to b during the “compression stroke.” Thus, an internal combustion engine produces mechanical power by repeatedly compressing a cold gas, heating that gas to a greater temperature, and then expanding that hot gas to extract energy from it.

At the conclusion of the power stroke, a valve opens to exhaust the hot gas and another valve opens to introduce cold gas. This places the piston and cylinder in the original condition, ready for another set of compression, ignition, and power strokes. This cycle is sometimes represented as a closed “loop” on the force/displacement graph, like this:



The amount of net energy output by the engine at the conclusion of each cycle is equivalent to the area enclosed by the loop. This is the difference in areas (integrals) between the “compression” and “power” strokes. Any design change to the engine resulting in a greater “loop” area (i.e. less energy required to compress the gas, and/or more energy extracted from its expansion) results in a more powerful engine. This is why heat engines output the most power when the difference in temperatures (cold gas versus heated gas) is greatest: a greater temperature shift results in the two curves being farther apart vertically, thus increasing the area enclosed by the “loop.”

References

Hague, Charles A. “The Recording Gauge Applied to Water Pressure and Other Uses”, *Cassier’s Magazine* Volume 8, London, England, 1895.

Keisler, H. Jerome, *Elementary Calculus – An Infinitesimal Approach*, Second Edition, University of Wisconsin, 2000.

Stewart, James, *Calculus: Concepts and Contexts*, 2nd Edition, Brooks/Cole, Pacific Grove, CA, 2001.

Thompson, Silvanus P. and Gardner, Martin, *Calculus Made Easy*, St. Martin’s Press, New York, NY, 1998.

Chapter 2

Physics

Everything there is to know about physics – excerpted from “Everything you need to know about school” in the September 16, 2008 edition of the Seattle periodical *The Stranger*:

If stuff is still, it doesn't like to move; if stuff is moving, it doesn't like to stop. The more stuff you are trying to move, the more you need to push it to speed it up. When you push on stuff, it pushes back on you. Stuff likes other stuff, from a distance at least. Stuff likes becoming more chaotic, but cannot be created or destroyed; stuff can be rearranged. (All that is Isaac Newton.) Energy, like stuff, cannot be created or destroyed, only changed from one kind to another. Energy can be stored in movement, bonds between stuff, and many other places. Changes in how energy is stored allow us to do things – like bake, drive, get up tall buildings, and kill each other. (That's Sadi Carnot.) Also, stuff *is* energy. Stuff is a lot of energy. (That's Albert Einstein.) Compress plutonium with explosives and the atoms fission, releasing the energy stored in stuff. When the energy is released in downtown Nagasaki, you kill about 40,000 people right away and another 40,000 over time. (Thanks, Enrico Fermi!)

There is a lot of truth to this quote, despite its sarcastic tone. Physics really is the study of matter and energy, and of the laws governing the interactions of both. The movement of matter largely follows the laws described by Isaac Newton (so much so that this category of physics is often referred to as *Newtonian physics*). The flow and transformation of energy is a category called *thermodynamics*. Albert Einstein went on to equate matter with energy through his famous equation $E = mc^2$ which helped usher in the age of nuclear physics.

Physics is important to the study of industrial instrumentation because its laws describe and explain so many applications of measurement and control. As a teacher, one of the things I like to tell students is that their chosen field is really nothing more than applied physics. Granted, there aren't too many applications of macroscopic projectile motion in industrial measurement and control, but on a microscopic level this is what all moving fluids are: an extremely large collection of independent molecules obeying Newtonian laws of motion. Physics is the study of *stuff*, while instrumentation is the discipline of measuring and controlling stuff. The better you understand stuff, the better you will be able to measure and control that stuff. Enjoy!

2.1 Terms and Definitions

Mass (m) is the opposition an object has to acceleration (changes in velocity). *Weight* is the force (F) imposed on a mass by a gravitational field. Mass is an intrinsic property of an object, regardless of the environment. Weight, on the other hand, depends on the strength of the gravitational field in which the object resides. A 20 kilogram slug of metal has the exact same mass whether it rests on Earth, or in the zero-gravity environment of outer space, or on the surface of the planet Jupiter. However, the *weight* of that mass depends on gravity: zero weight in outer space (where there is no gravity to act upon it), some weight on Earth, and a much greater amount of weight on the planet Jupiter (due to the much stronger gravitational field of that planet).

Since mass is the opposition of an object to changes in velocity (acceleration), it stands to reason force, mass, and acceleration for any particular object are directly related to one another:

$$F = ma$$

Where,

F = Force in newtons (metric) or pounds (British)

m = Mass in kilograms (metric) or slugs (British)

a = Acceleration in meters per second squared (metric) or feet per second squared (British)

If the force in question is the weight of the object, then the acceleration (a) in question is the acceleration constant of the gravitational field where the object resides. For Earth at sea level, $a_{gravity}$ is approximately 9.81 meters per second squared, or 32.2 feet per second squared. Earth's gravitational acceleration constant is usually represented in equations by the variable letter g instead of the more generic a .

Since acceleration is nothing more than the rate of velocity change with respect to time, the force/mass equation may be expressed using the calculus notation of the first derivative:

$$F = m \frac{dv}{dt}$$

Where,

F = Force in newtons (metric) or pounds (British)

m = Mass in kilograms (metric) or slugs (British)

v = Velocity in meters per second (metric) or feet per second (British)

t = Time in seconds

Since velocity is nothing more than the rate of position change with respect to time, the force/mass equation may be expressed using the calculus notation of the second derivative (acceleration being the derivative of velocity, which in turn is the derivative of position):

$$F = m \frac{d^2x}{dt^2}$$

Where,

F = Force in newtons (metric) or pounds (British)

m = Mass in kilograms (metric) or slugs (British)

x = Position in meters (metric) or feet (British)

t = Time in seconds


Mass density (ρ) for any substance is the proportion of mass to volume. *Weight density* (γ) for any substance is the proportion of weight to volume.

Just as weight and mass are related to each other by gravitational acceleration, weight density and mass density are also related to each other by gravity:

$$F_{weight} = mg \quad \text{Weight and Mass}$$

$$\gamma = \rho g \quad \text{Weight density and Mass density}$$

2.2 Metric prefixes




2.3 Areas and volumes

Area refers to the size of two-dimensional surface. *Volume* refers to the size of a three-dimensional space. To put both these measures into context; the question of how much paint will be required to adequately cover a house is one of area, while the question of how much water will be required to fill a pond is one of volume.


Some units of measurement for area and volume are nothing more than compounded linear units. Ten *centimeters* is an expression of distance, while ten *square centimeters* (cm^2) is an expression of area, and ten *cubic centimeters* (cm^3) is an expression of volume. It important to note that the modifiers “square” and “cubic” do not in any way imply the object in question is square or cubic in shape. It is perfectly reasonable to measure the area of a *circle*, for instance, using the unit of *square centimeters*.

Other units of spatial measurement are specific to area or to volume. The *acre*, for example, is a unit of area measurement developed for the purpose of quantifying the size of land plots, one acre being equivalent to 43560 square feet. An example of a unit specifically devoted to volume measurement is the *liter*, equivalent to 1000 cubic centimeters.


2.3.1 Common geometric shapes

Triangle

$$\text{Area } A = \frac{1}{2} bh$$


Trapezoid

$$\text{Area } A = \frac{1}{2} (b + B)h$$

Rectangle


$$\text{Perimeter } P = 2x + 2y$$

$$\text{Area } A = xy$$

Rectangular solid


$$\text{Surface area } A = 2xy + 2yz + 2xz$$

$$\text{Volume } V = xyz$$

Circle


$$\text{Circumference } C = \pi D = 2\pi r$$

$$\text{Area } A = \pi r^2$$

Sphere


$$\text{Surface area } A = 4\pi r^2$$

$$\text{Volume } V = \frac{4}{3} \pi r^3$$

Right circular cylinder

Surface area $A = 2\pi r^2 + 2\pi rh$


Volume $V = \pi r^2 h$

Cone

Surface area $A = \pi r \sqrt{r^2 + h^2} + \pi r^2$

Volume $V = \frac{1}{3} \pi r^2 h$

Note: the volume of any pyramid or cone is one-third the product of its height (h) and the area of its base.

Tetrahedron

Volume $V = \frac{1}{3} xyh$

2.4 Unit conversions and physical constants

Converting between disparate units of measurement is the bane of many science students. The problem is worse for students in the United States of America, who must work with British (“Customary”) units such as the pound, the foot, the gallon, etc. World-wide adoption of the metric system would go a long way toward alleviating this problem, but until then it is important for students to master the art of unit conversions¹.

It is possible to convert from one unit of measurement to another by use of tables designed expressly for this purpose. Such tables usually have a column of units on the left-hand side and an identical row of units along the top, whereby one can look up the conversion factor to multiply by to convert from any listed unit to any other listed unit. While such tables are undeniably simple to use, they are practically impossible to memorize.

A better way to convert between different units is shown in the next subsection.

¹An interesting point to make here is the United States did get something right when they designed their monetary system of dollars and cents. This is essentially a *metric* system of measurement, with 100 cents per dollar. The founders of the USA wisely decided to avoid the utterly confusing denominations of the British, with their pounds, pence, farthings, shillings, etc. The denominations of penny, dime, dollar, and eagle (\$10 gold coin) comprised a simple power-of-ten system for money. Credit goes to France for first adopting a metric system of general weights and measures as their national standard.

2.4.1 Unity fractions

An important principle in the physical sciences is to closely track all units of measurement when performing calculations of physical quantities. This practice is generally referred to as *dimensional analysis*. A brief example of dimensional analysis is shown here, used to analyze the simple formula $P = IV$ which describes the amount of power dissipated by an electrical load (P) given its current (I) and voltage drop (V):

$$P = IV$$

Substituting units of measurement for each variable in this formula (i.e. Watts for power, Amperes for current, and Volts for voltage), using bracket symbols to denote these as unit abbreviations rather than variables, we get this result:

$$[\text{Watts}] = [\text{Amperes}] \times [\text{Volts}] \quad \text{or} \quad [\text{W}] = [\text{A}][\text{V}]$$

If we happen to know that “watts” is equivalent to joules of energy dissipated per second, and that “amperes” is equivalent to coulombs of charge motion per second, and that “volts” is equivalent to joules of energy per coulomb of electrical charge, we may substitute these units of measurement into the formula and see that the unit of “coulomb” cancels just like identical variables in the numerator and denominator of multiplied fractions:

$$\left[\frac{\text{Joules}}{\text{Seconds}} \right] = \left[\frac{\text{Coulombs}}{\text{Seconds}} \right] \times \left[\frac{\text{Joules}}{\text{Coulombs}} \right] \quad \text{or} \quad \left[\frac{\text{J}}{\text{s}} \right] = \left[\frac{\text{C}}{\text{s}} \right] \left[\frac{\text{J}}{\text{C}} \right]$$

As it so happens, dimensional analysis may be employed in a similar manner to convert between different units of measurement via a technique I like to call *unity fractions*.

This technique involves setting up the original quantity as a fraction, then multiplying by a series of fractions having *physical* values of unity (1) so that by multiplication the original value does not change, but the units do. Let’s take for example the conversion of quarts into gallons, an example of a fluid volume conversion:

$$35 \text{ qt} = ??? \text{ gal}$$

Now, most people know there are four quarts in one gallon, and so it is tempting to simply divide the number 35 by four to arrive at the proper number of gallons. However, the purpose of this example is to show you how the technique of unity fractions works, not to get an answer to a problem.

To demonstrate the unity fraction technique, we will first write the original quantity as a fraction, in this case a fraction with 1 as the denominator:

$$\frac{35 \text{ qt}}{1}$$

Next, we will multiply this fraction by another fraction having a *physical* value of unity (1) so that we do not alter² the quantity. This means a fraction comprised of equal measures in the numerator and denominator, but having different units of measurement. This “unity” fraction must be arranged in such a way that the undesired unit cancels out and leaves only the desired unit(s) in the product. In this particular example, we wish to cancel out quarts and end up with gallons, so we must arrange a fraction consisting of quarts and gallons having equal quantities in numerator and denominator, such that quarts will cancel and gallons will remain:

$$\left(\frac{35 \text{ qt}}{1}\right) \left(\frac{1 \text{ gal}}{4 \text{ qt}}\right)$$

Now we see how the unit of “quarts” cancels from the numerator of the first fraction and the denominator of the second (“unity”) fraction, leaving only the unit of “gallons” left standing:

$$\left(\frac{35 \cancel{\text{qt}}}{1}\right) \left(\frac{1 \text{ gal}}{4 \cancel{\text{qt}}}\right) = 8.75 \text{ gal}$$

The reason this conversion technique is so powerful is it allows one to perform the largest range of unit conversions while memorizing the smallest possible set of conversion factors.

Here is a set of six equal volumes, each one expressed in a different unit of measurement:

$$1 \text{ gallon (gal)} = 231.0 \text{ cubic inches (in}^3\text{)} = 4 \text{ quarts (qt)} = 8 \text{ pints (pt)} = 128 \text{ fluid ounces (fl. oz.)} \\ = 3.7854 \text{ liters (l)}$$

Since all six of these quantities are physically equal, it is possible to build a “unity fraction” out of any two, to use in converting any of the represented volume units into any of the other represented volume units. Shown here are a few different volume unit conversion problems, using unity fractions built only from these factors (all canceled units shown using strike-out lines):

40 gallons converted into fluid ounces (using $128 \text{ fl. oz.} = 1 \text{ gal}$ in the unity fraction):

$$\left(\frac{40 \cancel{\text{gal}}}{1}\right) \left(\frac{128 \text{ fl. oz.}}{1 \cancel{\text{gal}}}\right) = 5120 \text{ fl. oz}$$

5.5 pints converted into cubic inches (using $231 \text{ in}^3 = 8 \text{ pt}$ in the unity fraction):

$$\left(\frac{5.5 \cancel{\text{pt}}}{1}\right) \left(\frac{231 \text{ in}^3}{8 \cancel{\text{pt}}}\right) = 158.8 \text{ in}^3$$

²A basic mathematical identity is that multiplication of any quantity by 1 does not change the value of that original quantity. If we multiply some quantity by a fraction having a physical value of 1, no matter how strange-looking that fraction may appear, the value of the original quantity will be left intact. The goal here is to judiciously choose a fraction with a physical value of 1 but with its units of measurement so arranged that we cancel out the original quantity’s unit(s) and replace them with the units we desire.

1170 liters converted into quarts:

$$\left(\frac{1170 \cancel{\text{l}}}{1}\right) \left(\frac{4 \text{ qt}}{3.7854 \cancel{\text{l}}}\right) = 1236 \text{ qt}$$

By contrast, if we were to try to memorize a 6×6 table giving conversion factors between *any two* of six volume units, we would have to commit 30 different conversion factors to memory! Clearly, the ability to set up “unity fractions” is a much more memory-efficient and practical approach.

This economy of conversion factors is very useful, and may also be extended to cases where linear units are raised to powers to represent two- or three-dimensional quantities. To illustrate, suppose we wished to convert 5.5 pints into cubic *feet* instead of cubic *inches*: with no conversion equivalence between pints and cubic feet included in our string of six equalities, what do we do?

We should know the equality between inches and feet: there are exactly 12 inches in 1 foot. This simple fact may be applied by incorporating *another* unity fraction in the original problem to convert cubic inches into cubic feet. We will begin by including another unity fraction comprised of 12 inches and 1 foot, just to see how this might work:

5.5 pints converted into cubic feet (*our first attempt!*):

$$\left(\frac{5.5 \cancel{\text{pt}}}{1}\right) \left(\frac{231 \text{ in}^3}{8 \cancel{\text{pt}}}\right) \left(\frac{1 \text{ ft}}{12 \cancel{\text{in}}}\right) = 13.23 \text{ in}^2 \cdot \text{ft}$$

Unfortunately, this yields a non-sensical unit of square inch-feet. Even though $\frac{1 \text{ ft}}{12 \text{ in}}$ is a valid unity fraction, it does not *completely* cancel out the unit of cubic inches in the numerator of the first unity fraction. Instead, the unit of “inches” in the denominator of the unity fraction merely cancels out one of the “inches” in the “cubic inches” of the previous fraction’s numerator, leaving square inches (in^2). What we need for full cancellation of cubic inches is a unity fraction relating *cubic* feet to *cubic* inches. We can get this, though, simply by *cubing* the $\frac{1 \text{ ft}}{12 \text{ in}}$ unity fraction:

5.5 pints converted into cubic feet (*our second attempt!*):

$$\left(\frac{5.5 \text{ pt}}{1}\right) \left(\frac{231 \text{ in}^3}{8 \text{ pt}}\right) \left(\frac{1 \text{ ft}}{12 \text{ in}}\right)^3$$

Distributing the third power to the interior terms of the last unity fraction:

$$\left(\frac{5.5 \text{ pt}}{1}\right) \left(\frac{231 \text{ in}^3}{8 \text{ pt}}\right) \left(\frac{1^3 \text{ ft}^3}{12^3 \text{ in}^3}\right)$$

Calculating the values of 1^3 and 12^3 inside the last unity fraction, then canceling units and solving:

$$\left(\frac{5.5 \cancel{\text{pt}}}{1}\right) \left(\frac{231 \cancel{\text{in}^3}}{8 \cancel{\text{pt}}}\right) \left(\frac{1 \text{ ft}^3}{1728 \cancel{\text{in}^3}}\right) = 0.0919 \text{ ft}^3$$

Now the answer makes sense: a volume expressed in units of cubic feet.

Once again, this unit conversion technique shows its power by minimizing the number of conversion factors we must memorize. We need not memorize how many cubic inches are in a cubic foot, or how many square inches are in a square foot, if we know how many linear inches are in a linear foot and we simply let the fractions “tell” us whether a power is needed for unit cancellation.

Unity fractions are also useful when we need to convert more than one unit in a given quantity. For example, suppose a flowmeter at a wastewater treatment facility gave us a flow measurement of 205 cubic feet per minute but we needed to convert this expression of water flow into units of cubic yards per day. Observe the following unit-fraction conversion to see how unity fractions serve the purpose of converting cubic feet into cubic yards, and minutes into days (by way of minutes to hours, and hours to days):

$$\left(\frac{205 \text{ ft}^3}{\text{min}}\right) \left(\frac{1^3 \text{ yd}^3}{3^3 \text{ ft}^3}\right) \left(\frac{60 \text{ min}}{1 \text{ hr}}\right) \left(\frac{24 \text{ hr}}{1 \text{ day}}\right) = 10933.3 \text{ yd}^3/\text{day}$$

Note how the only units left un-canceled on the left-hand side of the “equals” symbol are cubic yards (yd^3) and days, which therefore become the units of measurement for the final result.

A major caveat to this method of converting units is that the units must be *directly proportional* to one another, since this multiplicative conversion method is really nothing more than an exercise in mathematical proportions. Here are some examples (but not an exhaustive list!) of conversions that *cannot* be performed using the “unity fraction” method:

- Absolute / Gauge pressures, because one scale is *offset* from the other by 14.7 PSI (atmospheric pressure).
- Celsius / Fahrenheit, because one scale is *offset* from the other by 32 degrees.
- Wire diameter / gauge number, because gauge numbers grow smaller as wire diameter grows larger (inverse proportion rather than direct) and because there is no proportion relating the two.
- Power / decibels, because the relationship is logarithmic rather than proportional.

The following subsections give sets of physically equal quantities, which may be used to create unity fractions for unit conversion problems. Note that only those quantities shown in the same line (separated by = symbols) are truly equal to each other, not quantities appearing in different lines!

2.4.2 Conversion formulae for temperature

Note: all of the conversion factors given for temperature are *exact*, not approximations.

$$^{\circ}\text{F} = (^{\circ}\text{C})(9/5) + 32$$

$$^{\circ}\text{C} = (^{\circ}\text{F} - 32)(5/9)$$

$$^{\circ}\text{R} = ^{\circ}\text{F} + 459.67$$

$$\text{K} = ^{\circ}\text{C} + 273.15$$

2.4.3 Conversion factors for distance

Note: all of the conversion factors given for distance are *exact*, not approximations.

$$1 \text{ inch (in)} = \mathbf{2.54} \text{ centimeters (cm)}$$

$$1 \text{ foot (ft)} = \mathbf{12} \text{ inches (in)}$$

$$1 \text{ yard (yd)} = \mathbf{3} \text{ feet (ft)}$$

$$1 \text{ mile (mi)} = \mathbf{5280} \text{ feet (ft)}$$

2.4.4 Conversion factors for volume

Note: all conversion factors shown in **bold** type are *exact*, not approximations.

$$1 \text{ gallon (gal)} = 231.0 \text{ cubic inches (in}^3\text{)} = \mathbf{4} \text{ quarts (qt)} = \mathbf{8} \text{ pints (pt)} = \mathbf{16} \text{ cups} = \mathbf{128} \text{ fluid ounces (fl. oz.)} = 3.7854 \text{ liters (l)}$$

$$1 \text{ milliliter (ml)} = \mathbf{1} \text{ cubic centimeter (cm}^3\text{)}$$

2.4.5 Conversion factors for velocity

Note: all conversion factors shown in **bold** type are *exact*, not approximations.

$$1 \text{ mile per hour (mi/h)} = \mathbf{88} \text{ feet per minute (ft/m)} = 1.46667 \text{ feet per second (ft/s)} = 1.60934 \text{ kilometer per hour (km/h)} = 0.44704 \text{ meter per second (m/s)} = 0.868976 \text{ knot (knot - international)}$$

2.4.6 Conversion factors for mass

$$1 \text{ pound-mass (lbm)} = 0.4535924 \text{ kilogram (kg)} = 0.031081 \text{ slugs}$$

2.4.7 Conversion factors for force

$$1 \text{ pound-force (lbf)} = 4.448222 \text{ newtons (N)}$$

$$1 \text{ kilogram-force (kgf)} = \mathbf{9.80665} \text{ newtons (N)}$$

2.4.8 Conversion factors for area

Note: all conversion factors shown in **bold** type are *exact*, not approximations.

$$1 \text{ acre} = 43560 \text{ square feet (ft}^2\text{)} = 4840 \text{ square yards (yd}^2\text{)} = 4046.86 \text{ square meters (m}^2\text{)}$$

2.4.9 Conversion factors for pressure (either all gauge or all absolute)

Note: all conversion factors shown in **bold** type are *exact*, not approximations.

1 pounds per square inch (PSI) = 2.03602 inches of mercury at 0 °C (in. Hg) = 27.6799 inches of water at 4 °C (in. W.C.) = 6.894757 kilopascals (kPa) = 0.06894757 bar

1 bar = 100 kilopascals (kPa) = 14.504 pounds per square inch (PSI)

1 meter of water at 4 °C (m W.C.) = 9.80665 kilopascals (kPa)

2.4.10 Conversion factors for pressure (absolute pressure units only)

Note: all conversion factors shown in **bold** type are *exact*, not approximations.

1 standard atmosphere (Atm) = 14.7 pounds per square inch absolute (PSIA) = 101.325 kilopascals absolute (kPaA) = 1.01325 bar absolute = 760 millimeters of mercury absolute (mmHgA) = 760 torr (torr)

2.4.11 Conversion factors for energy or work

1 British thermal unit (Btu – “International Table”) = 251.996 calories (cal – “International Table”) = 1055.06 joules (J) = 1055.06 watt-seconds (W-s) = 0.293071 watt-hour (W-hr) = 1.05506×10^{10} ergs (erg) = 778.169 foot-pound-force (ft-lbf)

2.4.12 Conversion factors for power

Note: all conversion factors shown in **bold** type are *exact*, not approximations.

1 horsepower = 550 foot-pounds per second (ft-lbf/s) = 745.7 watts (W) = 2544.43 British thermal units per hour (Btu/h) = 0.0760181 boiler horsepower (hp – boiler)

2.4.13 Terrestrial constants

Acceleration of gravity at sea level = 9.806650 meters per second per second (m/s^2) = 32.1740 feet per second per second (ft/s^2)

Atmospheric pressure = 14.7 pounds per square inch absolute (PSIA) = 760 millimeters of mercury absolute (mmHgA) = 760 torr (torr) = 1.01325 bar (bar)

Atmospheric gas concentrations (by volume, not mass):

- Nitrogen = 78.084 %
- Oxygen = 20.946 %
- Argon = 0.934 %
- Carbon Dioxide (CO_2) = 0.033 %
- Neon = 18.18 ppm
- Helium = 5.24 ppm
- Methane (CH_4) = 2 ppm
- Krypton = 1.14 ppm
- Hydrogen = 0.5 ppm
- Nitrous Oxide (N_2O) = 0.5 ppm
- Xenon = 0.087 ppm

Density of dry air at 20 °C and 760 torr = 1.204 mg/cm^3 = 1.204 kg/m^3 = 0.075 lb/ft^3 = 0.00235 slugs/ ft^3

Absolute viscosity of dry air at 20 °C and 760 torr = 0.018 centipoise (cp) = 1.8×10^{-5} pascal-seconds (Pa·s)

2.4.14 Properties of water

Freezing point at sea level = 32 °F = 0 °C

Boiling point at sea level = 212 °F = 100 °C

Density of water at 4 °C = 1000 kg/m^3 = 1 g/cm^3 = 1 kg/liter = 62.428 lb/ft^3 = 1.94 slugs/ ft^3

Specific heat of water at 14 °C = 1.00002 $\text{calories/g}\cdot^\circ\text{C}$ = 1 $\text{BTU/lb}\cdot^\circ\text{F}$ = 4.1869 $\text{joules/g}\cdot^\circ\text{C}$

Specific heat of ice \approx 0.5 $\text{calories/g}\cdot^\circ\text{C}$

Specific heat of steam \approx 0.48 $\text{calories/g}\cdot^\circ\text{C}$

Absolute viscosity of water at 20 °C = 1.0019 centipoise (cp) = 0.0010019 pascal-seconds (Pa·s)

Surface tension of water (in contact with air) at 18 °C = 73.05 dynes/cm

pH of pure water at 25 °C = 7.0 (*pH scale = 0 to 14*)

2.4.15 Miscellaneous physical constants

Note: all constants shown in **bold** type are *exact*, not approximations. Parentheses show one standard deviation (σ) of uncertainty in the last digits: for example, Avogadro's number given as $6.02214179(30) \times 10^{23}$ means the center value ($6.02214179 \times 10^{23}$) plus or minus $0.00000030 \times 10^{23}$.

Avogadro's number (N_A) = $6.02214179(30) \times 10^{23}$ per mole (mol^{-1})

Boltzmann's constant (k) = $1.3806504(24) \times 10^{-23}$ joules per Kelvin (J/K)

Electronic charge (e) = $1.602176487(40) \times 10^{-19}$ Coulomb (C)

Faraday constant (F) = $9.64853399(24) \times 10^4$ Coulombs per mole (C/mol)

Gravitational constant (G) = $6.67428(67) \times 10^{-11}$ cubic meters per kilogram-seconds squared ($\text{m}^3/\text{kg}\cdot\text{s}^2$)

Molar gas constant (R) = $8.314472(15)$ joules per mole-Kelvin (J/mol-K) = $0.08205746(14)$ liters-atmospheres per mole-Kelvin

Planck constant (h) = $6.62606896(33) \times 10^{-34}$ joule-seconds (J-s)

Stefan-Boltzmann constant (σ) = $5.670400(40) \times 10^{-8}$ Watts per square meter-Kelvin⁴ ($\text{W}/\text{m}^2\cdot\text{K}^4$)

Speed of light in a vacuum (c) = **299792458 meters per second** (m/s) = 186282.4 miles per second (mi/s)

All constants taken from NIST data "Fundamental Physical Constants – Extensive Listing", published 2006.

2.4.16 Weight densities of common materials

All density figures approximate for samples at standard temperature and pressure³.

Liquids:

- Acetone: $\gamma = 49.4 \text{ lb/ft}^3$
- Alcohol, ethyl (ethanol): $\gamma = 49.4 \text{ lb/ft}^3$
- Alcohol, methyl (methanol): $\gamma = 50.5 \text{ lb/ft}^3$
- Benzene: $\gamma = 56.1 \text{ lb/ft}^3$
- Butane (liquid): $\gamma = 36.1 \text{ lb/ft}^3$
- Carbon disulfide: $\gamma = 80.7 \text{ lb/ft}^3$
- Carbon tetrachloride: $\gamma = 99.6 \text{ lb/ft}^3$
- Chloroform: $\gamma = 93 \text{ lb/ft}^3$
- Ethylene glycol (ethanediol): $\gamma = 69.22 \text{ lb/ft}^3$
- Gasoline: $\gamma = 41 \text{ lb/ft}^3$ to 43 lb/ft^3
- Glycerin: $\gamma = 78.6 \text{ lb/ft}^3$
- Isobutane (liquid): $\gamma = 34.8 \text{ lb/ft}^3$
- Kerosene: $\gamma = 51.2 \text{ lb/ft}^3$
- Mercury: $\gamma = 849 \text{ lb/ft}^3$
- Methanol (methyl alcohol): $\gamma = 50.5 \text{ lb/ft}^3$
- Milk: $\gamma = 64.2 \text{ lb/ft}^3$ to 64.6 lb/ft^3
- Naphtha, petroleum: $\gamma = 41.5 \text{ lb/ft}^3$
- Oil, castor: $\gamma = 60.5 \text{ lb/ft}^3$
- Oil, coconut: $\gamma = 57.7 \text{ lb/ft}^3$
- Oil, linseed (boiled): $\gamma = 58.8 \text{ lb/ft}^3$
- Oil, olive: $\gamma = 57.3 \text{ lb/ft}^3$
- Propane (liquid): $\gamma = 31.2 \text{ lb/ft}^3$
- Toluene: $\gamma = 54.1 \text{ lb/ft}^3$

³Density figures taken or derived from tables in the *CRC Handbook of Chemistry and Physics*, 64th Edition. Most liquid densities taken from table on page F-3 and solid densities taken from table on page F-1. Some liquid densities taken from tables on pages E-27 through E-31. All temperatures at or near 20 °C.

- Turpentine: $\gamma = 54.3 \text{ lb/ft}^3$
- Water, heavy: $\gamma = 68.97 \text{ lb/ft}^3$
- Water, light (normal): $\gamma = 62.4 \text{ lb/ft}^3$
- Water, sea: $\gamma = 63.99 \text{ lb/ft}^3$

Solids:

- Beryllium: $\gamma = 115.37 \text{ lb/ft}^3$
- Brass: $\gamma = 524.4 \text{ lb/ft}^3$
- Calcium: $\gamma = 96.763 \text{ lb/ft}^3$
- Carbon (diamond): $\gamma = 196.65 \text{ lb/ft}^3$ to 220.37 lb/ft^3
- Cement (set): $\gamma = 170 \text{ lb/ft}^3$ to 190 lb/ft^3
- Chromium: $\gamma = 448.86 \text{ lb/ft}^3$
- Copper: $\gamma = 559.36 \text{ lb/ft}^3$
- Cork: $\gamma = 14 \text{ lb/ft}^3$ to 16 lb/ft^3
- Gold: $\gamma = 1178.6 \text{ lb/ft}^3$
- Ice: $\gamma = 57.2 \text{ lb/ft}^3$
- Iron: $\gamma = 490.68 \text{ lb/ft}^3$
- Ivory: $\gamma = 114 \text{ lb/ft}^3$ to 120 lb/ft^3
- Lead: $\gamma = 708.56 \text{ lb/ft}^3$
- Leather: $\gamma = 54 \text{ lb/ft}^3$
- Magnesium: $\gamma = 108.50 \text{ lb/ft}^3$
- Molybdenum: $\gamma = 638.01 \text{ lb/ft}^3$
- Quartz: $\gamma = 165 \text{ lb/ft}^3$
- Rubber (soft): $\gamma = 69 \text{ lb/ft}^3$
- Rubber (hard): $\gamma = 74 \text{ lb/ft}^3$
- Salt, rock: $\gamma = 136 \text{ lb/ft}^3$
- Sugar: $\gamma = 99 \text{ lb/ft}^3$
- Tar: $\gamma = 66 \text{ lb/ft}^3$
- Wood, balsa: $\gamma = 7 \text{ lb/ft}^3$ to 9 lb/ft^3
- Wood, maple: $\gamma = 39 \text{ lb/ft}^3$ to 47 lb/ft^3

2.5 Dimensional analysis

An interesting parallel to the “unity fraction” unit conversion technique is something referred to in physics as *dimensional analysis*. Performing dimensional analysis on a physics formula means to set it up with units of measurement in place of variables, to see how units cancel and combine to form the appropriate unit(s) of measurement for the result.

For example, let’s take the familiar power formula used to calculate power in a simple DC electric circuit:

$$P = IV$$

Where,

P = Power (watts)

I = Current (amperes)

V = Voltage (volts)

Each of the units of measurement in the above formula (watt, ampere, volt) are actually comprised of more fundamental physical units. One watt of power is one joule of energy transferred per second. One ampere of current is one coulomb of electric charge moving by per second. One volt of potential is one joule of energy per coulomb of electric charge. When we write the equation showing these units in their proper orientations, we see that the result (power in watts, or joules per second) actually does agree with the units for amperes and volts because the unit of electric charge (coulombs) cancels out. In dimensional analysis we customarily distinguish unit symbols from variables by using non-italicized letters and surrounding each one with square brackets:

$$P = IV$$

$$[\text{Watts}] = [\text{Amperes}] \times [\text{Volts}] \quad \text{or} \quad [\text{W}] = [\text{A}][\text{V}]$$

$$\left[\frac{\text{Joules}}{\text{Seconds}} \right] = \left[\frac{\text{Coulombs}}{\text{Seconds}} \right] \times \left[\frac{\text{Joules}}{\text{Coulombs}} \right] \quad \text{or} \quad \left[\frac{\text{J}}{\text{s}} \right] = \left[\frac{\text{C}}{\text{s}} \right] \left[\frac{\text{J}}{\text{C}} \right]$$

Dimensional analysis gives us a way to “check our work” when setting up new formulae for physics- and chemistry-type problems.

2.6 The International System of Units

The very purpose of physics is to quantitatively describe and explain the physical world in as few terms as possible. This principle extends to units of measurement as well, which is why we usually find different units used in science actually defined in terms of more fundamental units. The *watt*, for example, is one joule of energy transferred per second of time. The joule, in turn, is defined in terms of three base units, the kilogram, the meter, and the second:

$$[J] = \frac{[\text{kg}][\text{m}^2]}{[\text{s}^2]}$$

Within the metric system of measurements, an international standard exists for which units are considered fundamental and which are considered “derived” from the fundamental units. The modern standard is called *SI*, which stands for *Système International*. This standard recognizes seven fundamental, or *base* units, from which all others are derived⁴:

Physical quantity	SI unit	SI symbol
Length	meter	m
Mass	kilogram	kg
Time	second	s
Electric current	ampere	A
Temperature	kelvin	K
Amount of substance	mole	mol
Luminous intensity	candela	cd

An older standard existed for base units, in which the *centimeter*, *gram*, and *second* comprised the first three base units. This standard is referred to as the *cgs* system, in contrast to the SI system⁵. You will still encounter some derived cgs units used in instrumentation, including the *poise* and the *stokes* (both used to express fluid viscosity). Then of course we have the *British engineering system* which uses such wonderful⁶ units as feet, pounds, and (thankfully) seconds. Despite the fact that the majority of the world uses the metric (SI) system for weights and measures, the British system is sometimes referred to as the *Customary* system.

⁴The only exception to this rule being units of measurement for angles, over which there has not yet been full agreement whether the unit of the *radian* (and its solid counterpart, the *steradian*) is a base unit or a derived unit.

⁵The older name for the SI system was “MKS,” representing meters, kilograms, and seconds.

⁶I’m noting my sarcasm here, just in case you are immune to my odd sense of humor.

2.7 Conservation Laws

The *Law of Mass Conservation* states that matter can neither be created nor destroyed. The *Law of Energy Conservation* states that energy can neither be created nor destroyed. However, both mass and energy may change forms, and even change into one another in the case of nuclear phenomena.

Conversion of mass into energy, or of energy into mass, is quantitatively described by Albert Einstein's famous equation:

$$E = mc^2$$

Where,

E = Energy (joules)

m = Mass (kilograms)

c = Speed of light (approximately 3×10^8 meters per second)

Conservation laws find practical context in many areas of science and life, but in the realm of process control we have the principles of *mass balance* and *energy balance* which are direct expressions of these Laws. "Mass balance" refers to the fact that the sum total of mass entering a process must equal the sum total of mass exiting the process, provided the process is in a steady-state condition (all variables remaining constant over time). To give a simple example of this, the mass flow rate of fluid entering a pipe *must* be equal to the mass flow rate of fluid exiting the pipe, provided the pipe is neither accumulating nor releasing mass within its internal volume. "Energy balance" is a parallel concept, stating that the sum total of energy entering a process must equal the sum total of energy exiting a process, provided a steady-state condition (no energy being stored or released from storage within the process).

2.8 Classical mechanics

Classical mechanics (often called *Newtonian* mechanics in honor of Isaac Newton) deal with forces and motions of objects in common circumstances. The vast majority of instrumentation applications deals with this realm of physics. Two other areas of physics, *relativistic* and *quantum*, will not be covered in this chapter because their domains lie outside the typical experience of industrial instrumentation⁷.

⁷Relativistic physics deals with phenomena arising as objects travel near the speed of light. Quantum physics deals with phenomena at the atomic level. Neither is germane to the vast majority of industrial instrument applications.

2.8.1 Newton's Laws of Motion

These laws were formulated by the great mathematician and physicist Isaac Newton (1642-1727). Much of Newton's thought was inspired by the work of an individual who died the same year Newton was born, Galileo Galilei (1564-1642).

1. An object at rest tends to stay at rest; an object in motion tends to stay in motion
2. The acceleration of an object is directly proportional to the net force acting upon it and inversely proportional to the object's mass
3. Forces between objects always exist in equal and opposite pairs

Newton's first law may be thought of as the *law of inertia*, because it describes the property of inertia that all objects having mass exhibit: resistance to change in velocity. This law is quite counter-intuitive for many people, who tend to believe that objects require continual force to keep moving. While this is true for objects experiencing friction, it is not for ideal (frictionless) motion. This is why satellites and other objects in space continue to travel with no mode of propulsion: they simply "coast" indefinitely on their own inertia because there is no friction in space to dissipate their kinetic energy and slow them down.

Newton's second law is the verbal equivalent of the force/mass/acceleration formula: $F = ma$. This law elaborates on the first, in that it mathematically relates force and motion in a very precise way. For a frictionless object, the *change in velocity* (i.e. its *acceleration*) is proportional to force. This is why a frictionless object may continue to move without any force applied: once moving, force would only be necessary for continued acceleration. If zero force is applied, the acceleration will likewise be zero, and the object will maintain its velocity indefinitely (again, assuming no friction at work).

Newton's third law describes how forces always exist in *pairs* between two objects. The rotating blades of a helicopter, for example, exert a downward force on the air (accelerating the air), but the air in turn exerts an upward force on the helicopter (suspending it in flight). A spider hanging on the end of a thread exerts a downward force (weight) on the thread, while the thread exerts an upward force of equal magnitude on the spider (tension). Force pairs are always equal in magnitude but opposite in direction.

2.8.2 Work, energy, and power

Two very fundamental and closely-related concepts in physics are *work* and *energy*. *Work* is simply what happens when any force acts through a parallel motion, such as when a weight is lifted against gravity or when a spring is compressed. *Energy* is a more abstract concept and therefore more difficult to define. One definition⁸ of energy is “that which permits or results in motion,” where the word “motion” is used in a very broad sense including even the motion of individual atoms within a substance. Energy exists in many different forms, and may be transferred between objects and/or converted from one form to another, but cannot be created from nothing or be destroyed and turned into nothing (this is the *Law of Energy Conservation*). *Power* is the rate at which work is done, or alternatively the rate at which energy transfer occurs.

First, just a little bit of math. Work (W) is mathematically defined as the dot-product of force (F) and displacement (x) vectors⁹, written as follows:

$$W = \vec{F} \cdot \vec{x}$$

Where,

W = Work, in newton-meters (metric) or foot-pounds (British)

\vec{F} = Force vector (force and direction exerted) doing the work, in newtons (metric) or pounds (British)

\vec{x} = Displacement vector (distance and direction traveled) over which the work was done, in meters (metric) or feet (British)

⁸A common definition of energy is the “ability to do work” which is not always true. There are some forms of energy which may *not* be harnessed to do work, such as the thermal motion of molecules in an environment where all objects are at the same temperature. Energy that has the ability to do work is more specifically referred to as *exergy*. While energy is always conserved (i.e. never lost, never gained), exergy is a quantity that can never be gained but can be lost. The inevitable loss of exergy is closely related to the concept of *entropy*, where energy naturally diffuses into less useful (more homogeneous) forms over time. This important concept explains why no machine can never be perfectly (100.0%) efficient, among other things.




⁹A *vector* is a mathematical quantity possessing both a magnitude and a direction. Force (F), displacement (x), and velocity (v) are all vector quantities. Some physical quantities such as temperature (T), work (W), and energy (E) possess magnitude but no direction. We call these directionless quantities “scalar.” It would make no sense at all to speak of a temperature being “79 degrees Celsius due North” whereas it *would* make sense to speak of a force being “79 Newtons due North”. Physicists commonly use a little arrow symbol over the variable letter to represent that variable as a vector, when both magnitude and direction matter. Thus \vec{F} represents a force vector with both magnitude and direction specified, while plain F merely represents the magnitude of that force without a specified direction. A “dot-product” is one way in which vectors may be multiplied, and the result of a dot-product is always a scalar quantity.

The fact that both force and displacement appear as *vectors* tells us their relative directions are significant to the calculation of work. If the force and displacement vectors point in exactly the same direction, work is simply the product of F and x magnitudes ($W = Fx$). If the force and displacement vectors point in opposite directions, work is the *negative* product ($W = -Fx$). Another way to express the calculation of work is as the product of the force and displacement magnitudes (F and x) and the cosine of the angle separating the two force vectors ($\cos \theta$):


$$W = Fx \cos \theta$$

When the two vectors \vec{F} and \vec{x} point the same direction, the angle θ between them is zero and therefore $W = Fx$ because $\cos 0^\circ = 1$. When the two vectors point in opposite directions, the angle between them is 180° and therefore $W = -Fx$ because $\cos 180^\circ = -1$.

This means if a force acts in the same direction as a motion (i.e. force and displacement vectors pointing in the same direction), the work being done by that force will be a *positive* quantity. If a force acts in the opposite direction of a motion (i.e. force and displacement vectors pointing in opposite directions), the work done by that force will be a *negative* quantity. If a force acts perpendicularly to the direction of a motion (i.e. force and displacement vectors at right angles to each other), that force will do *zero* work.



Illustrations are helpful in explaining these concepts. Consider a crane hoisting a 2380 pound weight 15 feet up into the air by means of an attached cable:



Crane does work on weight $W_{crane} = (2380 \text{ lb})(15 \text{ ft})(\cos 0^\circ) = +35700 \text{ ft-lbs}$
Work is done on the weight $W_{weight} = (2380 \text{ lb})(15 \text{ ft})(\cos 180^\circ) = -35700 \text{ ft-lbs}$


The amount of work done from the crane's perspective is $W_{crane} = +35700 \text{ ft-lbs}$, since the crane's force ($F_{crane} = 2380 \text{ lbs}$, up) points in the same upward direction as the cable's motion ($x = 15 \text{ ft}$, up) and therefore there is no angular difference between the crane force and motion vectors (i.e. $\theta = 0$). The amount of work done from weight's perspective, however, is $W_{weight} = -35700 \text{ ft-lbs}$ because the weight's force vector ($F_{weight} = 2380 \text{ lbs}$, down) points in the *opposite* direction as the cable's motion vector ($x = 15 \text{ feet}$, up), yielding an angular difference of $\theta = 180^\circ$. Another way of expressing these two work values is to state the crane's work in the active voice and the weight's work in the passive voice: *the crane did 35700 ft-lbs of work, while 35700 ft-lbs of work was done on the weight*. This language is truly appropriate, as the crane is indeed the *active* agent in this scenario, while the weight *passively opposes* the crane's efforts. In other words, the crane is the *motive source* of the work, while the weight is a *load*.

Now, how does *energy* fit into this illustration? Certainly we see that motion occurred, and therefore energy must have been involved. The energy lifting the 2380 pound weight 15 feet upward didn't come from nowhere – it must have been present somewhere in the universe prior to the weight's ascension if the Law of Energy Conservation is indeed true. If the crane is powered by an internal combustion engine, the energy came from the fuel stored in the crane's fuel tank. If the crane is electric, the energy to hoist the weight came from a battery or from an electrical generator somewhere sending power to the crane through an electric cable. The act of lifting the weight was actually an act of *energy transfer* from the crane's energy source, through the crane motor and cable, and finally to the weight. Along the way, some of the initial energy was also converted into heat through inefficiencies in the crane's motor and mechanism. However, if we were to calculate the sum of all the energy transferred to the lifted weight plus all energy "dissipated" in the form of heat, that total must be precisely equal to the initial energy of the fuel (or electricity) used by the crane in doing this lift. In other words, the sum total of all energy in the universe is the same after the lift as it was before the lift.

Power applies in this scenario to how *quickly* the weight rises. So far, all we know about the weight's lifting is that it took 35700 ft-lbs of energy to do that work. If we knew how long it took

the crane to do that work, we could calculate the crane's power output. For example, a crane with a power rating of 35700 ft-lbs per second could complete this 15-foot lift in only one second of time. Likewise, a crane with a power rating of only 3570 ft-lbs per second would require 10 seconds of time to execute the same lift.

An interesting thing happens if the crane moves sideways along the ground after lifting the weight 15 feet up into the air. No work is being done by the crane on the weight, or by the weight on the crane, because the displacement vector is now perpendicular ($\theta = 90^\circ$) to both force vectors:




No work done on or by the crane $W_{crane} = (2380 \text{ lb})(15 \text{ ft})(\cos 90^\circ) = 0 \text{ ft-lbs}$

No work done on or by the weight $W_{weight} = (2380 \text{ lb})(15 \text{ ft})(\cos -90^\circ) = 0 \text{ ft-lbs}$

It should be noted that the crane's engine *will* do work as it overcomes rolling friction in the wheels to move the crane along, *but this is not work done on or by the hoisted weight*. When we calculate work – as with all other calculations in physics – we must be very careful to keep in mind where the calculation(s) apply in the scenario. Here, the forces and displacement with regard to the hoisted weight are perpendicular to each other, and therefore no work is being done there. The only work done anywhere in this system as the crane rolls 15 feet horizontally involves the horizontal force required to roll the crane, which is unspecified in this illustration.

Similarly, there is no transfer of energy to or from the hoisted weight while the crane rolls along. Whatever energy comes through the crane's engine only goes into overcoming rolling friction at the wheels and ground, not to do any work with the weight itself. Generally this will be a very small amount of energy compared to the energy required to hoist a heavy load.

A good question to ask after hoisting the weight is “Where did that 35700 ft-lbs of energy go after the lift was complete?” The Law of Energy Conservation tells us that energy cannot be destroyed, and so the 35700 ft-lbs of work must be accounted for somehow. In this case, the energy is now *stored* in the weight where it may be released at some later time. We may demonstrate this fact by slowly lowering the weight back down to the ground and watching the energy transfer:



Work is done on the crane $W_{crane} = (2380 \text{ lb})(15 \text{ ft})(\cos 180^\circ) = -35700 \text{ ft-lbs}$
Weight does work on crane $W_{weight} = (2380 \text{ lb})(15 \text{ ft})(\cos 0^\circ) = +35700 \text{ ft-lbs}$

Notice how the weight is now the actively-working object in the system, doing work on the crane. The crane is now the passive element, opposing the work being done by the weight. Both the crane and the weight are still pulling the same directions as before (crane pulling up, weight pulling down), but now the direction of displacement is going down which means the weight is “winning” and therefore doing the work, while the crane is “losing” and opposing the work.

If we examine what is happening inside the crane as the weight descends, we see that energy is being transferred from the descending weight to the crane. In most cranes, the descent of a load is controlled by a *brake* mechanism, regulating the speed of descent by applying friction to the cable’s motion. This brake friction generates a great deal of heat, which is a form of energy transfer: energy stored in the elevated weight is now being converted into heat which exits the crane in the form of hot air (air whose molecules are now vibrating at a faster speed than they were at their previous temperature). If the crane is electric, we have the option of *regenerative braking* where we recapture this energy instead of dissipating it in the form of heat. This works by switching the crane’s electric motor into an electric generator on demand, so the weight’s descent turns the motor/generator shaft to generate electricity to re-charge the crane’s battery or be injected back into the electric power grid to do useful work elsewhere.

Referring back to the illustration of the crane hoisting the weight, it is clear that the weight stored energy while it was being lifted up by the crane, and released this energy back to the crane while it was being lowered down to the ground. The energy held by the elevated weight may therefore be characterized as *potential energy*, since it had the *potential* to do work even if no work was being done by that energy at that moment.

A special version of the general work formula $W = \vec{F} \cdot \vec{x}$ exists for calculating this *gravitational potential energy*. Rather than express force and displacement as vectors with arbitrary directions, we express the weight of the object as the product of its mass and the acceleration of gravity ($F = mg$) and the vertical displacement of the object simply as its height above the ground ($x = h$). The amount of potential energy stored in the lift is simply equal to the work done during the lift ($W = E_P$):

$$E_p = mgh$$

Where,

E_p = Gravitational potential energy in newton-meters (metric) or foot-pounds (British)

m = Mass of object in kilograms (metric) or slugs (British)


g = Acceleration of gravity in meters per second squared (metric) or feet per second squared (British)

h = Height of lift in meters (metric) or feet (British)

There is no need for vectors or cosine functions in the $E_p = mgh$ formula, because gravity and height are always guaranteed to act along the same axis. A positive height (i.e. *above* ground level) is assumed to yield a positive potential energy value for the elevated mass.

Many different forms of potential energy exist, although the standard “textbook” example is of a mass lifted to some height against the force of gravity. Compressed (or stretched) springs have potential energy, as do pressurized fluids, chemical bonds (e.g. fuel molecules prior to combustion), hot masses, electrically-charged capacitors, and magnetized inductors. Any form of energy with the potential to be released into a different form at some later time is, by definition, *potential* energy.

An important application of work, energy, and power is found in *rotational* motion. Consider the application of a truck towing some load by a rope. The truck exerts a linear (i.e. straight-line) pulling force on the load being towed, but it must do so by applying a *rotary* (i.e. turning) force through the axle shaft powering the drive wheels:



Clearly, the truck *does work* on the load by exerting a pulling force in the direction of motion. If we wish to quantify this work, we may consider the work done by the truck as it tows the load a distance of 40 feet:

$$W_{truck} = Fx \cos \theta$$

$$W_{truck} = (1200 \text{ lb})(40 \text{ ft}) \cos 0^\circ$$

$$W_{truck} = 48000 \text{ lb-ft}$$

It should also be clear that *work is being done on* the load, since the load's force on the tow rope points in the opposite direction of its motion:

$$W_{load} = Fx \cos \theta$$

$$W_{load} = (1200 \text{ lb})(40 \text{ ft}) \cos 180^\circ$$

$$W_{load} = -48000 \text{ lb-ft}$$

For the sake of this example, it matters not what happens to the energy delivered to the load. Perhaps it gets converted to heat through the mechanism of friction at the load (e.g. friction from road contact, friction from wind resistance), perhaps it gets converted into potential energy in the case of the road inclining to a greater altitude, or perhaps (most likely) it is some combination of all these.

In order to quantify this amount of work from the perspective of the truck's rotating wheel, we must cast the variables of pulling force and pulling distance into rotary terms. We will begin by first examining the distance traveled by the wheel. Any circular wheel has a radius, and the wheel must turn a certain number of revolutions in order to pull the load any distance. The obvious function of a wheel is to convert between linear and rotary motion, and the common measure between these two motions is the *circumference* of the wheel: each revolution of the wheel equates directly to one circumference's worth of linear travel. In our truck example, the wheel has a radius of 1.5 feet, which means it must have a circumference of 9.425 feet (i.e. $2 \times \pi \times 1.5$ feet, since $C = \pi D = 2\pi r$). Therefore, in order to travel 40 linear feet this wheel must rotate 4.244 times (i.e. $40 \text{ feet} \div 9.425 \text{ feet/revolution} = 4.244$ revolutions).

Next, we must relate pulling force to rotational force. The 1200 pounds of pulling force exerted by the wheel¹⁰ originates from the twisting force exerted by the axle at the wheel's center. The technical term for this twisting force is *torque* (symbolized by the Greek letter "tau", τ), and it is a function of both the linear pulling force and the wheel's radius:

$$\tau = rF$$

Solving for torque (τ) in this application, we calculate 1800 pound-feet given the wheel's linear pulling force of 1200 pounds and a radius of 1.5 feet. Please note that the unit of "pound-feet" for torque is *not* the same as the unit of "foot-pounds" for work. Work is the product of force and displacement (distance of motion), while torque is the product of force and radius length. Work requires motion, while torque does not¹¹.

The torque value of 1800 lb-ft and the turning of 4.244 revolutions should be sufficient to calculate work done by the wheel, since torque equates directly to pulling force and the number of revolutions equates directly to pulling distance. It should be obvious that the product of 1800 lb-ft and 4.244 revolutions does not, however, equal 48000 ft-lbs of work, and so our torque-revolutions-work formula must contain an additional multiplication factor k :

$$W = k\tau x$$

Substituting 48000 ft-lbs for work (W), 1800 lb-ft for torque (τ), and 4.244 for the number of revolutions (x), we find that k must be equal to 6.283, or 2π . Knowing the value for k , we may re-write our rotational work formula more precisely:

$$W = 2\pi\tau x$$


This formula is correct no matter the wheel's size. A larger-radius wheel will certainly travel farther for each revolution, but that larger radius will proportionately reduce the pulling force for any given amount of torque, resulting in an unchanged work value.

¹⁰Note that this calculation will assume all the work of towing this load is being performed by a *single* wheel on the truck. Most likely this will not be the case, as most towing vehicles have multiple driving wheels (at least two). However, we will perform calculations for a single wheel in order to simplify the problem.

¹¹Consider the example of applying torque to a stubbornly seized bolt using a wrench: the force applied to the wrench multiplied by the radius length from the bolt's center to the perpendicular line of force yields torque, but absolutely no work is done on the bolt until the bolt begins to move (turn).

Applying this work formula to another application, let us consider an electric *winch* where a load is lifted against the force of gravity. A “winch” is a mechanism comprised of a tubular drum and cable, the cable wrapping or unwrapping around the drum as the drum is turned by an electric motor. For comparison we will use the same weight and lifting distance of the crane example, merely replacing the crane with a winch located atop a tower. In this case we will make the winch drum 3 inches in diameter, and calculate both the torque required to rotate the drum as well as the number of rotations necessary to lift the weight 15 feet:

Winch lifts weight 15 feet up



From the crane example we already know the amount of work which must be done on the weight to hoist it 15 feet vertically: $35700 \text{ ft}\cdot\text{lbs}$. Our rotational work formula ($W = 2\pi\tau x$) contains *two* unknowns at this point in time, since we know the value of W but not τ or x . This means we cannot yet solve for either τ or x using this formula. If we also knew the value of τ we could solve for x and vice-versa, which means we must find some other way to solve for one of those unknowns.

A drum radius of 3 inches is equivalent to 0.25 feet, and since we know the relationship between radius, force, and torque we may solve for torque in that manner:

$$\tau = rF$$

$$\tau = (0.25 \text{ ft})(2380 \text{ lb})$$

$$\tau = 595 \text{ lb}\cdot\text{ft}$$

Now that we have a value for τ we may substitute it into the rotational work formula to solve for the number of necessary drum rotations:

$$W = 2\pi\tau x$$

$$x = \frac{W}{2\pi\tau}$$


$$x = \frac{35700 \text{ ft}\cdot\text{lb}}{(2\pi)(595 \text{ lb}\cdot\text{ft})}$$

$$x = 9.549 \text{ revolutions}$$

As always, energy is conserved in this electric winch system just as it is in any other system. The 35700 foot-pounds of energy invested in the gravitational potential energy of the weight had to originate from somewhere, which in the case of an electric winch is the electrical power source feeding the winch motor. If that winch motor were allowed to turn in reverse and act as a generator, it would convert the descending weight's loss in potential energy into electrical power to be returned to the source, whether that source by a rechargeable battery or an electrical power "grid" with other electrical loads that may use that energy.

The same, of course, is true for the tow truck example. The energy expended in towing the load must come from somewhere, and in the case of a combustion-type truck engine that source is the fuel powering the engine. For electric vehicles, the energy source is a rechargeable battery, and the reversible nature of electric motors means the vehicle may use its drive motor as a generator to "brake" (slow down), returning the vehicle's kinetic energy into electrical energy to recharge the battery for later use. This is why all-electric and hybrid-electric vehicles are remarkably efficient in stop-and-go traffic conditions: they utilize their drive motors as *regenerative brakes* to recover the vehicle's kinetic energy rather than dissipating that same energy in the form of heat using friction-based brake mechanisms.


Many applications exist within the industrial world for rotational work and electric motors. Conveyor belts, pumps, fans, compressors, and a host of other mechanisms may be powered via electric motors, the speed and torque of those electric motors controlled using electronic circuits called *motor drives*. Alternating-current (AC) electric motors are controlled by electronic devices called *Variable Frequency Drives*, or *VFDs*, which achieve precise speed control by varying the frequency of the AC power feeding the motor, and achieve precise torque control by varying the voltage and current feeding the motor. VFDs are very important "final control" devices used in a wide variety of industrial control systems.



Sophisticated VFDs control regeneration as well (i.e. using the motor as a generator to "brake" a rotating machine), allowing energy to be applied to and then extracted from the same mechanism.

Potential energy is an important principle not just in the study of physics, but also for workplace safety. When large amounts of potential energy are released, the effects may be hazardous to personnel and/or destructive to equipment. An industrial maintenance procedure known as *lock-out, tag-out* requires that all potential energy sources on a system must either be exhausted or otherwise secured to that there will be negligible risk to maintenance personnel as they perform work on a system. The most common way to ensure this is to place a padlock on each energy-disconnect device (e.g. switch, valve, etc.) to secure its position so that potential energy cannot be released in a hazardous or destructive way. Each maintenance worker places a padlock on that disconnect device to prevent its actuation, and also places a tag on the device explaining when and why it was secured. Each and every person performing work on the system must use their own personal padlock, so that the system cannot be re-activated without the active consent of all persons working on that system.

An efficient strategy for safely locking out a large number of safety-disconnect devices on a system with multiple personal locks is to use a sheet-metal box containing a numbered padlock (with matching key) for each energy-flow device to be secured on the equipment, as well as a list identifying which lock goes on which energy-flow device. Each energy-disconnect device is placed in the safe position and then locked in that position with a dedicated padlock. After that, all the padlock keys are placed back inside the sheet-metal box. The lid of this box is then shut and secured with a multi-lock device, permitting multiple peoples' personal locks to be applied so the lid cannot be opened unless *all* personal locks are removed from it:



None of the energy-securing devices may be altered unless *all* personal locks have been removed from the lock box, thereby ensuring the safety of all persons working on the system.

Procedures created and maintained at the worksite will identify the energy-flow devices in need of securing prior to commencement of work on a piece of equipment. These procedures are literally life-saving documents, as they ensure no energy-securing device is overlooked by personnel doing work on the equipment or system.

A photograph of such a document – appropriately titled an “Energy Control Procedure” – shows the steps mandated to secure all potential energy sources prior to commencing work on a large industrial engine. This procedure also serves to document which locks were used to secure which energy flow devices during the procedure, as well as who performed the procedure:

Energy Control Procedure			
Revision Number & Date one 7/23/2009			
Ingersol Rand KVS 412 Minor maintenance			
Items to be completed before proceeding			
Initial all steps			
		1	Only qualified /authorized employees are allowed to complete this task.
		2	Verify equipment has not changed before relying on this procedure
		3	Redline and update any changes & notify effected employee if required
		4	Notify all effected employees before starting
		5	Notify Gas Control
		6	If the crane is to be used complete crane inspection
		7	Complete all required documentation
		8	Identify required forms, Use updated forms
		9	Hot Work Permit WGP- 0059 Procedure 65.0020.05
		10	Crane pre use inspection. Form Attachment A Procedure 65.00.08.06
		11	Drawing ##### Not available at this time
Date & Time		9-25-11 10:00 a.m.	
Authorized Technician		Grieser Carlos	
Scope of work		spark plug wires expansion inlet	
Out of Service			
Step	Initials	Lock #	Task
1	JD	1	Tag out unit mode switch in off position
2	JD	2	Close starting air block valve. Lock and tag.
3	JD	3	Close fuel block valve. Lock and tag.
Return to Service			
Step	Initials	Lock # off	Task
1			Remove tag from unit mode switch. Place switch in local manual
2			Remove lock from starting air block valve open valve.
3			Remove tag from fuel block valve open valve
4			Reset unit panel alarms
5			Place unit panel selector switch in remote auto.
Affected employee log			
Lock	On	Off	Name
1	10:00		JD
CW	10:40		Carlos Grieser
CW		12:40	JD

Note the particular lock-out steps required in this procedure: switching the control mode to the “off” position and tagging it, closing the fuel gas valve supplying fuel to the engine and locking/tagging it, and finally closing the valve supplying high-pressure air for engine starting and locking/tagging it. The closure of the starting air valve prevents the engine from being pneumatically turned while personnel are performing work on it. The closure of the fuel gas valve eliminates hazards resulting from the pressure of the fuel gas as well as its chemical energy (i.e. fire hazard) and/or biological threats (poisoning or asphyxiation). Note also how this procedure lists steps of notification to be taken prior to locking or tagging anything on the engine, as well as any other procedures possibly necessary (e.g. inspecting the maintenance crane if that will be needed for work on the engine).

The following is a set of incomplete lists of various energy-securing devices and energy sources which may be “locked out” for maintenance on a larger system:

Electrical energy

- Circuit breaker (locked in “off” position, also “racked out” if possible)
- Grounding switch (locked in “on” position to positively ground power conductors)
- Power cord (plastic cap locked onto plug, covering prongs)

Mechanical energy


- Block valve (locked in “shut” position) to prevent pressurized fluid motion
- Flange blind (installed in line) to prevent pressurized fluid motion
- Vent valve (locked in “open” position) to prevent fluid pressure buildup
- Mechanical clutch (disengaged, and locked in that position) to prevent prime mover from moving something
- Mechanical coupling (disassembled, and locked in that state) to prevent prime mover from moving something
- Mechanical brake (engaged, and locked in that position) to prevent motion
- Mechanical locking pin (inserted, and held in position by a padlock) to prevent motion
- Raised masses lowered to ground level, with lifting machines locked out

Chemical energy

- Block valve (locked in “shut” position) to prevent chemical fluid entry
- Vent valve (locked in “open” position) to prevent chemical fluid pressure buildup
- Ventilation fan (locked in “run” state) to prevent chemical vapor buildup

With all these preventive measures, the hope is that no form of potential energy great enough to pose danger may be accidentally released.

Let us return to the crane illustration to explore another concept called *kinetic energy*. As you might guess by the word “kinetic,” this form of energy exists when an object is in the process of *moving*. Let’s imagine the crane lifting the 2380 pound weight 15 feet up into the air, and then the cable snapping apart so that the weight free-falls back to the ground:



The Conservation of Energy still (and always!) holds true: all the potential energy stored in the elevated weight must be accounted for, even when it free-falls. What happens, of course, is that the weight accelerates toward the ground at the rate determined by Earth’s gravity: 32.2 feet per second per second (32.2 ft/s^2). As the weight loses height, its potential energy decreases according to the gravitational potential energy formula ($E_p = mgh$), but since we know energy cannot simply disappear we must conclude it is taking some other form. This “other form” is based on the weight’s *velocity* (v), and is calculated by the kinetic energy formula:

$$E_k = \frac{1}{2}mv^2$$

Where,

E_k = Kinetic energy in joules or newton-meters (metric), or foot-pounds (British)

m = Mass of object in kilograms (metric) or slugs (British)

v = Velocity of mass in meters per second (metric) or feet per second (British)

Thus, the Conservation of Energy explains why a falling object must fall faster as it loses height: kinetic energy must increase by the same amount that potential energy decreases, if energy is to be conserved. A very small amount of this falling weight’s potential energy will be dissipated in the form of heat (as air molecules are disturbed) rather than get converted into kinetic energy. However, the vast majority of the initial 35700 ft-lbs of potential energy gets converted into kinetic energy, until the weight’s energy is all kinetic and no potential the moment it first touches the ground.

When the weight finally slams into the ground, all that (nearly 35700 ft-lbs) of kinetic energy once again gets converted into other forms. Compression of the soil upon impact converts much of the energy into heat (molecular vibrations). Chunks of soil ejected from the impact zone possess their own kinetic energy, carrying that energy to other locations where they slam into other stationary objects. Sound waves rippling through the air also convey energy away from the point of impact. All in all, the 35700 ft-lbs of potential energy which turned into (nearly) 35700 ft-lbs of kinetic energy at ground level becomes dispersed.

The Law of Energy Conservation is extremely useful in projectile mechanics problems, where we typically assume a projectile loses no energy and gains no energy in its flight. The velocity of a projectile, therefore, depends on its height above the ground, because the sum of potential and kinetic energies must remain constant:

$$E_p + E_k = \text{constant}$$

As a projectile rises in altitude (i.e. its gravitational potential energy increases), its velocity must slow down (i.e. its kinetic energy decreases) in order that its total energy remain the same in accordance with the Law of Energy Conservation.

In free-fall problems, where the only source of energy for a projectile is its initial height, the initial potential energy must be equal to the final kinetic energy:

$$E_p \text{ (initial)} = E_k \text{ (final)}$$

$$mgh_i = \frac{1}{2}mv_f^2$$

We can see from this equation that mass cancels out of both sides, leaving us with this simpler form:

$$gh_i = \frac{1}{2}v_f^2$$

It also leads to the paradoxical conclusion that the mass of a free-falling object is irrelevant to its velocity. That is, both a heavy object and a light object in free fall hit the ground with the same velocity, and fall for the same amount of time, if released from the same height under the influence of the same gravity¹².

¹²In practice, we usually see heavy objects fall faster than light objects due to the resistance of air. Energy losses due to air friction nullify our assumption of constant total energy during free-fall. Energy lost due to air friction never translates to velocity, and so the heavier object ends up hitting the ground faster (and sooner) because it had much more energy than the light object did to start.

Dimensional analysis confirms the common nature of energy whether in the form of potential, kinetic, or even mass (as described by Einstein's equation). First, we will set these three energy equations next to each other for comparison of their variables:

$$E_p = mgh \quad \text{Potential energy due to elevation}$$

$$E_k = \frac{1}{2}mv^2 \quad \text{Kinetic energy due to velocity}$$

$$E = mc^2 \quad \text{Mass-to-energy conversion}$$

Next, we will dimensionally analyze them using standard SI metric units (kilogram, meter, second). Following the SI convention, mass (m) is always expressed in kilograms [kg], distance (h) in meters [m], and time (t) in seconds [s]. This means velocity (v , or c for the speed of light) in the SI system will be expressed in meters per second [m/s] and acceleration (a , or g for gravitational acceleration) in meters per second squared [m/s²]:

$$\frac{[\text{kg}][\text{m}^2]}{[\text{s}^2]} = [\text{kg}] \left[\frac{\text{m}}{\text{s}^2} \right] [\text{m}] \quad \text{Potential energy due to elevation}$$

$$\frac{[\text{kg}][\text{m}^2]}{[\text{s}^2]} = [\text{kg}] \left[\frac{\text{m}}{\text{s}} \right]^2 \quad \text{Kinetic energy due to velocity}$$

$$\frac{[\text{kg}][\text{m}^2]}{[\text{s}^2]} = [\text{kg}] \left[\frac{\text{m}}{\text{s}} \right]^2 \quad \text{Mass-to-energy conversion}$$

In all three cases, the unit for energy is the same: kilogram-meter squared per second squared. This is the fundamental definition of a “joule” of energy (also equal to a “newton-meter” of energy), and it is the same result given by all three formulae.

Power is defined as the rate at which work is being done, or the rate at which energy is transferred. Mathematically expressed, power is the first time-derivative of work (W):

$$P = \frac{dW}{dt}$$

The metric unit of measurement for power is the *watt*, defined as one joule of work performed per second of time. The British unit of measurement for power is the *horsepower*, defined as 550 foot-pounds of work performed per second of time.

Although the term “power” is often colloquially used as a synonym for force or strength, it is in fact a very different concept. A “powerful” machine is not necessarily a machine capable of doing a great amount of work, but rather (more precisely) a great amount of work *in a short amount of time*. Even a “weak” machine is capable of doing a great amount of work given sufficient time to complete the task. The “power” of any machine is the measure of *how rapidly* it may perform work.

2.8.3 Mechanical springs

Many instruments make use of springs to translate force into motion, or vice-versa. The basic “Ohm’s Law” equation for a mechanical spring relating applied force to spring motion (displacement) is called *Hooke’s Law*¹³:

$$F = -kx$$

Where,

F = Force generated by the spring in newtons (metric) or pounds (British)

k = Constant of elasticity, or “spring constant” in newtons per meter (metric) or pounds per foot (British)

x = Displacement of spring in meters (metric) or feet (British)

Hooke’s Law is a linear function, just like Ohm’s Law is a linear function: doubling the displacement (either tension or compression) doubles the spring’s force. At least this is how springs behave when they are displaced a small percentage of their total length. If you stretch or compress a spring more substantially, the spring’s material will become strained beyond its elastic limit and either yield (permanently deform) or fail (break).

The amount of potential energy stored in a tensed spring may be predicted using calculus. We know that potential energy stored in a spring is the same as the amount of work done on the spring, and work is equal to the product of force and displacement (assuming parallel lines of action for both):

$$E_p = Fx$$

Thus, the amount of work done on a spring is the force applied to the spring ($F = kx$) multiplied by the displacement (x). The problem is, the force applied to a spring varies with displacement and therefore is not constant as we compress or stretch the spring. A mathematician would say that the spring’s force *is a function of x* because the force varies as x varies. Thus, in order to calculate the amount of potential energy stored in the spring ($E_p = Fx$), we must calculate the amount of energy stored over infinitesimal amounts of displacement ($F dx$, or $kx dx$) and then add those bits of energy up (\int) to arrive at a total:

$$E_p = \int kx dx$$

¹³Hooke’s Law may be written as $F = kx$ without the negative sign, in which case the force (F) is the force *applied* on the spring from an external source. Here, the negative sign represents the spring’s reaction force to being displaced (the *restoring* force). A spring’s reaction force always opposes the direction of displacement: compress a spring, and it pushes back on you; stretch a spring, and it pulls back. A negative sign is the mathematically symbolic way of expressing the opposing direction of a vector.

We may evaluate this integral using the power rule (x is raised to the power of 1 in the integrand):

$$E_p = \frac{1}{2}kx^2 + E_0$$

Where,

E_p = Energy stored in the spring in joules (metric) or foot-pounds (British)


k = Constant of elasticity, or “spring constant” in newtons per meter (metric) or pounds per foot (British)

x = Displacement of spring in meters (metric) or feet (British)

E_0 = The constant of integration, representing the amount of energy initially stored in the spring prior to our displacement of it

For example, if we take a very large spring with a constant k equal to 60 pounds per foot and displace it by 4 feet, we will store 480 foot-pounds of potential energy in that spring (i.e. we will do 480 foot-pounds of work on the spring).

Graphing the force-displacement function on a graph yields a straight line (as we would expect, because Hooke’s Law is a linear function). The area accumulated underneath this line from 0 feet to 4 feet represents the integration of that function over the interval of 0 to 4 feet, and thus the amount of potential energy stored in the spring:




Note how the geometric interpretation of the shaded area on the graph exactly equals the result predicted by the equation $E_p = \frac{1}{2}kx^2$: the area of a triangle is one-half times the base times the height. One-half times 4 feet times 240 pounds is 480 foot-pounds.

2.8.4 Rotational motion

Rotational motion may be quantified in terms directly analogous to linear motion, using different symbols and units.

The rotational equivalent of linear *force* (F) is *torque* (τ). Linear force and rotational torque are both vector quantities, mathematically related to one another by the radial distance separating the force vector from the centerline of rotation. To illustrate with a string pulling on the circumference of a wheel:



This relationship may be expressed mathematically as a *vector cross-product*, where the vector directions are shown by the *right-hand rule* (the first vector \vec{r} is the direction of the index finger, the second vector \vec{F} is the direction of the middle finger, and the product vector $\vec{\tau}$ is the direction of the thumb, with all three vectors perpendicular to each other):


$$\vec{\tau} = \vec{r} \times \vec{F}$$

Labeling force, radius, and torque as vectors is the formally correct way of noting the variables in a mechanical system such as this, and is the way college students studying physics typically learn the calculation of torque. In less academic settings, the force vector (\vec{F}) is typically labeled as a force along the *line of action*, and the radius vector (\vec{r}) is called the *moment arm*, with the line of action and moment arm always being perpendicular to each other.

The proper unit of measurement for torque is the product of the force unit and distance unit. In the metric system, this is customarily the *Newton-meter* (N-m). In the British system, this is customarily the *foot-pound* (ft-lb) or alternatively the *pound-foot* (lb-ft). Note that while these are the exact same *units* as those used to express work, they are not the same types of *quantities*. Torque is a vector cross-product, while work is a *dot-product* ($W = \vec{F} \cdot \vec{x}$). The cross-product of two vectors

is always another vector¹⁴, while the dot-product of two vectors is always a scalar (direction-less) quantity. Thus, torque always has a direction, whereas work or energy does not.


An example calculation applied to a hand wrench turning a bolt appears here:




With the radius and force vectors at right angles to each other, torque is simply the product of both. In many non-academic settings, torque is calculated this way as a scalar quantity, with the direction of rotation determined by observation rather than by strict adherence to the right-hand rule of vector cross products. In this example, we see the magnitude of torque as the simple product of 58 pounds force and 0.61 feet of moment arm (35.38 lb-ft of torque), with the torque direction obviously counter-clockwise as viewed from the head of the bolt.

¹⁴Technically, it is a *pseudovector*, because it does not exhibit all the same properties of a true vector, but this is a mathematical abstraction far beyond the scope of this book!

If we apply the same force to the wrench handle at a different angle (not perpendicular to the handle), the resulting torque will be less. The radius vector (moment arm), however, will still remain perpendicular to the force vector (line of action) – it just decreases in length. To determine the placement of the radius vector, all one must do is draw a line straight from the axis of rotation perpendicular to the line of action, then use trigonometry to calculate its magnitude:



A very practical example of torque is in the action of meshing gears, transferring mechanical power from one gear to another. Each gear effectively acts as a wheel, the point of contact between gear teeth acting to transfer force perpendicular to the radius of each gear (wheel). Thus, torque applied to one gear becomes a linear force at the meshing teeth, which translates into another torque at the second gear:



The ratio of torques between two meshing gears is equal to the ratio of gear teeth:

$$\frac{\tau_1}{\tau_2} = \frac{n_1}{n_2}$$

Where,

τ_1 = Torque of first gear

τ_2 = Torque of second gear

n_1 = Number of teeth on first gear

n_2 = Number of teeth on second gear

For example, if a small gear having 28 teeth meshes with a larger gear having 75 teeth, the torque multiplication factor from the small gear to the large gear will be 75:28, or 2.679 to 1. A torque of 40 lb-ft applied to the small gear will result in a torque of 107.1 lb-ft or torque generated at the large gear. This ratio of gear teeth is called the *gear ratio*.

As gears multiply torque (τ), they divide rotational speed (ω). Thus, the 75:28 tooth gear set creates a multiplication of torque from the small gear to the large gear, and an identical *reduction ratio* of speed from the small gear to the large gear. Given this ratio, the small gear will have to be turned 2.679 revolutions in order to make the large gear turn just one revolution.

We may express gear speeds as another ratio of gear teeth, reciprocated in relation to torque:

$$\frac{\omega_1}{\omega_2} = \frac{n_2}{n_1}$$

Where,

ω_1 = Rotational speed of first gear

ω_2 = Rotational speed of second gear

n_1 = Number of teeth on first gear

n_2 = Number of teeth on second gear

In a set of meshed gears, the smaller gear will have the least torque and the greatest speed; the larger gear will have the greatest torque and the least speed.

This is precisely how gear sets are used in industry: to transform torque and speed in mechanical power systems. The complementary effects of a gear set on torque and speed is analogous to the complementary effects that a transformer has on AC voltage and current: a *step-up transformer* (having more turns of wire in the secondary coil than in the primary coil) will multiply voltage but reduce (divide) current, both by the same turns ratio.

Every quantity of force and motion which may be expressed in linear form has a rotational equivalent. As we have seen, torque (τ) is the rotational equivalent of force (F). The following table contrasts equivalent quantities for linear and rotational motion (all units are metric, shown in *italic* font):

Linear quantity, symbol, and unit	Rotational quantity, symbol, and unit
Force (F) <i>N</i>	Torque (τ) <i>N-m</i>
Linear displacement (x) <i>m</i>	Angular displacement (θ) <i>radian</i>
Linear velocity (v) <i>m/s</i>	Angular velocity (ω) <i>rad/s</i>
Linear acceleration (a) <i>m/s²</i>	Angular acceleration (α) <i>rad/s²</i>
Mass (m) <i>kg</i>	Moment of Inertia (I) <i>kg-m²</i>

Familiar equations for linear motion have rotational equivalents as well. For example, Newton's Second Law of motion states, "The acceleration of an object is directly proportional to the net force acting upon it and inversely proportional to the object's mass." We may modify this law for rotational motion by saying, "The angular acceleration of an object is directly proportional to the net torque acting upon it and inversely proportional to the object's moment of inertia." The mathematical expressions of both forms of Newton's Second Law are as follows:

$$F = ma \qquad \tau = I\alpha$$


The calculus-based relationships between displacement (x), velocity (v), and acceleration (a) find parallels in the world of angular motion as well. Consider the following formula pairs, linear motion on the left and angular motion on the right:

$$v = \frac{dx}{dt} \quad (\text{Velocity as the time-derivative of displacement}) \qquad \omega = \frac{d\theta}{dt}$$

$$a = \frac{dv}{dt} \quad (\text{Acceleration as the time-derivative of velocity}) \qquad \alpha = \frac{d\omega}{dt}$$

$$a = \frac{d^2x}{dt^2} \quad (\text{Acceleration as the second time-derivative of displacement}) \qquad \alpha = \frac{d^2\theta}{dt^2}$$

An object's "moment of inertia" represents its angular inertia (opposition to changes in rotational velocity), and is proportional to the object's mass and to the square of its radius. Two objects having the same mass will have different moments of inertia if there is a difference in the distribution of their mass relative to radius. Thus, a hollow tube will have a greater moment of inertia than a solid rod of equal mass, assuming an axis of rotation in the center of the tube/rod length:



This is why *flywheels*¹⁵ are designed to be as wide as possible, to maximize their moment of inertia with a minimum of total mass.

The formula describing the amount of work done by a torque acting over an angular displacement is remarkably similar to the formula describing the amount of work done by a force acting over a linear displacement:

$$W = Fx \qquad W = \tau\theta$$

The formula describing the amount of kinetic energy possessed by a spinning object is also similar to the formula describing the amount of energy possessed by a linearly-traveling object:

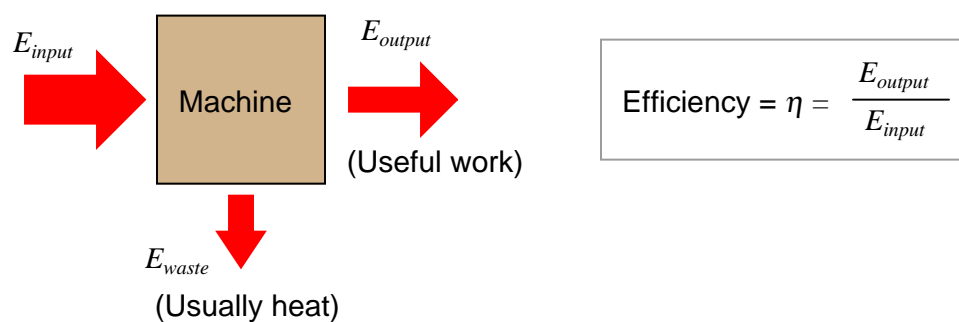
$$E_k = \frac{1}{2}mv^2 \qquad E_k = \frac{1}{2}I\omega^2$$

¹⁵A "flywheel" is a disk on a shaft, designed to maintain rotary motion in the absence of a motivating torque for the function of machines such as piston engines. The rotational kinetic energy stored by an engine's flywheel is necessary to give the pistons energy to compress the gas prior to the power stroke, during the times the other pistons are not producing power.

2.9 Simple machines

A *machine* in the broad sense of the word is any device designed to translate some form of energy into useful work. A “simple” machine is one where both the input energy and the output energy are mechanical in nature (i.e. both are forces acting along displacements). Examples of simple machines include levers, pulleys, ramps, wedges, gears, and chain/sprockets. More complex machines include such examples as electric motors, heat engines, pumps, compressors, and refrigerators.


The *efficiency* of any machine (symbolized by the Greek letter “eta” η) is defined as the ratio of output energy to input energy:




Ideally, these two will be equal, with all of the input energy translated losslessly into output energy. However, no machine is perfectly efficient although some simple machines come very close to achieving 100% efficiency. It is physically impossible to achieve an energy efficiency greater than 100%, as that would violate the Law of Energy Conservation.

2.9.1 Levers

Perhaps the most basic type of machine is the *lever*: a rigid beam pivoting on an axis. This axis may be something as simple as a round cylinder, a pointed wedge, or even a sophisticated bearing. In any case, the general term for the pivot point on a lever is *fulcrum*:




If we look at the lever's motion at each end, we see that the distance the “output” end moves is a function of how far the “input” end moves as well as the ratio of lengths from each end to the fulcrum. Showing examples using three different classes of lever, each one with an $\frac{8}{3}$ length ratio:



The ratio of output force to input force ($\frac{F_{out}}{F_{in}}$) is called the *mechanical advantage*¹⁶ of the machine. This ratio is always the reciprocal of the output versus input *motion*: if the output of the lever moves less than the input moves, the output force must be greater than the input force, and vice-versa. This makes perfect sense if you view a lever as a perfectly efficient machine where the output energy (work) must equal the input energy (work): since output energy is output force multiplied by output motion, and input energy is input force multiplied by input motion, in order for force to be multiplied, motion must be diminished.


Leverage abounds in everyday life. A shovel, for example, functions as either a first-class lever or a second-class lever, depending on its use. In either case, it is being used as a force multiplier, the trade-off being that the person must move the handle a farther distance than the rock moves, thus exchanging motion for force:



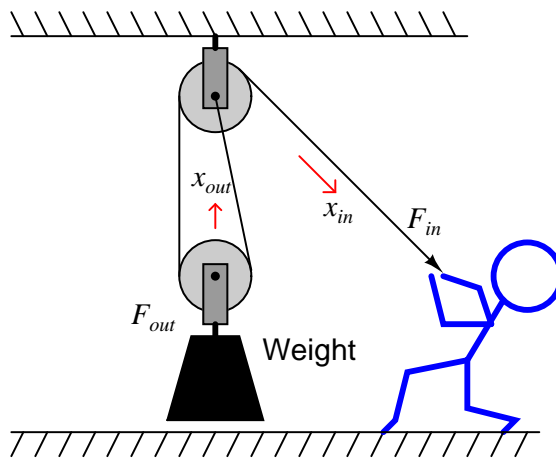
¹⁶Technically, mechanical advantage should be defined by the ratio of input *motion* to output *motion*, rather than being defined in terms of *force*. The reason for this is if friction happens to exist in the machine, it will cause a degradation of force but not of motion. Since “mechanical advantage” is supposed to represent the ideal ratio of the machine, it is always safest to define it in terms of motion where friction will not affect the calculation. For a frictionless machine, however, defining mechanical advantage in terms of force is perfectly legitimate, and in fact makes more intuitive sense, since a larger mechanical advantage always corresponds with force multiplication from input to output.

2.9.2 Pulleys

Another simple and useful machine is a *pulley* and rope. A “pulley” is nothing more than a wheel with a groove cut around its circumference to guide a rope or cable, a bearing and axle supporting the wheel and allowing it to freely turn. A single pulley hung from an overhead support has the ability to convert downward motion of a rope into upward motion to hoist a load:




A single-pulley system such as this exhibits no mechanical advantage, because $F_{out} = F_{in}$. If we get creative with multiple pulleys, however, we can achieve a mechanical advantage sufficient to hoist very heavy loads with modest input force:




Here, the weight is being supported by the tension within *two* ropes, not just one rope. Since the person’s force on the rope is what generates the rope’s tension, F_{in} is equal to rope tension, while F_{out} is equal to twice the rope’s tension. Thus, this simple machine has a mechanical advantage equal to 2. It also means the person’s motion while pulling the rope will be exactly *twice* the motion of the hoisted weight. Remember that we cannot cheat the Law of Energy Conservation: work

in cannot be less than work out. If the output force is twice as much as the input force due to mechanical advantage, the output motion can only be *half* as much as the input motion.

The mechanical advantage of a pulley system may be extended beyond two by adding even more pulleys. This pulley system has a mechanical advantage of 4, since the weight is being supported by the tension of *four* ropes, while the person pulling only feels the tension of a single rope:



Here is where one must be careful in analyzing pulley systems with regard to mechanical advantage. The mechanical advantage in each of these examples was based on the number of ropes supporting the weight. So far, this also happened to equal the number of pulleys in the system. Lest anyone be tempted to determine mechanical advantage by simply counting pulleys, here is an example that breaks the pattern:




Here there is only one pulley in the system, yet the weight is being supported by the tension in *two* ropes and the person pulling on the rope only feels the tension of one rope, which means the system has a mechanical advantage of 2.

This simple technology is commonly used on cranes to provide huge amounts of lifting force with modest amounts of cable tension. In this photograph you can see the multiple pulleys and lifting cable of a large industrial crane:



2.9.3 Inclined planes

A *wedge*, also referred to as an *inclined plane*, is another type of simple machine. A large enough wedge such as a ramp is useful for producing a mechanical advantage to lift a weight equipped with wheels. Instead of hoisting the weight vertically, the weight is rolled up the diagonal incline of the ramp:



In moving the heavy weight a short distance vertically, the person pushes with much less force over a longer distance. The mechanical advantage of this ramp, therefore, is equal to the ratio of the ramp's diagonal length (hypotenuse side) to its vertical height (opposite side). From the perspective of angle θ shown in the illustration, this equates to the *cosecant* function ($\csc \theta = \frac{\text{hypotenuse}}{\text{opposite}}$).

Another example of an inclined plane is a *screw conveyor* or *auger*, shown in the following photograph. The “fins” on the screw function as a long incline, wrapped around a central shaft:




Placed inside of a pipe and turned slowly, this simple machine moves semi-solid material linearly through the pipe.

In a similar fashion, this electric valve actuator uses the principle of an inclined plane to raise and lower a heavy gate to control the flow of wastewater through channels at a municipal wastewater treatment facility. The long threaded shaft pulls upward on the heavy gate (not shown), moved by the turning action of a nut engaged with the shaft's threads. The electric motor inside the blue-colored actuator turns the nut on command, raising or lowering the gate as needed:



2.9.4 Gears

A *gear set* is another type of simple machine, part of a whole class of simple machines converting one form of rotary motion into another forms of rotary motion. A set of *spur gears* are shown here:




Each gear rotates about a central axis (usually a rotating shaft), the teeth of each gear cut into shapes designed to smoothly “mesh” together as the gears rotate. Gears may be thought of as levers, with the radius of each gear equivalent to the distance between the fulcrum and the force point on a level. The gear with the largest radius turns the slowest, and with the most torque¹⁷.


The mechanical advantage of a gear set is simply the ratio of gear diameters. Another way to determine gear ratios is to count the number of teeth on each gear: since the teeth are all cut to the same size so as to smoothly mesh, the ratio of gear teeth will be proportional to the ratio of gear circumferences, which in turn must be proportional to the ratio of gear radii. If a gear set may be turned by hand, a simple counting of turns from input to output will also allow you to calculate the gear ratio. For example, if you turn one gear 15 revolutions to get the second gear to turn 4 revolutions, the mechanical advantage is $\frac{15}{4}$, or 3.75. As with levers, the gear that turns the farthest does so with less force (torque), and vice-versa. All simple machines work by trading motion for force, so that an increase in one necessarily results in a decrease of the other.

¹⁷“Torque” is to rotational motion as “force” is to linear motion. Mathematically, torque (τ) is defined as the cross-product of force acting on a radius ($\vec{\tau} = \vec{r} \times \vec{F}$).

A variety of spur gear designs appears on this page. In this illustration¹⁸, we see an external spur gear set with straight-cut teeth, perhaps the simplest style of gear:




Next, we see variations on this design where the gear teeth are cut at angles instead of being parallel with the gears' shafts. This causes the meshing of the gear teeth to be smoother and quieter, but also causes a *thrust* force to develop along the axis of the gear shaft, since the teeth act as inclined planes. A “double” helical gear pattern (also known as a *herringbone* gear due to its resemblance to a fish skeleton) cancels any thrust force by angling the teeth in opposite angles. Herringbone gear sets are quiet and strong, but tend to be more expensive to manufacture than single-helical gears:




¹⁸I am indebted to NASA for this and the rest of the black-and-white gear illustrations found in this section. All these illustrations were taken from NASA technical reports on gearing.

The exposed gear sets commonly found in antique machinery provide excellent visual examples of gear designs. This next photograph shows sets of external spur gears. The upper photograph shows a pair of meshing spur gears with parallel teeth, while the lower photograph shows a set of four meshing spur gears with single-helical teeth, both sets of gears found on antique gasoline engines:



Another style of gear set is called *planetary*, because its shape resembles circular orbits of planets around a star. Planetary gear sets are exceptionally strong, and are capable of delivering multiple gear ratios depending on which gear (the “sun” gear, the “ring” gear, or the “planet” gears) is being held stationary, which gear is the input, and which gear is the output. If any two sets of gears are locked together such that they rotate at the same speed, the third gear in a planetary mechanism must also rotate at that same speed, for a 1:1 ratio. Typically, the planet gears are all anchored in their respective positions by a rotating frame called a *carrier*:



The particular planetary gear set shown in the above illustration uses two sets of helical gears (much like a herringbone design, only with two single-helical gears placed back-to-back instead of one gear with double-helical teeth) in order to eliminate thrust forces on the shafts.

Planetary gear sets are the standard type of gears used in automatic transmissions for automobiles. Different gear ratios (e.g. “Low”, “Drive”, “Overdrive”) are achieved in an automatic transmission by selecting which gears in a planetary gear set are input, output, and stationary. A series of clutches engage and disengage shafts to control which gears are input versus output, while a series of bands act as brakes to hold different gears stationary in the planetary gear set. These clutches and bands are all operated by hydraulic oil pressure inside the transmission, controlled either by a series of hydraulic relays and/or by an electronic computer telling the transmission when to shift.

The SynergyTM gear drive system designed by Toyota for use in its line of hybrid gasoline-electric cars is a unique application of planetary gears. In the first-generation Toyota Prius, an electric motor/generator (“MG1”) is coupled to the sun gear, the internal-combustion (gasoline) engine is coupled to the planet carrier, and the driveshaft is coupled to the ring gear (as well as a second electric motor/generator “MG2”). This planetary gear set allows power to be transferred smoothly and with great flexibility between the engine, the motor/generator, and the driveshaft. Motor/generator MG1 functions as a kind of variable brake (in “generator” mode, passing its power


to either the battery or to MG2) to slow down the sun gear to achieve an infinite number of effective gear ratios for the engine, and the gasoline engine may also be locked to keep the planet gear carrier from turning during all-electric operation.

With a simple set of parallel-shaft spur gears, the ratio of the gear set is simply the ratio of gear teeth (or of effective diameters). For example, if a spur gear set has 15 teeth on the driving gear ($N_{driving} = 15$) and 45 teeth on the driven gear ($N_{driven} = 45$), the gear ratio will be a 45:15 or 3:1 reduction in speed (multiplication in torque). Planetary gear sets are more complicated than this, as shown by the following table:

Condition	Slow	Fast	Ratio ($x : 1$)
Ring gear held stationary	Planet carrier	Sun	$\frac{N_r}{N_s} + 1$
Sun gear held stationary	Planet carrier	Ring	$\frac{N_s}{N_r} + 1$
Planet carrier held stationary	Ring	Sun	$-\frac{N_r}{N_s}$
Any two locked together	—	—	1


It is interesting to note that the only gear teeth (diameters) values factoring into these ratio calculations belong to the ring and sun gears. The negative sign for the stationary-carrier condition refers to the reversed rotation of the ring gear compared to the sun gear.

As always, one should strive to *understand* rather than memorize when learning anything new, and planetary gear set ratios are no exception to this rule. An excellent exercise is to mentally visualize each of the conditions listed in the table above, applied to a graphic image of a planetary gear¹⁹ set. Run a series of “thought experiments” on the gear set, where you imagine one of the three pieces being held stationary while one of the free pieces is turned. Ask yourself whether the third piece turns faster or slower than the other free piece. Then, imagine the sun gear growing or shrinking in size, and ask yourself how this change in sun gear size affects the speed ratio:




¹⁹Here, each gear is shown simply as a toothless wheel for the sake of simplicity. Truth be told, your humble author has difficulty drawing realistic gear teeth!

A variety of gear set designs with perpendicular shafts exist to transfer mechanical power around corners. First, we see a *crossed helical spur gear*. Like parallel-shaft helical spur gears, crossed helical gears generate thrust forces due to the action of the gear teeth as inclined planes:




Next we see a *bevel gear* or *miter gear* set, where a *pinion* gear intersects with a *ring* gear to transfer mechanical power through perpendicular shafts. The left-hand illustration shows a *straight-toothed* bevel gear set, while the right-hand illustration shows a *spiral-toothed* bevel gear set.



These two styles of bevel gears are analogous to the straight- versus helical- toothed variants of the spur gear, with similar characteristics: spiral-toothed bevel gears provide smoother and quieter operation than straight-toothed bevel gears, but at the expense of generating large thrust forces on the pinion gear shaft, and radial forces on the ring gear shaft.

An interesting variation on the bevel gear concept is the *hypoid gear* system, where the two shaft axes do not intersect. In this gear set, the gear teeth actually rub against each other rather than merely touch, necessitating special lubricant to tolerate the dynamic pressures and stresses. Hypoid gear sets are exceptionally strong and quiet-running, and became popular for automotive axle drive systems because they allowed the driveshaft (attached to the pinion gear) to be lower than the axles (attached to the ring gear), providing more floor space in the vehicle. The non-intersecting shaft centerlines also make it possible to place support bearings on both ends of the pinion gear for extra strength, as seen in heavy-duty truck axle designs:




A photograph of a hypoid gear set inside the differential of an automobile is shown here, the differential housing cover removed for inspection:



An important type of gear set used for perpendicular-shaft applications with large speed-reduction ratios is the *worm gear*. A worm gear resembles a screw whose threads engage with matching helical-cut threads on another gear called a *worm wheel*:

Worm gear mechanism



An interesting and useful feature of a worm gear set is that power transfer occurs easily from the worm screw to the worm wheel, but not so easily from the worm wheel to the worm screw due to friction between the teeth of the two gears. This means when the worm screw is not being turned by an outside force, even small amounts of friction between the screw threads and wheel teeth will effectively “lock” the worm wheel in place such that it cannot turn when an outside force acts on it. A practical example of a worm gear exploiting this feature is a hand-crank winch, where we desire the winch drum to remain locked in position when we let go of the hand-crank.

Another interesting feature of worm gears is that the gear ratio is simply the number of teeth on the worm wheel, since the tooth pitch on the circumference of the worm wheel defines what the thread pitch *must* be on the worm screw. For example, a worm wheel having 40 teeth around its circumference will exhibit a 40:1 speed-reduction ratio regardless of worm screw size, since there is only one worm screw thread pitch that will engage with the teeth on this worm wheel.


A real worm gear assembly may be seen in this next photograph, as one component of an antique farm implement. Here, a hand-wheel turns the work screw, which then turns the worm wheel and either lifts or drops the height of the implement's blade:



The “one-way” action of a worm gear is advantageous in this application, so that the hand wheel does not spin on its own every time the implement's blade encounters uneven ground. Wherever the hand wheel is set to adjust blade height, that blade height remains fixed until the hand wheel is turned to some new position.

It should be noted that modern worm gear sets, like nearly all types of gears, are more commonly found encased in a housings where they operate in a lubricated environment, sealed from contamination by dust and other matter both solid and liquid. Exposed gears are more commonly seen on antique machinery, where the mechanical stresses were low enough (i.e. low torque forces, large gears) to permit reliable operation of gears with the only lubrication typically being a coat of heavy grease smeared over the gear teeth.

As with other types of simple machines, gear sets may be combined into larger assemblies, with the over-all mechanical advantage (i.e. gear ratio) being the product of all individual gear ratios in the system. An example of a compound gear train is this helicopter transmission, designed to reduce the high-speed shaft rotation of the helicopter's turbine engine down to a speed more suitable for the main rotor:



Mechanical shaft power flows from the turbine shaft (usually several thousand RPM) to the rotor (a few hundred RPM) in this transmission via several types of gears. First, the turbine shaft's speed is reduced through multiple sets of bevel gear operating in parallel (from the "driving spur pinion" gear to the "face gears"). These multiple face gears then couple to a "combining" gear through a spur gear reduction. This combining gear then feeds power to a central "sun" gear in a planetary gear train. The "ring" gear of this planetary set is fixed to the case of the transmission so that it does not turn. Finally, four "planet" gears running in a common carrier assembly drive the helicopter's rotor.

The following photograph shows a modern “gearbox” used to decrease the rotational speed of an electric motor, to turn an auger feeding grain out of a storage bin at a beer brewery. Like all modern gear sets, the gears inside this gearbox operate in a continuously-lubricated environment, sealed to prevent contaminants such as dust and water from entering:




A very large gearbox is shown in this next photograph, used in the head of a Vestas wind turbine to “step up” the slow rotational speed of the turbine to the much higher rotational speed of the electric generator. Planetary gears are used in wind turbine gearboxes due to their ruggedness and relatively compact size:




The ratio of this wind turbine’s gear set happens to be 111.5:1, with the generator turning 111.5 times faster than the turbine blades (but with only $\frac{1}{111.5}$ the torque of the turbine blades, of course)!

2.9.5 Belt drives

A *sheave* is very similar in form and function to a pulley, but designed to grip a flexible *belt* rather than a rope or a cable. Unlike a pulley which is designed to turn freely to re-direct the tension of a rope, a sheave works more like a gear to couple the belt's motion to a rotating shaft. The mechanical advantage of a pair of sheaves coupled by a common belt is simply the ratio of sheave radii, just like gears:



The following photograph shows a triple-belt drive from an electric motor to an agitator on the bottom of a sawdust storage bin:



As indicated by the respective sheave diameters, the electric motor turns much faster than the agitator, while the agitator spins with much greater torque than the motor.

Most modern belt drives are either *V-belt* or *toothed belt*, referring to the shapes of the belt and how they engage with the sheave. The triple-belt drive system for the sawdust agitator shown in the previous photograph used V-belts. A V-belt has a V-shaped cross-section, and sits in a V-shaped groove around the circumference of the sheave. Toothed belts almost resemble chains, in that their inner surface is characterized by regularly-spaced perpendicular ribs designed to engage with matching cavities machined into the circumference of the sheave. The advantage of a toothed belt is that it cannot slip on the sheave if overloaded, unlike a V-belt. A toothed belt is firmly “locked” into place on the sheave’s circumference so long as proper belt tension is maintained.


An older belt-and-sheave technology is the *flat belt*. Here, the sheave’s circumference is flat, and the belt itself is nothing more than a strip of flexible material with no special shape. The following photograph shows an antique flat belt drive in a workshop, where a central shaft ran along the ridge of the ceiling to power several machines in the shop, the shaft itself turned continuously by either a water turbine or a steam engine:



Flat belts are still used in modern times, but they tend to be much wider than V-belts or toothed belts of comparable rating in order to deliver adequate “grip” on the sheave. Also, sheave-to-sheave alignment is much more critical for a flat belt, which has no guides on the sheave to keep it centered²⁰.

²⁰An interesting feature of many flat-belt sheaves is a slight “crown” shape to the sheave, such that the diameter is slightly larger at the sheave’s center than it is at either side edge. The purpose of this crown is to help the belt center itself while in operation. As it turns out, a flat belt naturally tends to find the point at which it operates under maximum tension. If the belt happens to wander off-center, it will naturally find its way back to the center of the sheave as it rotates because that is where the tension reaches a maximum.

Like gear sets, industrial belt drive systems are typically shrouded for cleanliness and for personnel safety. Sheet-metal enclosures such as the one covering the top of this V-belt drive system on a “walking-beam” style of oil field pump. The sheet-metal enclosure protects the belts and sheaves from rain and snow. You will also note a large gearbox following the belt drive, further reducing rotational speed from the electric motor to the pump’s counter-weighted crank:



Belts of all styles are subject to wear and fatigue, and as such must be periodically replaced. Some belt drive systems employ *tensioner* mechanisms which maintain consistent belt tension by applying a constant force to the belt. Small tensioners are usually spring-loaded, while large belt tensioners (particularly conveyor belts) are loaded by the weight of a large mass. Minimum belt tension is extremely important for belt drives, as loose belts will begin to “slip” under load and quickly fail if the problem is not remedied.

When multiple belts are used to distribute loading between belts in high-power drive systems, it is important that all belts be replaced simultaneously, never partially. If a new belt is installed next to an old belt on the same sheave, the old belt will run “loose” and not bear its full share of the load, thus overloading the other (new) belt(s) in the drive system.

2.9.6 Chain drives

Sprockets are identical in function to sheaves, using *link chain* rather than belt to couple the two rotating pieces together. Bicycles are perhaps the best-known example of sprockets and chains from everyday life, being the most efficient simple machine for the purpose of coupling a person's leg power to a rotating wheel for propulsion. Like gear sets, the mechanical advantage ratio of a sprocket set may be determined by counting teeth on each sprocket and then dividing one tooth count by the other, or empirically by rotating one sprocket by hand and counting the number of turns (revolutions) each sprocket makes.

The following photograph shows a pair of sprockets linked together with a roller chain. The sprocket ratio here is 1:1, as both sprockets share the same number of teeth:



Bicycles use sprockets and a chain to transfer power from the crank to the rear wheel. Here, a multi-speed sprocket assembly allows the rider to select the best ratio (i.e. mechanical advantage) for riding at different speeds and in different conditions. Three sprockets on the crank and eight sprockets on the wheel give a theoretical²¹ maximum of 24 different “speeds” or “gears” from which to select:



Chain drive systems require thorough lubrication and freedom from dirt and other abrasive particles in order to deliver full service life. Open-chain systems such as the two shown in the above photographs are challenging to maintain in good working order for these reasons.

²¹In practice, not all of these 24 “speeds” are recommended, because some of the front/rear sprocket selections would place the chain at an extreme angle as it engaged with both sprockets. In the interest of extending chain life, it should run as “straight” on each sprocket as possible.

2.10 Elementary thermodynamics

Thermodynamics is the study of heat, temperature, and their related effects in physical systems. As a subject, thermodynamics is quite complex and expansive, usually taught as a course in itself at universities. The coverage in this book is limited to some of the more elementary and immediately practical facets of thermodynamics rather than a comprehensive overview.

2.10.1 Heat versus Temperature

Most people use the words *heat* and *temperature* interchangeably. This is unfortunate for every student of thermodynamics, because it means they must first deconstruct this false conception and replace it with one more scientifically accurate before any progress may be made. While “heat” and “temperature” are related concepts, they are not identical.

When people say something is “hot,” what they really mean is that the object has a high temperature. Temperature is a direct function of random molecular motion within an object or a fluid sample. This is usually easiest to visualize for a gas, where unattached molecules have great freedom to vibrate, collide, and otherwise move about. The molecules of a substance at high temperature are moving more vigorously (higher velocity) than the molecules of the same substance at low temperature.

Heat, by contrast, is an expression of thermal energy transfer. By placing a pot of water over a fire, we are *adding heat* to that pot (transferring thermal energy to the water), the effect of which is to *raise its temperature* (making the water molecules’ motions more vigorous). If that same pot is taken away from the fire and allowed to cool, its *loss of heat* (transferring energy out of the water to the surrounding air) will result in its *temperature lowering* (the individual water molecules slowing down).

Heat gain or loss often results in temperature change, but not always. In some cases heat may be gained or lost with negligible temperature change – here, the gain or loss of heat manifests as physical changes to the substance other than temperature. One example of this is the boiling of water at constant pressure: no matter how much heat is transferred to the water, its temperature will remain constant at the boiling point (100 degrees Celsius at sea level) until all the water has boiled to vapor. The addition of thermal energy to the boiling water does not raise its temperature (i.e. make the molecules move faster), but rather goes into the work of disrupting inter-molecular bonds so that the liquid turns into vapor. Another example is the heating of chemical reactants in an endothermic (heat-absorbing) reaction: much of the thermal energy added to the chemical solution goes into the work of separating chemical bonds, resulting in molecular changes but not (necessarily) increased temperature.

Heat transfer can *only* happen, though, where there is a difference of temperature between two objects. Thermal energy (heat) naturally flows from the “hotter” (higher-temperature) substance to the “colder” (lower-temperature) substance. To use the boiling water example, the only way to get heat transfer into the water is to subject the water to a hotter substance (e.g., a flame, or a hot electric heating element). If you understand temperature as being molecular motion within a substance, with a hotter object’s molecules vibrating more vigorously than a colder object’s molecules, this natural transfer of heat from hot to cold makes perfect sense: the molecular vibrations of the higher-temperature object literally transfer to the molecules of the lower-temperature object. As those respective molecules touch each other, with fast-vibrating molecules colliding against slow-vibrating molecules, the inter-molecular collisions transfer energy away from the fast-vibrating molecules (so they aren’t vibrating as fast anymore) and toward the slow-moving molecules (so they begin to vibrate faster than before). It’s like a vibrating tuning fork touched to a non-vibrating tuning fork: the vibrating fork loses some of its vibration by transferring energy to the (formerly) quiet tuning fork.

Much more attention will be directed to the concepts of heat and temperature in subsequent subsections.

2.10.2 Temperature

In an ideal, monatomic²² gas (one atom per molecule), the mathematical relationship between average molecular velocity and temperature is as follows:

$$\frac{1}{2}m\overline{v^2} = \frac{3}{2}kT$$

Where,

m = Mass of each molecule

v = Velocity of a molecule in the sample

\overline{v} = Average (“mean”) velocity of all molecules in the sample

$\overline{v^2}$ = Mean-squared molecular velocities in the sample

k = Boltzmann’s constant (1.38×10^{-23} J / K)

T = Absolute temperature (Kelvin)

Non-ideal gases, liquids, and solids are more complex than this. Not only can the atoms of complex molecules move to and fro, but they may also twist and oscillate with respect to each other. No matter how complex the particular substance may be, however, the basic principle remains unchanged: temperature is an expression of how rapidly molecules move within a substance.

There is a temperature at which all molecular motion ceases. At that temperature, the substance cannot possibly become “colder,” because there is no more motion to halt. This temperature is called *absolute zero*, equal to -273.15 degrees Celsius, or -459.67 degrees Fahrenheit. Two temperature scales based on this absolute zero point, *Kelvin* and *Rankine*, express temperature relative to absolute zero. That is, zero Kelvin and zero degrees Rankine is equal to absolute zero temperature. Any temperature greater than absolute zero will be a positive value in either the Kelvin or the Rankine scales. A sample of freezing water at sea level, for example, is 0 degrees Celsius (32 degrees Fahrenheit) but could also be expressed as 273.15 Kelvin²³ (0 plus 273.15) or 491.67 degrees Rankine (32 plus 459.67).

²²Helium at room temperature is a close approximation of an ideal, monatomic gas, and is often used as an example for illustrating the relationship between temperature and molecular velocity.

²³Kelvin is typically expressed without the customary “degree” label, unlike the three other temperature units: (degrees) Celsius, (degrees) Fahrenheit, and (degrees) Rankine.

A table of melting and boiling points (at sea-level atmospheric pressure) for various substances appears in this table, labeled in these four different units of temperature measurement:

Melting or boiling substance	°C	°F	K	°R
Melting point of water (H ₂ O)	0	32	273.15	491.67
Boiling point of water (H ₂ O)	100	212	373.15	671.67
Melting point of ammonia (NH ₃)	-77.7	-107.9	195.45	351.77
Boiling point of ammonia (NH ₃)	-33.6	-28.5	239.55	431.17
Melting point of gold (Au)	1063	1945	1336	2405
Melting point of magnesium (Mg)	651	1203.8	924.2	1663.5
Boiling point of acetone (C ₃ H ₆ O)	56.5	133.7	329.65	593.37
Boiling point of propane (C ₃ H ₈)	-42.1	-43.8	231.05	415.87
Boiling point of ethanol (C ₂ H ₆ O)	78.4	173.1	351.55	632.77

Note how degrees Celsius and Kelvin for each point on the table differ by a constant (offset) of 273.15, while each corresponding degree Fahrenheit and degree Rankine value differs by 459.67 (note that many of the figures in this table are slightly rounded, so the offset might not be *exactly* that much). You might think of Kelvin as nothing more than the Celsius scale zero-shifted by 273.15 degrees, and likewise degrees Rankine as nothing more than the Fahrenheit scale zero-shifted by 459.67 degrees.

Note also how *increments* in temperature measured in degrees Fahrenheit are the same as *increments* of temperature measured in degrees Rankine. The same is true for degrees Celsius and Kelvin. The difference between the melting point of ammonia (-77.7 degrees C) and the melting point of water (0 degrees C) is the same difference as that between the melting points of ammonia and water expressed in Kelvin: 195.45 and 273.15, respectively. Either way, the difference in temperature between these two substances' melting points is 77.7 degrees (C or K). This is useful to know when dealing with temperature changes over time, or temperature differences between points in a system – if an equation asks for a temperature difference (ΔT) in Kelvin, it is the same value as the temperature difference expressed in Celsius. Likewise, a ΔT expressed in degrees Rankine is identical to a ΔT expressed in degrees Fahrenheit. This is analogous to differences between two fluid pressures expressed in PSIG versus PSIA: the differential pressure value (PSID) will be the same.

Most people are familiar with the Fahrenheit and Celsius temperature scales used to express temperature in common applications, but the absolute scales of Rankine and Kelvin have special significance and purpose in scientific endeavors. The fact that Rankine and Kelvin are *absolute* scales in the same manner that *atmospheres* and *torr* are units of absolute pressure measurement makes them uniquely suited for expressing temperature (molecular motion) in relation to the absence of thermal energy. Certain scientific laws such as the *Ideal Gas Law* and the *Stefan-Boltzmann Law* relate other physical quantities to absolute temperature, and so require the use of these absolute units of measurement.

2.10.3 Heat

Heat, being the transfer of energy in thermal (molecular motion) form, may be measured in the same units as any other form of energy is measured: *joules* (metric) and *foot-pounds* (British). However, special units of measurement are often used for heat instead:

- calorie
- kilocalorie (or “dietary Calorie”)
- British Thermal Unit (BTU)

A *calorie* of heat is defined as the amount of thermal energy transfer required to change the temperature of one gram of water by one degree Celsius ($\Delta T = 1\text{ }^{\circ}\text{C} = 1\text{ K}$). One calorie is equivalent to 4.186 joules.

A *British Thermal Unit*, or *BTU* is defined as the amount of thermal energy transfer required to change the temperature of one pound of water by one degree Fahrenheit ($\Delta T = 1\text{ }^{\circ}\text{F} = 1\text{ }^{\circ}\text{R}$). One BTU is equivalent to 778.2 foot-pounds.

The unit of “dietary” calories is used to express the amount of thermal energy available in a sample of food by combustion²⁴. Since the official unit of the “calorie” is so small compared to the typical amounts of energy contained in a meal, nutritionists use the unit of the kilocalorie (1000 calories, or 4186 joules) and call it “Calorie” (with a capital letter “C”).

Just as “Calories” are used to rate the energy content of food, the heat units of “calories” and “BTU” are very useful in describing the potency of various industrial fuels. The following table shows the *heat of combustion* for a few common fuels, in units of kilocalories per gram and BTU per pound:

Fuel	Combustion heat (kcal/g)	Combustion heat (BTU/lb)
Methane (CH ₄)	13.3	23940
Methanol (CH ₄ O)	5.43	9767
Ethanol (C ₂ H ₆ O)	7.10	12783
Propane (C ₃ H ₈)	12.1	21700
Carbon monoxide (CO)	2.415	4347

For example, if exactly one gram of methane gas were completely burnt, the resulting heat liberated in the fire would be sufficient to warm 13.3 kilograms of water from 20 degrees Celsius to 21 degrees Celsius (a temperature rise, or ΔT , of one degree Celsius).

If a meal rated at 900 Calories (900 “dietary calories,” or 900 kilocalories) of energy were completely metabolized, the resulting heat would be sufficient to warm a pool of water 900 kilograms in mass (900 liters, or about 237 gallons) by one degree Celsius. This same amount of heat could raise half the amount of water twice the temperature rise: 450 liters of water warmed two degrees Celsius.

²⁴Animals process food by performing a very slow version of combustion, whereby the carbon and hydrogen atoms in the food join with oxygen atoms inhaled to produce water and carbon dioxide gas (plus energy). Although it may seem strange to rate the energy content of food by measuring how much heat it gives off when *burnt*, burning is just a faster method of energy extraction than the relatively slow processes of biological metabolism.

2.10.4 Heat transfer

Heat spontaneously²⁵ flows from higher-temperature substances to lower-temperature substances. This is the phenomenon you experience standing next to a fire on a cold day. Your body is cold (low temperature), but the fire is much hotter (high temperature), and your proximity to the fire aids in heat transfer from the fire to you.

Three principal methods exist for heat to transfer from one substance to another:

- Radiation²⁶ (by light waves)
- Conduction (by direct contact)
- Convection (by intermediate contact with a moving fluid)

Practical examples of heat transfer often involve multiple modes rather than just one. For example, the transfer of heat to a person's body by sunlight obviously involves *radiation* from the Sun, but it also involves *conduction* through layers of clothing and *convection* by air passing from sun-warmed objects to the person.

Temperature-sensing instruments used to measure temperature in industrial applications likewise rely on multiple heat-transfer modes to sample thermal energy from a process fluid or object(s). Infrared thermometers detect temperature by sensing the intensity of infrared light *radiated* by hot objects. A thermocouple directly touching a hot object relies on *conduction* to sense the temperature of that object. An RTD inserted into a pipe carrying a hot fluid relies on *convection* to measure the average temperature of that fluid. A filled-bulb thermometer inserted into a thermowell, inserted into a fluid-filled process vessel relies on both *convection* (from the process fluid to the thermowell) and *conduction* (from the thermowell to the bulb) to sense process temperature.

²⁵Heat may be forced to flow from cold to hot by the use of a machine called a *heat pump*, but this direction of heat flow does not happen naturally, which is what the word “spontaneous” implies. In truth, the rule of heat flowing from high-temperature to cold-temperature applies to heat pumps as well, just in a way that is not obvious from first inspection. Mechanical heat pumps cause heat to be drawn from a cool object by placing an even cooler object (the *evaporator*) in direct contact with it. That heat is then transferred to a hot object by placing an even hotter object (the *condenser*) in direct contact with it. Heat is moved against the natural (spontaneous) direction of flow from the evaporator to the condenser by means of mechanical compression and expansion of a refrigerant fluid.

²⁶In this context, we are using the word “radiation” in a very general sense, to mean thermal energy radiated away from the hot source via photons. This is quite different from nuclear radiation, which is what some may assume this term means upon first glance.

Radiation

If you have ever experienced the immediate sensation of heat from a large fire or explosion some distance away, you know how *radiation* works to transfer thermal energy. Radiation is also the method of heat transfer experienced in the Earth's receiving of heat from the Sun (and also the mechanism of heat loss from Earth to outer space). Radiation is the least efficient of the three heat transfer mechanisms. It may be quantified by the Stefan-Boltzmann Law, which states the rate of heat lost by an object ($\frac{dQ}{dt}$) is proportional to the *fourth power* of its absolute temperature, and directly proportional to its radiating area:

$$\frac{dQ}{dt} = e\sigma AT^4$$

Where,

$\frac{dQ}{dt}$ = Radiant heat loss rate (watts)

e = Emissivity factor (unitless)

σ = Stefan-Boltzmann constant (5.67×10^{-8} W / m² · K⁴)

A = Surface area (square meters)

T = Absolute temperature (Kelvin)

Here is one of the scientific applications where temperature expressed in *absolute* units is truly necessary. Radiant energy is a direct function of molecular motion, and so we would logically expect objects to radiate energy at any temperature above absolute zero. The temperature value used in this formula *must* be in units of Kelvin²⁷ in order for the resulting $\frac{dQ}{dt}$ value to be correct. If degrees Celsius were used for T instead of Kelvin, the formula would predict zero thermal radiation at the freezing point of water (0 °C) and *negative* radiation at any temperature below freezing, which is not true. Remember that the “zero” points of the Celsius and Fahrenheit scales were arbitrarily set by the inventors of those scales, but that the “zero” points of the Kelvin and Rankine scales reflect a fundamental limit of nature.

The emissivity factor varies with surface finish and color, ranging from one (ideal) to zero (no radiation possible). Dark-colored, rough surfaces offer the greatest emissivity factors, which is why heater elements and radiators are usually painted black. Shiny (reflective), smooth surfaces offer the least emissivity, which is why thermally insulating surfaces are often painted white or silver.

Like all heat-transfer modes, radiation is two-way. Objects may emit energy in the form of radiation, and they may also receive energy in the form of radiation. Everyone knows white-colored shirts are cooler than black-colored shirts worn on a hot, sunny day – this is an example of how emissivity affects heat absorption by radiant transfer. A black-colored shirt (high emissivity value) enhances the receiving of radiant energy by your body from the sun. What is not as obvious, though, is that a white-colored shirt will keep you warmer than a black-colored shirt on a cold, dark day because that same decreased emissivity inhibits body heat *loss* by radiation. Thus, high-emissivity objects both heat *and* cool more readily by radiation than low-emissivity objects.

²⁷Or in degrees Rankine, provided a suitably units-corrected value for the Stefan-Boltzmann constant were used.

Conduction

If you have ever accidentally touched a hot iron or stove heating element, you possess a very vivid recollection of heat transfer through *conduction*. In conduction, fast-moving molecules in the hot substance transfer some of their kinetic energy to slower-moving molecules in the cold substance. Since this transfer of energy requires collisions between molecules, it only applies when the hot and cold substances directly contact each other.

Perhaps the most common application of heat conduction in industrial processes is through the walls of a furnace or some other enclosure containing an extreme temperature. In such applications, the desire is usually to *minimize* heat loss through the walls, so those walls will be “insulated” with a substance having poor thermal conductivity.

Conductive heat transfer rate is proportional to the difference in temperature between the hot and cold points, the area of contact, the distance of heat travel from hot to cold, and the thermal conductivity of the substance:

$$\frac{dQ}{dt} = \frac{kA\Delta T}{l}$$

Where,

$\frac{dQ}{dt}$ = Conductive heat transfer rate

k = Thermal conductivity


A = Surface area

ΔT = Difference of temperature between “hot” and “cold” sides

l = Length of heat flow path from “hot” to “cold” side

Note the meaning of “ ΔT ” in this context: it refers to the *difference* in temperature between two different locations in a system. Sometimes the exact same symbology (“ ΔT ”) refers to a *change in temperature over time* in the study of thermodynamics. Unfortunately, the only way to distinguish one meaning of ΔT from the other is by context.

An illustration showing heat conduction through a wall gives context to the variables in the previous equation. As we see here, A refers to the surface area of the wall, ΔT refers to the difference of temperature between either surface of the wall, and l refers to the thickness of the wall:



In the United States, a common measure of insulating ability used for the calculation of conductive heat loss in shelters is the *R-value*. The greater the R-value of a thermally insulating material, the less conductive it is to heat (lower k value). “R-value” mathematically relates to k and l by the following equation:

$$R = \frac{l}{k}$$

Rearranging this equation, we see that $l = kR$, and this allows us to substitute kR for l in the conduction heat equation, then cancel the k terms:

$$\frac{dQ}{dt} = \frac{kA\Delta T}{kR}$$

$$\frac{dQ}{dt} = \frac{A\Delta T}{R}$$

R is always expressed in the compound unit of square feet · hours · degrees Fahrenheit per BTU. This way, with a value for area expressed in square feet and a temperature difference expressed in degrees Fahrenheit, the resulting heat transfer rate ($\frac{dQ}{dt}$) will naturally be in units of BTU per hour, which is the standard unit in the United States for expressing heat output for combustion-type heaters. Dimensional analysis shows how the units cancel to yield a heat transfer rate in BTUs per hour:

$$\frac{[\text{BTU}]}{[\text{h}]} = \frac{[\text{ft}^2][^\circ\text{F}]}{\frac{[\text{ft}^2][\text{h}][^\circ\text{F}]}{[\text{BTU}]}}$$

The utility of R-value ratings may be shown by a short example. Consider a contractor trailer, raised up off the ground on a mobile platform, with a total skin surface area of 2400 square feet (walls, floor, and roof) and a uniform R-value of 4 for all surfaces. If the trailer's internal temperature must be maintained at 70 degrees Fahrenheit while the outside temperature averages 40 degrees Fahrenheit, the required output of the trailer's heater will be:

$$\frac{dQ}{dt} = \frac{(2400 \text{ ft}^2)(30^\circ \text{ F})}{4 \text{ ft}^2 \cdot \text{h} \cdot ^\circ \text{F}/\text{BTU}} = 18000 \text{ BTU per hour}$$

If the trailer's heater is powered by propane and rated at 80% efficiency (requiring 22500 BTU per hour of fuel heating value to produce 18000 BTU per hour of heat transfer into the trailer), the propane usage will be just over one pound per hour, since propane fuel has a heating value of 21700 BTU per pound.

Convection

Most industrial heat-transfer processes occur through *convection*, where a moving fluid acts as an intermediary substance to transfer heat from a hot substance (heat *source*) to a cold substance (heat *sink*). Convection may be thought of as two-stage heat conduction on a molecular scale: fluid molecules come into direct contact with a hot object and absorb heat from that object through conduction, then those molecules later release that heat energy through conduction by direct contact with a cooler object. If the fluid is recycled in a piping loop, the two-stage conduction process repeats indefinitely, individual molecules heating up as they absorb heat from the heat source and then cooling down as they release heat to the heat sink.

Special process devices called *heat exchangers* perform this heat transfer function between two different fluids, the two fluids circulating past each other on different sides of tube walls. A simple example of a heat exchanger is the radiator connected to the engine of an automobile, being a water-to-air exchanger, the engine's hot water transferring heat to cooling air entering the grille of the car as it moves.

Another example of a liquid-to-air heat exchanger is the *condenser* on a heat pump, refrigerator, or air conditioner, a photograph appearing here:




Another common style of heat exchanger works to transfer heat between two liquids. A small example of this design used to transfer heat from a boat engine is shown here:



The purpose for this heat exchanger is to exchange heat between the liquid coolant of the boat engine and sea water, the latter being quite corrosive to most metals. An engine would soon be damaged if sea water were used directly as the coolant fluid, and so heat exchangers such as this provide a means to release excess heat to the sea without subjecting the engine block to undue corrosion. The heat exchanger, of course, *does* suffer from the corrosive effects of sea water, but at least it is less expensive and more convenient to replace than an entire engine when it reaches the end of its service life.


This marine engine heat exchanger is an example of a *shell-and-tube* design, where one fluid passes inside small tubes and a second fluid passes outside those same tubes, the tube bundle being contained in a shell. The interior of such an exchanger looks like this when cut away:



The tubes of this particular heat exchanger are made of copper, a metal with extremely high thermal conductivity (k), to facilitate conductive heat transfer.

Liquid-to-liquid heat exchangers are quite common in industry, where a set of tubes carry one process liquid while a second process liquid circulates on the outside of those same tubes. The metal walls of the tubes act as heat transfer areas for conduction to occur. Metals such as copper with very high k values (very low R values) encourage heat transfer, while long lengths of tube ensure ample surface area for heat exchange.


A common application of liquid-to-liquid heat exchangers is in exothermic (heat-releasing) chemical reaction processes where the reactants must be pre-heated before entering a reaction vessel (“reactor”). Since the chemical reaction is exothermic, the reaction itself may be used as the heat source for pre-heating the incoming feed. A simple P&ID shows how a heat exchanger accomplishes this transfer of heat:



Another industrial application of heat exchangers is in *distillation* processes, where mixed components are separated from each other by a continuous process of boiling and condensation. Alcohol purification is one example of distillation, where a mixture of alcohol and water are separated to yield a purer (higher-percentage) concentration of alcohol. Distillation (also called *fractionation*) is a very energy-intensive²⁸ process, requiring great inputs of heat to perform the task of separation. Any method of energy conservation typically yields significant cost savings in a distillation process, and so we often find heat exchangers used to transfer heat from outgoing (distilled, or fractionated) products to the incoming feed mixture, pre-heating the feed so that less heat need be added to the distillation process from an external source.

²⁸Jim Cahill of Emerson wrote in April 2010 (“Reducing Distillation Column Energy Usage” Emerson Process Expert weblog) about a report estimating distillation column energy usage to account for approximately 6% of the total energy used in the United States. This same report tallied the number of columns in US industry to be approximately 40000 total, accounting for about 19% of all energy used in manufacturing processes!


The following P&ID shows a simple distillation process complete with heat exchangers for reboiling (adding heat to the bottom of the distillation column), condensing (extracting heat from the “overhead” product at the top of the column), and energy conservation (transferring heat from the hot products to the incoming feed):



Distillation “columns” (also called *fractionating towers* in the industry) are tall vessels containing sets of “trays” where rising vapors from the boiling process contact falling liquid from the condensing process. Temperatures increase toward the bottom of the column, while temperatures decrease toward the top. In this case, steam through a “reboiler” drives the boiling process at the bottom of the column (heat input), and cold water through a “condenser” drives the condensing process at the top of the column (heat extraction). Products coming off the column at intermediate points are hot enough to serve as pre-heating flows for the incoming feed. Note how the “economizing” heat exchangers expose the cold feed flow to the cooler Product A before exposing it to the warmer Product B, and then finally the warmest “Bottoms” product. This sequence of cooler-to-warmer maximizes the efficiency of the heat exchange process, with the incoming feed flowing past products of increasing temperature as it warms up to the necessary temperature for distillation entering the


column.

Some heat exchangers transfer heat from hot gases to cool(er) liquids. An example of this type of heat exchanger is the construction of a steam boiler, where hot combustion gases transfer heat to water flowing inside metal tubes:



Here, hot gases from the combustion burners travel past the metal “riser” tubes, transferring heat to the water within those tubes. This also serves to illustrate an important convection phenomenon: a *thermal siphon* (often written as *thermosiphon*). As water heats in the “riser” tubes, it becomes less dense, producing less hydrostatic pressure at the bottom of those tubes than the colder water in the “downcomer” tubes. This difference of pressure causes the colder water in the downcomer tubes to flow down to the mud drum, and hot water in the riser tubes to flow up to the steam drum. This natural *convection current* will continue as long as heat is applied to the riser tubes by the burners, and an unobstructed path exists for water to flow in a loop.

Natural convection also occurs in heated air, such as in the vicinity of a lit candle:



This thermally forced circulation of air helps *convect* heat from the candle to all other points within the room it is located, by carrying heated air molecules to colder objects.

2.10.5 Specific heat and enthalpy

Earlier, we saw how units of heat measurement were defined in terms of the amount of energy gain or loss required to alter the temperature of a water sample by one degree. In the case of the *calorie*, it was the amount of heat gain/loss required to heat/cool one gram of water one degree Celsius. In the case of the *BTU*, it was the amount of heat gain/loss required to heat/cool one pound of water one degree Fahrenheit.

As one might expect, one heat unit might be similarly defined as the amount of heat gain or loss to alter the temperature one-half of a degree for twice as much water, or two degrees for half as much water. We could express this as a proportionality:

$$Q \propto m\Delta T$$

Where,

Q = Heat gain or loss

m = Mass of sample

ΔT = Temperature change (rise or fall) over time

The next logical question to ask is, “How does the relationship between heat and temperature change work for substances other than water?” Does it take the same amount of heat to change the temperature of one gram of *iron* by one degree Celsius as it does to change the temperature of one gram of *water* by one degree Celsius? The answer to this question is a resounding *no!* Different substances require vastly different amounts of heat gain/loss to alter their temperature by the same degree, even when the masses of those substances happen to be identical.

We have a term for this ability to absorb or release heat, called *heat capacity* or *specific heat*, symbolized by the variable c . Thus, our heat/mass/temperature change relationship may be described as a true formula instead of a mere proportionality:

$$Q = mc\Delta T$$

Where,

Q = Heat gain or loss (metric calories or British BTU)

m = Mass of sample (metric grams or British pounds)

c = Specific heat of substance

ΔT = Temperature change (metric degrees Celsius or British degrees Fahrenheit)

Pure water, being the standard by which all other substances are measured, has a specific heat value of 1. The smaller the value for c , the less heat gain or loss is required to alter the substance’s temperature by a set amount. That substance (with a low value of c) has a low “heat capacity” because each degree of temperature rise or fall represents a relatively small amount of energy gained or lost. Substances with low c values are easy to heat and cool, while substances having high c values require much heat in order to alter their temperatures, assuming equal masses.

A table of specific heat values (at room temperature, 25 degrees Celsius²⁹) for common substances appears here:

Substance	Specific heat value (c) cal/g·°C or BTU/lb·°F
Aluminum (solid)	0.215
Iron (solid)	0.108
Copper (solid)	0.092
Lead (solid)	0.031
Ice (solid)	0.50
Water (liquid)	1.00
Methanol (liquid)	0.609
Ethanol (liquid)	0.587
Acetone (liquid)	0.521
Hydrogen (gas)	3.41
Helium (gas)	1.24
Nitrogen (gas)	0.249
Oxygen (gas)	0.219
Steam (gas)	0.476


If a liquid or a gas is chosen for use as a coolant (a substance to efficiently convect heat away from an object), greater values of c are better. Water is one of the best liquid coolants with its relatively high c value of one: it has more capacity to absorb heat than other liquids, for the same rise in temperature. The ideal coolant would have an infinite c value, being able to absorb an infinite amount of heat without itself rising in temperature at all.

As you can see from the table, the light gases (hydrogen and helium) have extraordinarily high c values. Consequently, they function as excellent media for convective heat transfer. This is why large electric power generators often use hydrogen gas as a coolant: hydrogen has an amazing ability to absorb heat from the wire windings of a generator without rising much in temperature. In other words, hydrogen absorbs a lot of heat while still remaining “cool” (i.e. remains at a low temperature). Helium, although not quite as good a coolant as hydrogen, has the distinct advantage of being chemically inert (non-reactive), in stark contrast to hydrogen’s extreme flammability. Some nuclear reactors use helium gas as a coolant rather than a liquid such as water or molten sodium metal.

Lead has an extraordinarily low c value, being a rather “easy” substance to heat up and cool down. Anyone who has ever cast their own lead bullets for a firearm knows how quickly a new lead bullet cools off after being released from the mold, especially if that same person has experience casting other metals such as aluminum.

²⁹An important detail to note is that specific heat does *not* remain constant over wide temperature changes. This complicates calculations of heat required to change the temperature of a sample: instead of simply multiplying the temperature change by mass and specific heat ($Q = mc\Delta T$ or $Q = mc[T_2 - T_1]$), we must *integrate* specific heat over the range of temperature ($Q = m \int_{T_1}^{T_2} c dT$), summing up infinitesimal products of specific heat and temperature change ($c dT$) over the range starting from temperature T_1 through temperature T_2 then multiplying by the mass to calculate total heat required. So, the specific heat values given for substances at 25 °C only hold true for relatively small temperature changes deviating from 25 °C. To accurately calculate heat transfer over a large temperature change, one must incorporate values of c for that substance at different temperatures along the expected range.

Numerical examples are helpful to better understand specific heat. Consider a case where a copper pot filled with water receives heat from a small gas burner operating at an output of 5000 BTU per hour (350 calories per second):



A reasonable question to ask would be, “How much will the temperature of this water-filled pot rise after 40 seconds of heating?” With the burner’s heat output of 350 calories per second and a heating time of 40 seconds, we may assume³⁰ the amount of heat absorbed by the water-filled pot will be the simple product of heat rate times time:

$$Q = \left(\frac{dQ}{dt} \right) t = \left(\frac{350 \text{ cal}}{\text{s}} \right) 40 \text{ s} = 14000 \text{ calories}$$

This amount of heat not only goes into raising the temperature of the water, but it also raises the temperature of the copper pot. Each substance (water, copper) has its own specific heat and mass values (c and m), but they will share the same temperature rise (ΔT), so we must sum their heats as follows:

$$Q_{total} = Q_{pot} + Q_{water}$$

$$Q_{total} = m_{pot}c_{pot}\Delta T + m_{water}c_{water}\Delta T$$

Since both the pot and the water start at the same temperature and end at the same temperature, ΔT is a common variable to both terms and may therefore be factored out:

$$Q_{total} = (m_{pot}c_{pot} + m_{water}c_{water})\Delta T$$

³⁰In reality, the amount of heat actually absorbed by the pot will be less than this, because there will be heat losses from the warm pot to the surrounding (cooler) air. However, for the sake of simplicity, we will assume *all* the burner’s heat output goes into the pot and the water it holds.

Solving this equation for temperature rise, we get:


$$\Delta T = \frac{Q_{total}}{m_{pot}c_{pot} + m_{water}c_{water}}$$

$$\Delta T = \frac{14000 \text{ cal}}{(1100 \text{ g})(0.092 \frac{\text{cal}}{\text{g}^\circ\text{C}}) + (3700 \text{ g})(1 \frac{\text{cal}}{\text{g}^\circ\text{C}})}$$

$$\Delta T = 3.68 \text{ }^\circ\text{C}$$

So, if the water and pot began at a temperature of 20 degrees Celsius, they will be at a temperature of 23.68 degrees Celsius after 40 seconds of heating over this small burner.

Another example involves the mixing of two substances at different temperatures. Suppose a heated mass of iron drops into a cool container³¹ of water. Obviously, the iron will lose heat energy to the water, causing the iron to decrease in temperature while the water rises in temperature. Suppose the iron's mass is 100 grams, and its original temperature is 65 degrees Celsius. Suppose the water's mass is 500 grams, and its original temperature is 20 degrees Celsius:



³¹We will assume for the sake of this example that the container holding the water is of negligible mass, such as a Styrofoam cup. This way, we do not have to include the container's mass or its specific heat into the calculation.

What will the equilibrium temperature be after the iron falls into the water and both their temperatures equalize? We may solve this by setting two heat equations equal to each other³²: the heat lost by the iron and the heat gained by the water, with the final equilibrium temperature being T :

$$Q_{iron} = Q_{water}$$

$$m_{iron}c_{iron}\Delta T_{iron} = m_{water}c_{water}\Delta T_{water}$$

$$m_{iron}c_{iron}(65^{\circ}\text{C} - T) = m_{water}c_{water}(T - 20^{\circ}\text{C})$$

Note how the ΔT term is carefully set up for each side of the equation. In order to make the iron's heat loss a positive value and the water's heat gain a positive value, we must ensure the quantity within each set of parentheses is positive. For the iron, this means ΔT will be 65 degrees minus the final temperature. For the water, this means ΔT will be the final temperature minus its starting temperature of 20 degrees.

In order to solve for the final temperature (T), we must distribute the terms, collecting all T -containing terms to one side of the equation, then factor and isolate T :

$$m_{iron}c_{iron}(65) - m_{iron}c_{iron}T = m_{water}c_{water}T - m_{water}c_{water}(20)$$

$$m_{iron}c_{iron}(65) + m_{water}c_{water}(20) = m_{iron}c_{iron}T + m_{water}c_{water}T$$

$$m_{iron}c_{iron}(65) + m_{water}c_{water}(20) = T(m_{iron}c_{iron} + m_{water}c_{water})$$

$$T = \frac{m_{iron}c_{iron}(65) + m_{water}c_{water}(20)}{m_{iron}c_{iron} + m_{water}c_{water}}$$

$$T = \frac{(100 \text{ g})(0.108 \text{ cal/g}^{\circ}\text{C})(65^{\circ}\text{C}) + (500 \text{ g})(1 \text{ cal/g}^{\circ}\text{C})(20^{\circ}\text{C})}{(100 \text{ g})(0.108 \text{ cal/g}^{\circ}\text{C}) + (500 \text{ g})(1 \text{ cal/g}^{\circ}\text{C})}$$


$$T = 20.95^{\circ}\text{C}$$

Thus, the iron's temperature falls from 65 degrees Celsius to 20.95 degrees Celsius, while the water's temperature rises from 20 degrees Celsius to 20.95 degrees Celsius. The water's tremendous specific heat value compared to the iron (nearly 10 times as much!), as well as its superior mass (5 times as much) results in a much larger temperature change for the iron than for the water.

³²An alternative way to set up the problem would be to calculate ΔT for each term as $T_{final} - T_{start}$, making the iron's heat loss a negative quantity and the water's heat gain a positive quantity, in which case we would have to set up the equation as a zero-sum balance, with $Q_{iron} + Q_{water} = 0$. I find this approach less intuitive than simply saying the iron's heat loss will be equal to the water's heat gain, and setting up the equation as two positive values equal to each other.

An analogy to help grasp the concept of specific heat is to imagine heat as a fluid³³ that may be “poured” into vessels of different size, those vessels being objects or substances to be heated. The amount of liquid held by any vessel represents the total amount of thermal energy, while the *height* of the liquid inside any vessel represents its temperature:

Fluid analogy for heat and temperature



The factor determining the relationship between liquid volume (heat) and liquid height (temperature) is of course the cross-sectional area of the vessel. The wider the vessel, the more heat will be required to “fill” it up to any given temperature. In this analogy, the area of the vessel is analogous to the term mc : the product of mass and specific heat. Objects with larger mass require more heat to raise their temperature to any specific point, specific heats being equal. Likewise, objects with large specific heat values require more heat to raise their temperature to any specific point, masses being equal.

In the first numerical calculation example where we determined the temperature of a pot of water after 40 seconds of heating, the analogous model would be to determine the height of liquid in a vessel after pouring liquid into it for 40 seconds at a fixed rate. A model for the second numerical example would be to calculate the equilibrium height (of liquid) after connecting two vessels together at their bottoms with a tube. Although the liquid heights of those vessels may be different at first, the levels will equalize after time by way of liquid passing through the tube from the higher-level vessel to the lower-level vessel.

³³This is not far from the hypotheses of eighteenth-century science, where heat was thought to be an invisible fluid called *caloric*.

Many industrial processes use fluids to convectively transfer thermal energy from one object (or fluid) to another. In such applications, it is important to know how much thermal energy will be carried by a specific quantity of that fluid over a specified temperature drop. One common way to express this quantity is called *enthalpy*. Enthalpy is the amount of heat lost by a unit mass (one gram metric, or one pound British) of a substance as it cools from a given temperature all the way down to the freezing point of water (0 degrees Celsius, or 32 degrees Fahrenheit). In other words, enthalpy is a measure³⁴ of a substance's thermal energy using the freezing temperature of water as a reference. A sample of water at a temperature of 125 degrees Fahrenheit, for example, has an enthalpy of 93 BTU per pound (or 93 calories per gram), because 93 BTU of thermal energy would be lost if one pound of that water happened to cool from its given temperature (125 °F) down to 32 °F:

$$Q = mc\Delta T$$

$$Q = (1 \text{ lb}) \left(1 \frac{\text{BTU}}{\text{lb}^\circ\text{F}} \right) (125^\circ\text{F} - 32^\circ\text{F})$$

$$Q = 93 \text{ BTU}$$

Even if the substance in question does not cool down to the freezing temperature of water, enthalpy is a useful figure for comparing the thermal energy “content” of hot fluids (per unit mass). For example, if one were given the enthalpy values for a substance before and after heat transfer, it would be easy to calculate the amount of heat transfer that transpired simply by subtracting those enthalpy values³⁵. If water at 125 °F has an enthalpy value of 93 BTU/lb and water at 170 °F has an enthalpy of value 138 BTU/lb, we may calculate the amount of heat needed to increase the temperature of a sample of water from 125 °F to 170 °F simply by subtracting 93 BTU/lb from 138 BTU/lb to arrive at 45 BTU/lb.

In this rather trivial example, it would have been just as easy for us to calculate the heat necessary to increase water's temperature from 125 °F to 170 °F by using the specific heat formula ($Q = mc\Delta T$)³⁶, and so it might appear as though the concept of enthalpy sheds no new light on the subject of heat transfer. However, the ability to calculate heat transfer based on a simple subtraction of enthalpy values proves quite useful in more complex scenarios where substances change phase, as we will see next.

³⁴A useful analogy for enthalpy is the *maximum available balance* of a bank account. Suppose you have a bank account with a minimum balance requirement of \$32 to maintain that account. Your maximum available balance at any time would be the total amount of money in that account minus \$32, or to phrase this differently your maximum available balance is the most money you may spend from this account while still keeping that account open. Enthalpy is much the same: the amount of thermal energy a sample may “spend” (i.e. lose) before its temperature reaches 32 degrees Fahrenheit.

³⁵Appealing to the *maximum available balance* analogy, if we compared the maximum available balance in your bank account before and after a transaction, we could determine how much money was deposited or withdrawn from your account simply by subtracting those two values.

³⁶Following the formula $Q = mc\Delta T$, we may calculate the heat as $(1)(1)(170 - 125) = 45$ BTU. This is obviously the same result we obtained by subtracting enthalpy values for water at 170 °F and 125 °F.

2.10.6 Phase changes

Scientists often speak of four *phases* of matter: *solid*, *liquid*, *gas* (or *vapor*), and *plasma*. Of these four, the first three are common to everyday life. Plasma is a phase of matter where the atoms of a gas are excited (energized) to the point where they become electrically ionized, such as neon gas in an electric tube light, or the gas comprising stars in space.

Phase changes are very important in thermodynamics, principally because energy transfer (heat loss or heat gain) must occur for a substance to change states, often with negligible change in temperature. To cite an example, consider the case of water (a liquid) turning into steam (a vapor) at atmospheric pressure. At sea level, this phase change will occur at a temperature of 100 degrees Celsius, or 212 degrees Fahrenheit. The amount of energy required to increase the temperature of water from ambient up to its boiling point is a simple function of the sample's mass and its original temperature. For instance, a sample of water 70 grams in mass starting at 24 degrees Celsius will require 5320 calories of heat to reach the boiling point:

$$Q = mc\Delta T$$

$$Q = (70 \text{ g}) \left(1 \frac{\text{cal}}{\text{g}^\circ\text{C}} \right) (100^\circ\text{C} - 24^\circ\text{C})$$

$$Q = 5320 \text{ cal}$$

However, actually boiling the 70 grams of water into 70 grams of steam (both at 100 degrees Celsius) requires a comparatively enormous input of heat: *37734 calories* – over seven times as much heat to turn the water to steam as what is required to warm the water to its boiling point! Furthermore, this additional input of 37734 calories does not increase the temperature of the water at all: the resulting steam is still at a temperature of (only) 100 degrees Celsius. If further heat is added to the 70 gram steam sample, its temperature will rise, albeit at a rate proportional to the value of steam's specific heat (0.476 calories per gram degree Celsius, or BTU per pound degree Fahrenheit).

What we see here is a fundamentally different phenomenon than we saw with specific heat. Here, we are looking at the thermal energy required to transition a substance from one phase to another, not to change its temperature. We call this quantity *latent heat* rather than *specific heat*, because no temperature change occurs³⁷. Conversely, if we allow the steam to condense back into liquid water, it must release the same 37734 calories of heat energy we invested in it turning the water into steam before it may cool at all below the boiling point (100 degrees Celsius).

Latent heat has the effect of greatly increasing a substance's enthalpy. Recall that “enthalpy” is the amount of heat lost by one pound (mass) of a substance if it happened to cool from its given temperature all the way down to the freezing temperature of water (0 °C, or 32 °F). Hot water has an enthalpy of 1 BTU/lb for every degree of temperature above freezing. Steam, however, possesses far greater enthalpy because of the latent heat released in the phase change from vapor to liquid before it releases heat as water cooling down to 32 °F.

³⁷The word “latent” refers to something with potential that is not yet realized. Here, heat exchange takes place without there being any realized change in temperature. By contrast, heat resulting in a temperature change ($Q = mc\Delta T$) is called *sensible heat*.

As with specific heat, there is a formula relating mass, latent heat, and heat exchange:

$$Q = mL$$

Where,

Q = Heat of transition required to completely change the phase of a sample (metric calories or British BTU)

m = Mass of sample (metric grams or British pounds)

L = Latent heat of substance

Each substance has its own set of latent heat values, one³⁸ for each phase-to-phase transition. Water, for example, exhibits a latent heat of vaporization (boiling/condensing) of 539.1 calories per gram, or 970.3 BTU per pound, at atmospheric pressure (boiling point = 100 °C = 212 °F). Water also exhibits a latent heat of fusion (melting/freezing) of 79.7 calories per gram, or 143.5 BTU per pound. Both figures are enormous compared to water's specific heat value of 1 calorie per gram-degree Celsius (or 1 BTU per pound-degree Fahrenheit³⁹): it takes only one calorie of heat to warm one gram of water one degree Celsius, but it takes *539.1 calories* of heat to boil that same gram of water into one gram of steam, and *79.7 calories* of heat to melt one gram of ice into one gram of water. The lesson here is simple: phase changes involve huge amounts of heat.

A table showing various latent heats of vaporization (all at room temperature, 70 degrees Fahrenheit) for common industrial fluids appears here, contrasted against their specific heat values (as liquids). In each case you will note how much larger L is than c :


Fluid (@ 70 °F)	$L_{vaporization}$, BTU/lb	$L_{vaporization}$, cal/g	c_{liquid}
Water	970.3	539.1	1
Ammonia	508.6	282.6	1.1
Carbon dioxide	63.7	35.4	0.66
Butane	157.5	87.5	0.56
Propane	149.5	83.06	0.6

One of the most important, and also non-intuitive, consequences of latent heat is the relative stability of temperature during the phase-change process. Referencing the table of latent heats of vaporization, we see how much more heat is needed to boil a liquid into a vapor than is needed to warm that same liquid by one degree of temperature. During the process of boiling, all heat input to the liquid goes into the task of phase change (latent heat) and none of it goes into increased temperature. In fact, until all the liquid has been vaporized, the liquid's temperature *cannot* rise above its boiling point! The requirement of heat input to vaporize a liquid forces temperature to stabilize (not rise further) until *all* the liquid has evaporated from the sample.

³⁸Latent heat of vaporization also varies with pressure, as different amounts of heat are required to vaporize a liquid depending on the pressure that liquid is subject to. Generally, increased pressure (increased boiling temperature) results in less latent heat of vaporization.

³⁹The reason specific heat values are identical between metric and British units, while latent heat values are not, is because latent heat does not involve temperature change, and therefore there is one less unit conversion taking place between metric and British when translating latent heats. Specific heat in both metric and British units is *defined* in such a way that the three different units for heat, mass, and temperature all cancel each other out. With latent heat, we are only dealing with mass and heat, and so we have a proportional conversion of $\frac{5}{9}$ or $\frac{9}{5}$ left over, just the same as if we were converting between degrees Celsius and Fahrenheit alone.


If we take a sample of ice and add heat to it over time until it melts, warms, boils, and then becomes steam, we will notice a temperature profile that looks something like this:



The flat areas of the graph during the melting and boiling periods represents times where the sample's temperature does not change at all, but where all heat input goes into the work of changing the sample's phase. Only where we see the curve rising does the temperature change. So long as there is a *mixture* of different phases, the temperature remains "locked" at one value. Only when there is a single phase of material is the temperature "allowed" to rise or fall.

The sloped areas of the graph reveal the specific heat of the substance in each particular phase. Note how the liquid (water) portion of the graph has a relatively shallow slope, due to the specific heat value (c) of water being equal to 1. Both the ice and the steam portions of the graph have steeper slopes because both of those phases possess smaller values of specific heat ($c = 0.5$ and $c = 0.476$, respectively). The smaller the value of c , the more a sample's temperature will rise for any given input of thermal energy. For any given rate of heat transfer, smaller c values result in more rapid temperature changes.

We may employ our liquid-filled vessel analogy to the task of explaining latent heat. Any point of phase change is analogous to a point along the vessel's height equipped with a large expansion chamber, so that the vessel "acts" as if its area were much larger at one point, requiring *much* more fluid volume (heat) to change height (temperature) past that one point:



Liquid poured into this vessel will fill it at a rate proportional to the volume added and inversely proportional to the vessel's cross-sectional area at the current liquid height. As soon as the liquid level reaches the expansion chamber, a great deal more liquid must be added to cause level to increase, since this chamber must completely fill before the liquid level may rise above it. Once that happens, the liquid level rises at a different rate with addition introduced volume, because now the phase is different (with a different specific heat value).

Remember that the filling of a vessel with liquid is merely an analogy for heat and temperature, intended to provide an easily visualized process mimicking another process not so easily visualized. The important concept to realize with latent heat and phase change is that it constitutes a discontinuity in the temperature/heat function for any given substance.

A vivid demonstration of this phenomenon is to take a paper⁴⁰ cup filled with water and place it in the middle of a roaring fire⁴¹. "Common sense" might tell you the paper will burn through with the fire's heat, so that the water runs out of the cup through the burn-hole. This does not happen, however. Instead, the water in the cup will rise in temperature until it boils, and there it will maintain that temperature no matter how hot the fire burns. The boiling point of water happens to be substantially below the burning point of paper, and so the boiling water keeps the paper cup too cool to burn. As a result, the paper cup remains intact so long as water remains in the cup. The *rim* of the cup above the water line will burn up because the steam does not have the same temperature-stabilizing effect as the water, leaving a rimless cup that grows shorter as the water boils away.

⁴⁰Styrofoam and plastic cups work as well, but paper exhibits the furthest separation between the boiling point of water and the burning point of the cup material, and it is usually thin enough to ensure good heat transfer from the outside (impinging flame) to the inside (water).

⁴¹This is a lot of fun to do while camping!

The point at which a pure substance changes phase not only relates to temperature, but to pressure as well. We may speak casually about the boiling point of water being 100 degrees Celsius (212 degrees Fahrenheit), but that is only if we assume the water and steam are at atmospheric pressure (at sea level). If we reduce the ambient air pressure⁴², water will boil at a lesser temperature. Anyone familiar with cooking at high altitudes knows you must generally cook for longer periods of time at altitude, because the decreased boiling temperature of water is not as effective for cooking. Conversely, anyone familiar with *pressure cooking* (where the cooking takes place inside a vessel pressurized by steam) knows how comparatively little cooking time is required because the pressure raises water's boiling temperature. In either of these scenarios, where pressure influences⁴³ boiling temperature, the latent heat of water acts to hold the boiling water's temperature stable until all the water has boiled away. The only difference is the temperature at which the water begins to boil (or when the steam begins to condense).

Many industrial processes use boiling liquids to convectively transfer heat from one object (or fluid) to another. In such applications, it is important to know how much heat will be carried by a specific quantity of the vapor as it condenses into liquid over a specified temperature drop. The quantity of *enthalpy* (heat content) used for rating the heat-carrying capacity of liquids applies to condensing vapors as well. Enthalpy is the amount of heat lost by a unit mass (one gram metric, or one pound British) of the fluid as it cools from a given temperature all the way down to the freezing point of water (0 degrees Celsius, or 32 degrees Fahrenheit)⁴⁴. When the fluid's initial state is vapor, and it condenses into liquid as it cools down to the reference temperature (32 °F), the heat content (enthalpy) is not just a function of specific heat, but also of latent heat.

Water at its atmospheric boiling point has an enthalpy of approximately 180 BTU per pound. Steam at atmospheric pressure and 212 °F, however, has an enthalpy of about *1150 BTU per pound*: more than six times as much heat as water at the same temperature. 970 of that 1150 BTU/lb is due to the phase change from steam to water, while the rest is due to water's specific heat as it cools from 212 °F to 32 °F.

Many technical reference books contain a set of data known as a *steam table* showing various properties of steam at different temperatures and pressures. Enthalpy is one of the most important parameters given in a steam table, showing how much available energy resides in steam under different pressure and temperature conditions. For this reason, enthalpy is sometimes referred to as *total heat* (h_g). Steam tables also show saturation temperature (the condensing temperature for steam at that pressure) and steam density.

⁴²This may be done in a vacuum jar, or by traveling to a region of high altitude where the ambient air pressure is less than at sea level.

⁴³The mechanism of this influence may be understood by considering what it means to boil a liquid into a vapor. Molecules in a liquid reside close enough to each other that they cohere, whereas molecules in a vapor or gas are relatively far apart and act as independent objects. The process of boiling requires that cohesion between liquid molecules be broken, so the molecules may drift apart. Increased pressure encourages cohesion in liquid form by helping to hold the molecules together, while decreased pressure encourages the separation of molecules into a vapor/gas.

⁴⁴As mentioned previously, a useful analogy for enthalpy is the *maximum available balance* for a bank account with a \$32 minimum balance requirement: that is, how much money may be spent from that account without closing it out.

If the vapor in question is at a temperature greater than its boiling point at that pressure, the vapor is said to be *superheated*. The enthalpy of superheated vapor comes from three different heat-loss mechanisms:

- Cooling the vapor down to its condensing temperature (specific heat of vapor)
- Phase-changing from vapor to liquid (latent heat of phase change)
- Cooling the liquid down to the reference temperature (specific heat of liquid)

Using steam as the example once more, a sample of superheated steam at 500 °F and atmospheric pressure (boiling point = 212 °F) has an enthalpy of approximately 1287 BTU per pound. We may calculate the heat lost by one pound of this superheated steam as it cools from 500 °F to 32 °F in each of the three steps previously described. Here, we will assume a specific heat for steam of 0.476, a specific heat for water of 1, and a latent heat of vaporization for water of 970:

Heat loss mechanism	Formula	Quantity
Cooling vapor	$Q = mc\Delta T$	$(1)(0.476)(500 - 212) = 137 \text{ BTU}$
Phase change	$Q = mL$	$(1)(970) = 970 \text{ BTU}$
Cooling liquid	$Q = mc\Delta T$	$(1)(1)(212 - 32) = 180 \text{ BTU}$
TOTAL		1287 BTU

Enthalpy values are very useful⁴⁵ in steam engineering to predict the amount of thermal energy delivered to a load given the steam's initial temperature, its final (cooled) temperature, and the mass flow rate. Although the definition of enthalpy – where we calculate heat value by supposing the vapor cools all the way down to the *freezing point* of water – might seem a bit strange and impractical (how common is it for steam to lose so much heat to a process that it reaches freezing temperature?), it is not difficult to shift the enthalpy value to reflect a more practical final temperature. Since we know the specific heat of liquid water is very nearly one, all we have to do is offset the enthalpy value by the amount that the final temperature differs from freezing in order to calculate how much heat the steam will lose (per pound) to its load⁴⁶.


Furthermore, the *rate* at which heat is delivered to a substance by steam (or conversely, the rate at which heat is required to boil water into steam) may be easily calculated if we take this heat value in units of BTU per pound and multiply it by the mass flow rate in pounds per minute: as the unit of “pound” cancels in the multiplication, we arrive at a result for heat transfer rate in units of BTU per minute.

⁴⁵At first it may seem as though the enthalpy of steam is so easy to calculate it almost renders steam tables useless. If the specific heats of water and steam were constant, and the latent heat of vaporization for water likewise constant, this would be the case. However, both these values (c and L) are not constant, but rather change with pressure and with temperature. Thus, steam tables end up being quite valuable to engineers, allowing them to quickly reference heat content of steam across a broad range of pressures and temperatures without having to account for changing c and L values (performing integral calculus in the form of $Q = m \int_{T_1}^{T_2} c dT$ for specific heat) in their heat calculations.


⁴⁶This is not unlike calculating the voltage dropped across an electrical load by measuring the voltage at each of the load's two terminals with respect to *ground*, then subtracting those two measured voltage values. In this analogy, electrical “ground” is the equivalent of water at freezing temperature: a common reference point for energy level.

For example, suppose we were to employ the same 500 °F superheated steam used in the previous example to heat a flow of oil through a heat exchanger, with the steam condensing to water and then cooling down to 170 degrees Fahrenheit as it delivers heat to the flowing oil. Here, the steam's enthalpy value of 1287 BTU per pound may simply be offset by 138 (170 degrees minus 32 degrees) to calculate how much heat (per pound) this steam will deliver to the oil: $1287 - 138 = 1149$ BTU per pound:

Heat exchanger application



Thermal diagram



Here we see how 500 °F steam has an enthalpy (total heat) of 1287 BTU/lb, but since the steam does not in fact cool all the way down to 32 °F in the act of heating oil in the heat exchanger, we must subtract the enthalpy of the 170 °F water (138 BTU/lb) to determine⁴⁷ the amount of heat actually delivered to the oil by the steam (1149 BTU/lb). Calculating heat transfer rate is a simple matter of multiplying this heat per pound of steam by the steam's mass flow rate: for example, if the mass flow rate of this steam happened to be 2 pounds per minute, the heat transfer rate would be 2298 BTU per minute.

If we happen to be dealing with a situation where steam gives up some heat energy to a process fluid but not enough to cool to the point of condensation, all we need to do to calculate the amount of heat liberated by the superheated steam as it cools is subtract the enthalpy values between its hot and cool(er) states.


For example, suppose we have a heat-exchange process where superheated steam enters at 105 PSIG and 600 °F, exiting at 75 PSIG and 360 °F. The enthalpy of the steam under those two sets of conditions as given by a superheated steam table are 1328 BTU/lb and 1208 BTU/lb, respectively. Thus, the heat lost by the steam as it goes through this heat exchanger is the difference in enthalpy values: $1328 \text{ BTU/lb} - 1208 \text{ BTU/lb} = 120 \text{ BTU/lb}$. Once again, calculating heat transfer rate is a simple matter of multiplication: if the mass flow rate of this steam happened to be 80 pounds per hour, the heat transfer rate would be $120 \text{ BTU/lb} \times 80 \text{ lb/hr} = 9600 \text{ BTU/hr}$.

⁴⁷Applying the *maximum available balance* analogy to this scenario, it would be as if your bank account began with a maximum available balance of \$1287 and then finished with a maximum available balance of \$138 after an expenditure: the amount of money you spent is the different between the initial and final maximum available balances ($\$1287 - \$138 = \$1149$).

By encompassing both specific heat and latent heat into one quantity, enthalpy figures given in steam tables greatly simplify heat transfer calculations, as compared to evaluating specific heat and latent heat formulae ($Q = mc\Delta T$ and $Q = mL$, respectively) for water. Calculations based on steam tables are also more accurate than those derived from the specific and/or latent heat formulae, because steam tables take into account the changing values of c and L over wide temperature and pressure ranges. This is the power of empirical data: steam tables were developed by taking actual calorimetric measurements of steam under those temperature and pressure conditions, and as such are a *record* of water's true behavior rather than a *prediction* of water's theoretical behavior.

2.10.7 Phase diagrams and critical points

A comprehensive way of describing the relationship between pressure, temperature, and substance phase is with something called a *phase diagram*. With pressure shown on one axis, and temperature on the other, a phase diagram describes the various phases of a substance in possible equilibrium at certain pressure/temperature combinations.



This phase diagram (for water) illustrates some of the features common to all phase diagrams: curved lines define the boundaries between solid, liquid, and vapor phases; the point of intersection of these three curves is where the substance may exist in all three phases simultaneously (called the *triple point*⁴⁸) and points where a curve simply ends within the span of the graph indicate critical points, where the certain phases cease to exist.

The curved line from the triple point up and to the right defines the boundary between liquid water and water vapor. Each point on that line represents a set of unique pressure and temperature conditions for boiling (changing phase from liquid to vapor) or for condensation (changing phase from vapor to liquid). As you can see, increased pressure results in an increased boiling point (i.e. at higher pressures, water must be heated to greater temperatures before boiling may take place).

⁴⁸When H_2O is at its triple point, vapor (steam), liquid (water), and solid (ice) of water will co-exist in the same space. One way to visualize the triple point is to consider it the pressure at which the boiling and freezing temperatures of a substance become the same.

In fact, the whole concept of a singular boiling *point* for water becomes quaint in the light of a phase diagram: boiling is seen to occur over a wide range of temperatures⁴⁹, the exact temperature varying with pressure.

Something interesting happens when the temperature is raised above a certain value called the *critical temperature*. At this value (approximately 374 degrees Celsius for water), no amount of pressure will maintain it in a liquid state. Water, once heated beyond 374 degrees Celsius, is no longer a liquid and may only exist in a stable condition as a vapor. The *critical pressure* of any substance is the pressure value at the liquid/vapor boundary at the point of critical temperature.

A vivid example of critical temperature is this photograph of an ultra-high pressure storage vessel for oxygen gas, at a rocket engine testing facility:




The critical temperature for oxygen is 154.58 Kelvin, which is equal to -118.57 degrees Celsius or -181.43 degrees Fahrenheit. Since this pressure vessel is completely uninsulated, we know the temperature of the oxygen inside of it will be the same (or nearly the same) as ambient temperature, which is obviously much warmer than -118.57 °C. Since the oxygen's temperature is well above the critical temperature for the element oxygen, we may safely conclude that the oxygen inside this vessel must exist as a gas. Even at the extremely high pressure this vessel is capable of holding (15000 PSIG), the oxygen cannot liquefy.

The slightly curved line from the triple point up and to the left defines the boundary between solid ice and liquid water. As you can see, the near-vertical pitch of this curve suggests the freezing temperature of water is quite stable over a wide pressure range.

⁴⁹Anywhere between the triple-point temperature and the critical temperature, to be exact.

Carbon dioxide exhibits a different set of curves than water on its phase diagram, with its own unique critical temperature and pressure values:



Note how the triple-point pressure of carbon dioxide is well above ambient conditions on Earth. This means carbon dioxide is not stable in its liquid state unless put under substantial pressure (at least 60.4 PSIG). This is why solid carbon dioxide is referred to as *dry ice*: it does not liquefy with the application of heat, rather it *sublimates* directly into its vapor phase.

Another interesting difference between the carbon dioxide and water phase diagrams is the slope of the solid/liquid boundary line. With water, this boundary drifts to the left (lower temperature) as pressure increases. With carbon dioxide, this boundary drifts to the right (higher temperature) as pressure increases. Whether the fusion temperature increases or decreases with increasing pressure marks whether that substance contracts or expands as it transitions from liquid to solid. Carbon dioxide, like most pure substances, contracts to a smaller volume when it goes from liquid to solid, and its fusion curve drifts to the right as pressure increases. Water is unusual in this regard, expanding to a larger volume when freezing, and its fusion curve drifts to the left as pressure increases.

2.10.8 Saturated steam table

A saturated steam table shows temperatures and pressures for water at the liquid/vapor transition (i.e. points lying along the liquid/vapor interface shown in a phase change diagram), as well as enthalpy values for the water and steam under those conditions. The sensible heat of water is the amount of thermal energy per pound necessary to raise water's temperature from the freezing point to the boiling point. The latent heat of vapor is the amount of energy per pound necessary to convert water (liquid) into steam (vapor). The total heat is the enthalpy of steam (thermal energy per pound) between the listed condition in the table and the freezing temperature of water.

By definition a saturated steam table does *not* describe steam at temperatures greater than the boiling point. For such purposes, a *superheated steam table* is necessary.

Data for this saturated steam table was taken from *Thermal Properties of Saturated and Superheated Steam* by Lionel Marks and Harvey Davis, published in 1920 by Longmans, Green, and Company.

Saturated Steam Table

Temperature (Deg F)	Pressure (PSIA)	Sensible heat of liquid (BTU/lb)	Latent heat of vapor (BTU/lb)	Total heat (BTU/lb)
32	0.0886	0.00	1073.4	1073.4
40	0.1217	8.05	1068.9	1076.9
50	0.1780	18.08	1063.3	1081.4
60	0.2562	28.08	1057.8	1085.9
70	0.3626	38.06	1052.3	1090.3
80	0.505	48.03	1046.7	1094.8
90	0.696	58.00	1041.2	1099.2
100	0.946	67.97	1035.6	1103.6
110	1.271	77.94	1030.0	1108.0
120	1.689	87.91	1024.4	1112.3
130	2.219	97.89	1018.8	1116.7
140	2.885	107.87	1013.1	1121.0
150	3.714	117.86	1007.4	1125.3
160	4.737	127.86	1001.6	1129.5
170	5.992	137.87	995.8	1133.7
180	7.51	147.88	989.9	1137.8
190	9.34	157.91	983.9	1141.8
200	11.52	167.94	977.8	1145.8
210	14.13	177.99	971.6	1149.6
212	14.70	180.00	970.4	1150.4

Saturated Steam Table (continued)


Temperature (Deg F)	Pressure (PSIA)	Sensible heat of liquid (BTU/lb)	Latent heat of vapor (BTU/lb)	Total heat (BTU/lb)
220	17.19	188.1	965.2	1153.3
230	20.77	198.2	958.7	1156.9
240	24.97	208.3	952.1	1160.4
250	29.82	218.5	945.3	1163.8
260	35.42	228.6	938.4	1167.0
270	41.85	238.8	931.4	1170.2
280	49.18	249.0	924.3	1173.3
290	57.55	259.3	916.9	1176.2
300	67.00	269.6	909.5	1179.1
310	77.67	279.9	901.9	1181.8
320	89.63	290.2	894.2	1184.4
330	103.0	300.6	886.3	1186.9
340	118.0	311.0	878.3	1189.3
350	134.6	321.4	870.1	1191.5
360	153.0	331.9	861.8	1193.7
370	173.3	342.4	853.4	1195.8
380	195.6	352.9	844.8	1197.7
390	220.2	363.5	836.1	1199.6
400	247.1	374.1	827.2	1201.3
410	276.4	384.7	818.2	1202.9
420	308.4	395.4	809.0	1204.4
430	343.2	406.2	799.6	1205.8
440	380.8	417.0	790.1	1207.1

2.10.9 Thermodynamic degrees of freedom

If we look at the areas bounded by phase transition curves in a phase diagram (solid area, liquid area, and vapor area), we see that both pressure and temperature may change independent of one another. A vessel filled with liquid water, for instance, may be at 30 degrees Celsius and 2 atmospheres, or at 50 degrees Celsius and 2 atmospheres, or at 50 degrees Celsius and 1 atmosphere, all equally stable. A more technical way to state this is to say the liquid water has *two degrees of freedom*. Here, the word “degree” has a completely different meaning than it does when used to denote a unit of temperature or angle. In this context, “degree” may be thought of loosely as “dimension.” A cube has three physical dimensions, a square has two and a line has one. A point within a cube has three degrees of freedom (motion), while a point within a square only has two, and a point along a line only has one. Here, we use the word “degree” to denote the number of independent ways a system may change. For areas bounded by phase transition curves in a phase diagram, *pressure* and *temperature* are the two “free” variables, because within those bounded areas we may freely alter pressure without altering temperature, and vice-versa.

Such is not the case at any point lying along one of the phase transition curves. Any point on a curve is geometrically defined by a pair of coordinates, which means that for a two-phase mixture in equilibrium there will be exactly one temperature value valid for each unique pressure value. At any point along a phase transition curve, pressure and temperature are not independent variable, but rather are *related*. This means that for any single substance, there is only one degree of freedom along any point of a phase transition curve.

To illustrate this concept, suppose we equip a closed vessel containing water with both a thermometer and a pressure gauge. The thermometer measures the temperature of this water, while the pressure gauge measures the pressure of the water. A burner beneath the vessel adds heat to alter the water’s temperature, and a pump adds water to the vessel to alter the pressure inside:



So long as the water is all liquid (one phase), we may adjust its pressure and temperature independently. In this state, the system has two thermodynamic degrees of freedom.

However, if the water becomes hot enough to boil, creating a system of two phases in direct contact with each other (equilibrium), we will find that pressure and temperature become linked: one cannot alter one without altering the other. For a steam boiler, operation at a given steam

pressure thus *defines* the temperature of the water, and vice-versa. In a single-component, two-phase system, there is only one degree of thermodynamic freedom.

Our freedom to alter pressure and temperature becomes even more restricted if we ever reach the *triple point*⁵⁰ of the substance. For water, this occurs (only) at a pressure of -14.61 PSIG (0.006 atmospheres) and a temperature of 0.01 degrees Celsius: the coordinates where all three phase transition curves intersect on the phase diagram. In this state, where solid (ice), liquid (water), and vapor (steam) coexist, there are zero degrees of thermodynamic freedom. Both the temperature and pressure are *locked* at these values until one or more of the phases disappears.

The relationship between degrees of freedom and phases is expressed neatly by *Gibbs' Phase Rule* – the sum of phases and degrees of freedom equals the number of substances (“components”) plus two:

$$N_{\text{freedom}} + N_{\text{phase}} = N_{\text{substance}} + 2$$

We may simplify Gibbs' rule for systems of just one substance (1 “component”) by saying the number of degrees of freedom plus phases in direct contact with each other is always equal to three. So, a vessel filled with nothing but liquid water (one component, one phase) will have two thermodynamic degrees of freedom: we may change pressure or temperature independently of one another. A vessel containing nothing but boiling water (two phases – water and steam, but still only one component) has just one thermodynamic degree of freedom: we may change pressure and temperature, but just not independently of one another. A vessel containing water at its triple point (three phases, one component) has no thermodynamic freedom at all: both temperature and pressure are fixed⁵¹ so long as all three phases coexist in equilibrium.

2.10.10 Applications of phase changes


Applications of phase changes abound in industrial and commercial processes. Some of these applications exploit phase changes for certain production goals, such as the storage and transport of energy. Others merely serve to illustrate certain phenomena such as latent heat and degrees of thermodynamic freedom. This subsection will highlight several different processes for your learning benefit.

⁵⁰The triple point for any substance is the pressure at which the boiling and freezing temperatures become one and the same.

⁵¹The non-freedom of both pressure and temperature for a pure substance at its triple point means we may exploit different substances' triple points as *calibration standards* for both pressure and temperature. Using suitable laboratory equipment and samples of sufficient purity, anyone in the world may force a substance to its triple point and calibrate pressure and/or temperature instruments against that sample.

Propane storage tanks

A common example of a saturated liquid/vapor (two-phase) system is the internal environment of a propane storage tank, such as the kind commonly used to store propane fuel for portable stoves and gas cooking grills. If multiple propane storage tanks holding different volumes of liquid propane are set side by side, pressure gauges attached to each tank will all register the exact same pressure:



This is counter-intuitive, as most people tend to think the fullest tank should register the highest pressure (having the least space for the vapor to occupy). However, since the interior of each tank is a liquid/vapor system in equilibrium, the pressure is defined by the point on the liquid/vapor transition curve on the phase diagram for pure propane matching the tanks' *temperature*. Thus, the pressure gauge on each tank actually functions as a *thermometer*⁵², since pressure is a direct function of temperature for a saturated liquid/vapor system and therefore cannot change without a corresponding change in temperature. This is a thermodynamic system with just *one* degree of freedom.

Storage tanks containing liquid/vapor mixtures in equilibrium present unique safety hazards. If ever a rupture were to occur in such a vessel, the resulting decrease in pressure causes the liquid to spontaneously boil, halting any further decrease in pressure. Thus, a punctured propane tank does not lose pressure in the same manner than a punctured compressed air tank loses pressure. This gives the escaping vapor more “power” to worsen the rupture, as its pressure does not fall off over time the way it would in a simple compressed-gas application. As a result, relatively small punctures can and often do grow into catastrophic ruptures, where all liquid previously held inside the tank escapes and flashes into vapor, generating a vapor cloud of surprisingly large volume⁵³.

Compounding the problem of catastrophic tank rupture is the fact that propane happens to be highly flammable. The thermodynamic properties of a boiling liquid combined with the chemical


⁵²To be more precise, a propane tank acts like a *Class II filled-bulb* thermometer, with liquid and vapor coexisting in equilibrium.

⁵³Steam boilers exhibit this same explosive tendency. The expansion ratio of water to steam is on the order of a thousand to one (1000:1), making steam boiler ruptures very violent even at relatively low operating pressures.

property of flammability in air makes propane tank explosions particularly violent. Fire fighters often refer to this as a *BLEVE*: a *Boiling Liquid Expanding Vapor Explosion*.

Class II Filled-bulb thermometers

This same pressure-temperature interdependence finds application in a type of temperature measurement instrument called a *Class II filled-bulb*, where a metal bulb, tube, and pressure-sensing element are all filled with a saturated liquid/vapor mixture:



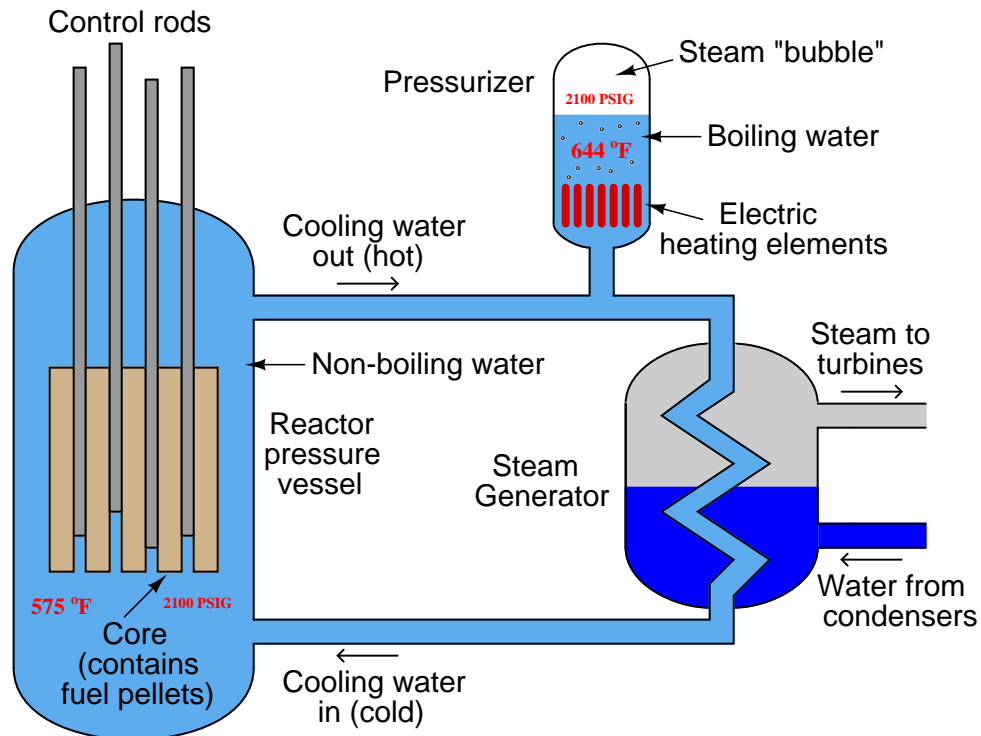
Heat applied to the bulb literally “boils” the liquid inside until its pressure reaches the equilibrium point with temperature. As the bulb’s temperature increases, so does the pressure throughout the sealed system, indicating at the operator display where a bellows (or some other pressure-sensing element) moves a pointer across a calibrated scale.

The only difference between the two filled-bulb thermometers shown in the illustration is which end of the instrument is warmer. The Class IIA system on the left (where liquid fills the pressure-indicating element) is warmer at the bulb than at the indicating end. The Class IIB system on the right (where vapor fills the indicating bellows) has a cooler bulb than the indicating bellows. The long length and small internal diameter of the connecting tube prevents any substantial heat transfer from one end of the system to the other, allowing the sensing bulb to easily be at a different temperature than the indicating bellows. Both types of Class II thermometers work the same⁵⁴, the indicated pressure being a strict function of the bulb’s temperature where the liquid and vapor coexist in equilibrium.

⁵⁴Class IIA systems do suffer from *elevation error* where the indicator may read a higher or lower temperature than it should due to hydrostatic pressure exerted by the column of liquid inside the tube connecting the indicator to the sensing bulb. Class IIB systems do not suffer from this problem, as the gas inside the tube exerts no pressure over an elevation.

Nuclear reactor pressurizers

Nuclear reactors using pressurized water as the moderating and heat-transfer medium must maintain the water coolant in liquid form despite the immense heat output of the reactor core, to avoid the formation of steam bubbles within the reactor core which could lead to destructive “hot spots” inside the reactor. The following diagram shows a simplified⁵⁵ pressurized water reactor (PWR) cooling system:



In order to maintain a liquid-only cooling environment for the reactor core, the water is held at a pressure too high for boiling to occur inside the reactor vessel. Typical operating conditions for a pressurized water reactor are 575 °F and 2100 PSIG. A steam table shows the boiling point of water at 2100 PSIG to be over 640 °F, which means the water inside the reactor cannot boil if the reactor only operates at 575 °F. Referencing the phase diagram for water, the operating point of the reactor core is maintained *above* the liquid/vapor phase transition line by an externally supplied pressure.

This excess pressure comes from a device in the primary coolant loop called a *pressurizer*. Inside the pressurizer is an array of immersion-style electric heater elements. The pressurizer is essentially an electric boiler, purposely boiling the water inside at a temperature greater⁵⁶ than that reached

⁵⁵Circulation pumps and a multitude of accessory devices are omitted from this diagram for the sake of simplicity.

⁵⁶This is another example of an important thermodynamic concept: the distinction between *heat* and *temperature*. While the temperature of the pressurizer heating elements exceeds that of the reactor core, the total heat output of course does not. Typical comparative values for pressurizer power versus reactor core power are 1800 kW versus

by the reactor core itself. For the example figure of 2100 PSIG, the pressurizer elements would have to operate at a temperature of approximately 644 °F to maintain a boiling condition inside the pressurizer.

By maintaining the water temperature inside the pressurizer greater than at the reactor core, the water flowing through the reactor core literally *cannot* boil. The water/vapor equilibrium inside the pressurizer is a system with one degree of freedom (pressure and temperature linked), while the water-only environment inside the reactor core has two degrees of freedom (temperature may vary to any amount below the pressurizer's temperature without water pressure changing at all). Thus, the pressurizer functions like the temperature-sensing bulb of a *gigantic* Class IIA filled-bulb thermometer, with a liquid/vapor equilibrium inside the pressurizer vessel and liquid only inside the reactor vessel and all other portions of the primary coolant loop. Reactor pressure is then controlled by the temperature inside the pressurizer, which is in turn controlled by the amount of power applied to the heating element array⁵⁷.


Steam boilers

Boilers in general (the nuclear reactor system previously described being just one example of a large “power” boiler) are outstanding examples of phase change applied to practical use. The purpose of a boiler is to convert water into steam, sometimes for heating purposes, sometimes as a means of producing mechanical power (through a steam engine), sometimes for chemical processes requiring pressurized steam as a reactant, sometimes for utility purposes (maintenance-related cleaning, process vessel purging, sanitary disinfection, fire suppression, etc.) or all of the above. Steam is a tremendously useful substance in many industries, so you will find boilers in use at almost every industrial facility.

3800 MW, respectively: a ratio exceeding three orders of magnitude. The pressurizer heating elements don't have to dissipate much power (compared to the reactor core) because the pressurizer is not being cooled by a forced convection of water like the reactor core is.

⁵⁷In this application, the heaters are the *final control element* for the reactor pressure control system.

A simplified diagram of a *water tube* boiler appears here:




Water enters the boiler through a heat exchanger in the stack called an *economizer*. This allows cold water to be pre-heated by the warm exhaust gases before they exit the stack. After pre-heating in the economizer, the water enters the boiler itself, where water circulates by natural convection (“thermosiphon”) through a set of tubes exposed to high-temperature fire. Steam collects in the “steam drum,” where it is drawn off through a pipe at the top. Since this steam is in direct contact with the boiling water, it will be at the same temperature as the water, and the steam/water environment inside the steam drum represents a two-phase system with only one degree of freedom. With just a single degree of freedom, steam temperature and pressure are direct functions of each other – coordinates at a single point along the liquid/vapor phase transition line of water’s phase diagram. One cannot change one variable without changing the other.

Consulting a steam table⁵⁸, you will find that the temperature required to boil water at a pressure of 120 PSIG is approximately 350 degrees Fahrenheit. Thus, the temperature of the steam drum will be fixed at 350 °F while generating steam pressure at 120 PSIG. The only way to increase pressure in that boiler is to increase its temperature, and vice-versa.

When steam is at the same temperature as the boiling water it came from, it is referred to as *saturated* steam. Steam in this form is very useful for heating and cleaning, but not as much for operating mechanical engines or for process chemistry. If saturated steam loses any temperature at all (by losing its latent heat), it immediately condenses back into water. Liquid water can cause major mechanical problems inside steam engines (although “wet” steam works wonderfully well as a cleaning agent!), and so steam must be made completely “dry” for some process applications.

⁵⁸Since the relationship between saturated steam pressure and temperature does not follow a simple mathematical formula, it is more practical to consult published tables of pressure/temperature data for steam. A great many engineering manuals contain steam tables, and in fact entire books exist devoted to nothing but steam tables.

The way this is done is by a process known as *superheating*. If steam exiting the steam drum of a boiler is fed through another heat exchanger inside the firebox so it may receive more heat, its temperature will rise beyond the saturation point. This steam is now said to be *superheated*:



Superheated steam is absolutely dry, containing no liquid water at all. It is therefore safe to use as a fluid medium for engines (piston and turbine alike) and as a process reactant where liquid water is not tolerable. The difference in temperature between superheated steam and saturated steam at any given pressure is the amount of *superheat*. For example, if saturated steam at 350 degrees Fahrenheit and 120 PSI drawn from the top of the steam drum in a boiler is heated to a higher temperature of 380 degrees Fahrenheit (at the same pressure of 120 PSI), it is said to have 30 degrees (Fahrenheit) of superheat.


Fruit crop freeze protection

An interesting application of phase changes and latent heat occurs in agriculture. Fruit growers, needing to protect their budding crop from the damaging effects of a late frost, will spray water over the fruit trees to maintain the sensitive buds above freezing temperature. As cold air freezes the water, the water's latent heat of fusion prevents the temperature at the ice/water interface from dropping below 32 degrees Fahrenheit. So long as liquid water continues to spray over the trees, the buds' temperature *cannot* fall below freezing. Indeed, the buds cannot even freeze in this condition, because once they cool down to the freezing point, there will be no more temperature difference between the freezing water and the buds. With no difference of temperature, no heat will transfer out of the buds. With no heat loss, water inside the buds cannot change phase from liquid to solid (ice) even if held at the freezing point for long periods of time, thus preventing freeze damage⁵⁹. Only if the buds are exposed to cold air (below the freezing point), or the water turns completely to ice and no longer holds stable at the freezing point, can the buds themselves ever freeze solid.

⁵⁹An experiment illustrative of this point is to maintain an ice-water mixture in an open container, then to insert a sealed balloon containing liquid water into this mixture. The water inside the balloon will eventually equalize in temperature with the surrounding ice-water mix, but it will not itself freeze. Once the balloon's water reaches 0 degrees Celsius, it stops losing heat to the surrounding ice-water mix, and therefore cannot make the phase change to solid form.

Evaporative cooling towers

A very common use of a liquid-to-vapor phase change is for cooling purposes: taking hot water and mechanically forcing that hot water to evaporate in order to remove large quantities of heat energy from the water, thus cooling it to a lower temperature. Devices called *evaporative cooling towers* accomplish this task by causing ambient air to flow past droplets of water. As the rising air contacts the falling water droplets, some of the water is forced to evaporate, the heat required of this evaporation being provided by sensible heat within the remaining liquid water. As a result, the still-liquid water must cool in temperature as it gives up heat energy to the newly-formed water vapor.



Smaller evaporative cooling towers use fans to force air upward through the tower, employing inert “fill” material to provide large amounts of surface area for the liquid water and the air to contact each other. Some large evaporative cooling towers are self-drafting, the heat of the water providing enough convective force to the air that no fans are needed.

The following photograph shows a pair of induced-draft evaporative cooling towers used at a beer brewery:



This next photograph shows a forced-draft evaporative cooling tower used at a coal-fired electric power plant. Note the large air fans located around the circumference of the cooling tower, blowing cool air into the tower from outside. This fan placement eliminates problems arising from having the fan blades and motor located within the moist air stream:



2.11 Fluid mechanics


A *fluid* is any substance having the ability to *flow*: to freely change shape and move under the influence of a motivating force. Fluid motion may be analyzed on a microscopic level, treating each fluid molecule as an individual projectile body. This approach is extraordinarily tedious on a practical level, but still useful as a model of fluid behavior.

Some fluid properties are accurately predicted by this model, especially predictions dealing with potential and kinetic energies. However, the ability of a fluid's molecules to independently move give it unique properties that solids do not possess. One of these properties is the ability to effortlessly transfer *pressure*, defined as force applied over area.

2.11.1 Pressure


The common phases of matter are *solid*, *liquid*, and *gas*. Liquids and gases are fundamentally distinct from solids in their intrinsic inability to maintain a fixed shape. In other words, liquids and gases tend to fill whatever solid containers they are held in. Similarly, both liquids and gases both have the ability to flow, which is why they are collectively called *fluids*.

Due to their lack of definite shape, fluids tend to disperse any force applied to them. This stands in marked contrast to solids, which tend to transfer force only in the applied direction. Take for example the force transferred by a nail, from a hammer to a piece of wood:



The impact of the hammer's blow is directed straight through the solid nail into the wood below – nothing surprising here.

Now consider what a *fluid* would do when subjected to the same hammer blow:



Given the freedom of a fluid's molecules to move about, the impact of the hammer blow becomes directed *everywhere* against the inside surface of the container (the cylinder). This is true for all fluids: liquids and gases alike. The only difference between the behavior of a liquid versus a gas in the same scenario is that the gas will compress (i.e. the piston will move down as the hammer struck it), whereas the liquid will not compress (i.e. the piston will remain in its resting position). Gases yield under pressure, liquids do not.

It is very useful to quantify force applied to a fluid in terms of force per unit area, since the force applied to a fluid becomes evenly dispersed in all directions to the surface containing it. This is the definition of *pressure*⁶⁰ (P): the amount of force (F) distributed across a given area (A).



$$P = \frac{F}{A}$$

In the metric system, the standard unit of pressure is the *pascal* (Pa), defined as one Newton (N) of force per square meter (m^2) of area. In the British system of measurement, the standard unit of pressure is the *PSI*: pounds (lb) of force per square inch (in^2) of area. Pressure is often expressed in units of kilopascals (kPa) when metric units are used because one pascal is a rather small⁶¹ pressure for most engineering applications.

⁶⁰The concept of pressure is also applicable to solid materials: applying either a compressive or tensile force to a solid object of given cross-sectional area generates a *pressure* within that object, also referred to as *stress*.

⁶¹To give some perspective on this, 1 pascal of pressure is equal to (only) 0.000145 pounds per square inch!

The even distribution of force throughout a fluid has some very practical applications. One application of this principle is the *hydraulic lift*, which functions somewhat like a fluid lever:



Force applied to the small piston creates a pressure throughout the fluid. That pressure exerts a greater force on the large piston than what is exerted on the small piston, by a factor equal to the ratio of piston areas. Since area for a circular piston is proportional to the square of the radius ($A = \pi r^2$), even modest ratios of piston diameter yield large ratios of area and therefore of force. If the large piston has five times the area of the small piston (i.e. the large piston's diameter is 2.236 times greater than the small piston's diameter), force will be multiplied five-fold. Just as with the lever, however, there must be a trade-off so as to not violate the Conservation of Energy. The trade-off for increased force is decreased distance, whether in the lever system or in the hydraulic lift system. If the large piston generates a force five times greater than what is applied to the small piston, it must move only one-fifth as far as the small piston's motion. In this way, energy in equals energy out (remember that *work*, which is equivalent to energy, is calculated by multiplying force by parallel distance traveled).


For those familiar with electricity, what you see here in either the lever system or the hydraulic lift is analogous to a *transformer*: we can step AC voltage up, but only by reducing AC current. Being a passive device, a transformer cannot boost power. Therefore, power out can never be greater than power in, and given a perfectly efficient transformer, power out will always be precisely equal to power in:

$$\text{Power} = (\text{Voltage in})(\text{Current in}) = (\text{Voltage out})(\text{Current out})$$


$$\text{Work} = (\text{Force in})(\text{Distance in}) = (\text{Force out})(\text{Distance out})$$

Fluid may be used to transfer power just as electricity is used to transfer power. Such systems are called *hydraulic* if the fluid is a liquid (usually oil), and *pneumatic* if the fluid is a gas (usually air). In either case, a machine (pump or compressor) is used to generate a continuous fluid pressure, pipes are used to transfer the pressurized fluid to the point of use, and then the fluid is allowed to exert a force against a piston or a set of pistons to do mechanical work:

Hydraulic power system




Pneumatic power system



To learn more about fluid power systems, refer to section [10.2](#) beginning on page [692](#).


An interesting use of fluid we see in the field of instrumentation is as a *signaling medium*, to transfer information between places rather than to transfer power between places. This is analogous to using electricity to transmit voice signals in telephone systems, or digital data between computers along copper wire. Here, fluid pressure represents some other quantity, and the principle of force being distributed equally throughout the fluid is exploited to transmit that representation to some distant location, through piping or tubing:



This illustration shows a simple temperature-measuring system called a *filled bulb*, where an enclosed bulb filled with fluid is exposed to a temperature that we wish to measure. A rise in temperature makes the fluid expand and thereby increases pressure sensed at the gauge. The purpose of the fluid here is two-fold: first to sense temperature, and second to relay this temperature measurement a long distance away to the gauge. The principle of even pressure distribution allows the fluid to act as a signal medium to convey the information (bulb temperature) to a distant location.

2.11.2 Pascal's Principle and hydrostatic pressure

We learned earlier that fluids tend to evenly distribute any applied force. This fundamental principle is the basis of fluid power and fluid signaling systems, where pressure is assumed to be transferred equally to all points in a confined fluid. In the example of a hydraulic lift given earlier, we assume that the pressure throughout the fluid pathway is equal:




If additional force is applied to the small piston (say, 160 lbs instead of 150 lbs), the fluid pressure throughout the system will increase, not just the fluid pressure in the vicinity of the piston. The effect of this additional force will be immediately⁶² “felt” at all points of the system. The phenomenon of pressure changes being evenly distributed throughout an enclosed fluid is called *Pascal's principle*.

Pascal's principle is really nothing more than the direct consequence of fluids' ability to *flow*. The only way an additional applied pressure would *not* be transmitted to all points within a confined fluid volume is if the fluid molecules were somehow not free to move. Since they are mobile, any compression applied to one region of that fluid will propagate to all other regions within that fluid volume. As fluid molecules are subjected to greater pressure, they naturally try to migrate to regions of lower pressure where they “bump up” against other fluid molecules, distributing that increased pressure in doing so.

⁶²There is actually a speed of propagation to this increase in pressure, and it is the *speed of sound* within that particular fluid. This makes sense, since sound waves are nothing more than rapidly-changing regions of pressure within a material.

Pascal's principle tells us any *change* in applied pressure to a confined fluid will be distributed evenly throughout, but it does not say *pressure* will be the same throughout all points. If forces other than those applied to pistons exert pressure on the fluid, we may indeed experience gradients of pressure throughout a confined fluid volume.


In cases where we are dealing with tall columns of dense fluid, there is another force we must consider: the weight of the fluid itself. Suppose we took a cubic foot of water which weighs approximately 62.4 pounds, and poured it into a very tall vertical tube with a cross-sectional area of 1 square inch:



Naturally, we would expect the pressure measured at the bottom of this tall tube to be 62.4 pounds per square inch⁶³, since the entire column of water (weighing 62.4 pounds) has its weight supported by one square inch of area.

⁶³Interestingly, the amount of pressure generated by the weight of a fluid depends only on the *height* of that fluid column, not its cross-sectional area. Suppose we had a column of water the same height (144 feet) but in a tube having an area twice as large: 2 square inches instead of 1 square inch. Twice the area means twice the volume of water held in the tube, and therefore twice the weight (124.8 lbs). However, since this greater weight is distributed over a proportionately greater area at the bottom of the tube, the pressure there remains the same as before: 124.8 pounds \div 2 square inches = 62.4 pounds per square inch.

If we placed another pressure gauge mid-way up the tube, though, how much pressure would it register? At first you might be inclined to say 62.4 PSI as well, because you learned earlier in this lesson that fluids naturally distribute force throughout their bulk. However, in this case the pressure is *not* the same mid-way up the column as it is at the bottom:




The reason for this apparent discrepancy is that the source of pressure in this fluid system comes from the weight of the water column itself. Half-way up the column, the water only experiences half the total weight (31.2 pounds), and so the pressure is half of what it is at the very bottom. We did not consider this effect before, because we assumed the force exerted by the piston in the hydraulic lift was so large it “swamped” the weight of the fluid itself. Here, with our very tall column of water (144 feet tall!), the effect of gravity upon the water’s mass is quite substantial. Indeed, without a piston to exert an external force on the water, weight is the *only* source of force we have to consider when calculating pressure.

This fact does not invalidate Pascal’s principle. Any *change* in pressure applied to the fluid column will still be distributed equally throughout. For example, if we were to place a piston at the top of this fluid column and apply a force to the fluid, pressure at all points in that fluid column would increase by the same amount⁶⁴. This is not the same as saying all pressures will be equal throughout the column, however.

⁶⁴Suppose a 1 square inch piston were set on the top of this tall fluid column, and a downward force of 20 lbs were applied to it. This would apply an *additional* 20 PSI pressure to the fluid molecules at all points within the column. The pressure at the bottom would be 82.4 PSI, and the pressure at the middle would be 51.2 PSI.

An interesting fact about pressure generated by a column of fluid is that the width or shape of the containing vessel is irrelevant: the *height* of the fluid column is the only dimension we need to consider. Examine the following tube shapes, all connected at the bottom:



Since the force of fluid weight is generated only along the axis of gravitational attraction (straight down), that is the only axis of measurement important in determining “hydrostatic” fluid pressure.

The fixed relationship between the vertical height of a water column and pressure is such that sometimes water column height is used as a unit of measurement for pressure. That is, instead of saying “30 PSI,” we could just as correctly quantify that same pressure as 830.4 inches of water (“W.C. or ”H₂O), the conversion factor being approximately 27.68 inches of vertical water column per PSI.

As one might guess, the *density* of the fluid in a vertical column has a significant impact on the hydrostatic pressure that column generates. A liquid twice as dense as water, for example, will produce twice the pressure for a given column height. For example, a column of this liquid (twice as dense as water) 14 inches high will produce a pressure at the bottom equal to 28 inches of water (28 ”W.C.), or just over 1 PSI. An extreme example is liquid mercury, which is over 13.5 times as dense as water. Due to its exceptional density and ready availability, the height of a mercury column is also used as a standard unit of pressure measurement. For instance, 25 PSI could be expressed as 50.9 inches of mercury (“Hg), the conversion factor being approximately 2.036 inches of vertical mercury column per PSI.

The mathematical relationship between vertical liquid height and hydrostatic pressure is quite simple, and may be expressed by either of the following formulae:

$$P = \rho gh$$

$$P = \gamma h$$

Where,

P = Hydrostatic pressure in units of weight per square area unit: pascals (N/m²) or lb/ft²

ρ = Mass density of liquid in kilograms per cubic meter (metric) or slugs per cubic foot (British)

g = Acceleration of gravity (9.81 meters per second squared or 32.2 feet per second squared)

γ = Weight density of liquid in newtons per cubic meter (metric) or pounds per cubic foot (British)

h = Vertical height of liquid column


Dimensional analysis – where we account for all units of measurement in a formula – validates the mathematical relationship between pressure, density, and height. Taking the second formula as an example:

$$P = \gamma h$$

$$\left[\frac{\text{lb}}{\text{ft}^2} \right] = \left[\frac{\text{lb}}{\text{ft}^3} \right] \left[\frac{\text{ft}}{1} \right]$$

As you can see, the unit of “feet” in the height term cancels out one of the “feet” units in the denominator of the density term, leaving an answer for pressure in units of pounds per *square* foot. If one wished to set up the problem so the answer presented in a more common pressure unit such as pounds per square *inch*, both the liquid density and height would have to be expressed in appropriate units (pounds per cubic *inch* and *inches*, respectively).

Applying this to a realistic problem, consider the case of a tank filled with 8 feet (vertical) of castor oil, having a weight density of 60.5 pounds per cubic foot:



This is how we would set up the formula to calculate for hydrostatic pressure at the bottom of the tank:

$$P = \gamma h$$

$$P = \left(\frac{60.5 \text{ lb}}{\text{ft}^3} \right) (8 \text{ ft})$$

$$P = \frac{484 \text{ lb}}{\text{ft}^2}$$

If we wished to convert this result into a more common unit such as PSI (pounds per square inch), we could do so using an appropriate fraction of conversion units:

$$P = \left(\frac{484 \text{ lb}}{\text{ft}^2} \right) \left(\frac{1 \text{ ft}^2}{144 \text{ in}^2} \right)$$

$$P = \frac{3.36 \text{ lb}}{\text{in}^2} = 3.36 \text{ PSI}$$

2.11.3 Fluid density expressions

The *density* of any substance is defined as the ratio of its mass or weight to the volume occupied by that mass or weight. Common expressions of density include pounds per cubic foot (British units) and kilograms per cubic meter (metric units). When the substance in question is a liquid, a common form of expression for density is a ratio of the liquid's density to the density of pure water at standard temperature⁶⁵. This ratio is known as *specific gravity*. For example, the specific gravity of glycerin may be determined by dividing the density of glycerin by the density of water:

$$\begin{aligned}\text{Specific gravity of any liquid} &= \frac{D_{\text{liquid}}}{D_{\text{water}}} \\ \text{Specific gravity of glycerin} &= \frac{D_{\text{glycerin}}}{D_{\text{water}}} = \frac{78.6 \text{ lb/ft}^3}{62.4 \text{ lb/ft}^3} = 1.26\end{aligned}$$

The density of gases may also be expressed in ratio form, except the standard of comparison is ambient air instead of water. Chlorine gas, for example, has a specific gravity of 2.47 (each volumetric unit of chlorine having 2.47 times the mass of the same volume of air under identical temperature and pressure conditions). Specific gravity values for gases are sometimes called *relative gas densities* to avoid confusion with “specific gravity” values for liquids.

As with all ratios, specific gravity is a unitless quantity. In our example with glycerine, we see how the identical units of pounds per cubic foot cancel out of both numerator and denominator, to leave a quotient with no unit at all.

An alternative to expressing fluid density as a ratio of mass (or weight) to volume, or to compare it against the density of a standard fluid such as pure water or air, is to express it as the ratio of volume to mass. This is most commonly applied to vapors such as steam, and it is called *specific volume*. The relationship between specific volume and density is one of mathematical reciprocation: the reciprocal of density (e.g. pounds per cubic foot) is specific volume (e.g. cubic feet per pound). For example, consulting a table of saturated steam properties, we see that saturated steam at a pressure of 60 PSIA has a specific volume of 7.175 cubic feet per pound. Translating this into units of pounds per cubic feet, we reciprocate the value 7.175 to arrive at 0.1394 pounds per cubic foot.

Industry-specific units of measurement also exist for expressing the relative density of a fluid. These units of measurement all begin with the word “degree” much the same as for units of temperature measurement, for example:

- Degrees API (*used in the petroleum industries*)
- Degrees Baumé (*used in a variety of industries including paper manufacture and alcohol production*)
- Degrees Twaddell (*used in the textile industry for tanning solutions and the like*)

⁶⁵Usually, this standard temperature is 4 degrees Celsius, the point of maximum density for water. However, sometimes the specific gravity of a liquid will be expressed in relation to the density of water at some other temperature. In some cases specific gravity is expressed for a liquid at one temperature compared to water at another temperature, usually in the form of a superscript such as 20/4 (liquid at 20 degrees Celsius compared to water at 4 degrees Celsius).

The mathematical relationships between each of these “degree” units of density versus specific gravity⁶⁶ is as follows:

$$\text{Degrees API} = \frac{141.5}{\text{Specific gravity}} - 131.5$$

$$\text{Degrees Twaddell} = 200 \times (\text{Specific gravity} - 1)$$

Two different formulae exist for the calculation of degrees Baumé, depending on whether the liquid in question is heavier or lighter than water. For lighter-than-water liquids:

$$\text{Degrees Baumé (light)} = \frac{140}{\text{Specific gravity}} - 130$$

Note that pure water would measure 10° Baumé on the light scale. As liquid density decreases, the light Baumé value increases. For heavier-than-water liquids:

$$\text{Degrees Baumé (heavy)} = 145 - \frac{145}{\text{Specific gravity}}$$

Note that pure water would measure 0° Baumé on the heavy scale. As liquid density increases, the heavy Baumé value increases.

Just to make things confusing, there are different standards for the heavy Baumé scale. Instead of the constant value 145 shown in the above equation (used throughout the United States of America), an older Dutch standard used the same formula with a constant value of 144. The *Gerlach* heavy Baumé scale uses a constant value of 146.78:

$$\text{Degrees Baumé (heavy, old Dutch)} = 144 - \frac{144}{\text{Specific gravity}}$$


$$\text{Degrees Baumé (heavy, Gerlach scale)} = 146.78 - \frac{146.78}{\text{Specific gravity}}$$

There exists a seemingly endless array of “degree” scales used to express liquid density, scattered throughout the pages of history. For the measurement of sugar concentrations in the food industries, the unit of degrees *Balling* was invented. This scale was later revised to become the unit of degrees *Brix*, which is directly proportional to the percent concentration of sugar in the liquid. Another density scale used for expressing sugar concentration is degrees *Plato*. The density of tanning liquor may be measured in degrees *Bark*. Milk density may be measured in degrees *Soxhlet*. Vegetable oil density (and in older times, the density of oil extracted from sperm whales) may be measured in degrees *Oleo*.

⁶⁶For each of these calculations, specific gravity is defined as the ratio of the liquid’s density at 60 degrees Fahrenheit to the density of pure water, also at 60 degrees Fahrenheit.

2.11.4 Manometers

Expressing fluid pressure in terms of a vertical liquid column makes perfect sense when we use a very simple kind of motion-balance pressure instrument called a *manometer*. A manometer is nothing more than a piece of clear (glass or plastic) tubing filled with a liquid of known density, situated next to a scale for measuring distance. The most basic form of manometer is the *U-tube* manometer, shown here:



The basis for all manometers is the mathematical relationship between a liquid's density (ρ in mass units or γ in weight units) and vertical height. The diameter of the manometer tubes is irrelevant:


$$P = \rho gh$$

$$P = \gamma h$$

Pressure is read on the scale as the difference in height (h) between the two liquid columns. One nice feature of a manometer is it really cannot become "uncalibrated" so long as the fluid is pure and the assembly is maintained in an upright position. If the fluid used is water, the manometer may be filled and emptied at will, and even rolled up for storage if the tubes are made of flexible plastic.

We may create even more sensitive manometers by purposely inclining one or more of the tubes, so that the liquid must travel a farther distance along the tube length to achieve the same vertical shift in height. This has the effect of “amplifying” the liquid’s motion to make it easier to resolve small pressures:

Inclined manometer




This way, a greater motion of liquid (x) is required to generate the same hydrostatic pressure (vertical liquid displacement, h) than in an upright manometer, making the inclined manometer more sensitive. As the similar triangle in the illustration shows, x and h are related trigonometrically by the sine function:

$$\sin \theta = \frac{h}{x}$$

The difference in fluid column positions measured diagonally along the scale (x) must always be greater than the vertical height difference between the two columns (h) by a factor of $\frac{1}{\sin \theta}$, which will always be greater than one for angles less than 90° . The smaller the angle θ , the greater the ratio between x and h , leading to more sensitivity.

If even more sensitivity is desired, we may construct something called a *micromanometer*, consisting of a gas bubble trapped in a clear horizontal tube between two large vertical manometer chambers:

A simple micromanometer



Pressure applied to the top of either vertical chamber will cause the vertical liquid columns to shift just the same as any U-tube manometer. However, the bubble trapped in the clear horizontal tube will move much farther than the vertical displacement of either liquid column, owing to the huge difference in cross-sectional area between the vertical chambers and the horizontal tube. This amplification of motion is analogous to the amplification of motion in a hydraulic piston system (where the smaller piston moves farther than the larger piston), and makes the micromanometer exceptionally sensitive to small pressures.

The movement of the gas bubble within the clear horizontal viewing tube (x) relates to applied pressure by the following formula:


$$x = \frac{\gamma h A_{large}}{2A_{small}}$$

Using water as the working liquid in a standard U-tube manometer, 1 PSI of applied gas pressure results in approximately 27.7 inches of vertical liquid column displacement (i.e. 27.7 inches of height *difference* between the two water columns). This relatively large range of motion limits the usefulness of water manometers to modest pressures only. If we wished to use a water manometer to measure the pressure of compressed air in an industrial pneumatic supply system at approximately 100 PSI, the manometer would have to be in excess of 230 feet tall! Clearly, a water manometer would not be the proper instrument to use for such an application.

However, water is not the only viable liquid for use in manometers. We could take the exact same clear U-tube and fill it partially full of liquid *mercury* instead, which is substantially denser than water. In a mercury manometer, 1 PSI of applied gas pressure results in very slightly more than 2 inches of liquid column displacement. A mercury manometer applied to the task of measuring air pressure in an industrial pneumatic system would only have to be 17 feet tall – still quite large and cumbersome⁶⁷ for a measuring instrument, but not impossible to construct or to use.

⁶⁷A colleague of mine told me once of working in an industrial facility with a very old steam boiler, where boiler

A common form of manometer seen in industrial instrument calibration shops is the *well* type, consisting of a single vertical tube and a relatively large reservoir (called the “well”) acting as the second column:



Due to the well’s much larger cross-sectional area, liquid motion inside of it is negligible compared to the motion of liquid inside the clear viewing tube. For all practical purposes⁶⁸, the liquid level inside the “well” is constant, and so the liquid inside the tube moves the full distance equivalent to the applied pressure. Thus, the well manometer provides an easier means of reading pressure: no longer does one have to measure the difference of height between *two* liquid columns, only the height of a single column.

steam pressure was actually indicated by tall mercury manometers reaching from floor to ceiling. Operations personnel had to climb a ladder to accurately read pressure indicated by these manometers!

⁶⁸To give some perspective on just how little the liquid level changes in the well, consider a well-type manometer with a 1/4 inch (inside) diameter viewing tube and a 4-inch diameter circular well. The ratio of diameters for these two liquid columns is 16:1, which means their ratio of areas is 256:1. Thus, for every inch of liquid motion in the viewing tube, the liquid inside the well moves *only* $\frac{1}{256}$ of an inch. Unless the viewing tube is quite tall, the amount of error incurred by interpreting the tube’s liquid height directly as pressure will be minimal – quite likely less than what the human eye is able to discern on a ruler scale anyway. If the utmost accuracy is desired in a well manometer, however, we may compensate for the trifling motion of liquid in the well by building a custom ruler for the vertical tube – one with a $\frac{255}{256}$ reduced scale (so that $\frac{255}{256}$ of an inch of liquid motion in the tube reads as exactly 1 inch of liquid column) in the case of the 1/4 inch tube and 4 inch well dimensions.

2.11.5 Systems of pressure measurement

Pressure measurement is often a relative thing. When we say there is 35 PSI of air pressure in an inflated car tire, what we mean is that the pressure inside the tire is 35 pounds per square inch *greater than* the surrounding, ambient air pressure. It is a fact that we live and breathe in a pressurized environment. Just as a vertical column of liquid generates a hydrostatic pressure, so does a vertical column of gas. If the column of gas is very tall, the pressure generated by it will be substantial. Such is the case with Earth's atmosphere, the pressure at sea level caused by the weight of the atmosphere being approximately 14.7 PSI.

You and I do not perceive this constant air pressure around us because the pressure inside our bodies is equal to the pressure outside our bodies. Thus our eardrums, which serve as differential pressure-sensing diaphragms, detect no *difference* of pressure between the inside and outside of our bodies. The only time the Earth's air pressure becomes perceptible to us is if we rapidly ascend or descend, where the pressure inside our bodies does not have time to equalize with the pressure outside, and we feel the force of that differential pressure on our eardrums.

If we wish to speak of a fluid pressure in terms of how it compares to a perfect vacuum (absolute zero pressure), we specify it in terms of *absolute* units. For example, when I said earlier that the atmospheric pressure at sea level was 14.7 PSI, what I really meant is it is 14.7 PSIA (pounds per square inch *absolute*), meaning 14.7 pounds per square inch *greater than a perfect vacuum*. When I said earlier that the air pressure inside an inflated car tire was 35 PSI, what I really meant is it was 35 PSIG (pounds per square inch *gauge*), meaning 35 pounds per square inch *greater than ambient air pressure*. The qualifier "gauge" implies the pressure indicated by a pressure-measuring gauge, which in most cases works by comparing the sample fluid's pressure to that of the surrounding atmosphere. When units of pressure measurement are specified without a "G" or "A" suffix, "gauge" pressure is usually⁶⁹ assumed.

⁶⁹With few exceptions!

Gauge and absolute pressure values for some common fluid pressures are shown in this table:

Gauge pressure	Fluid example	Absolute pressure
90 PSIG	Bicycle tire air pressure	104.7 PSIA
35 PSIG	Automobile tire air pressure	49.7 PSIA
0 PSIG	Atmospheric pressure at sea level	14.7 PSIA
-9.8 PSIG (9.8 PSI vacuum)	Engine manifold vacuum under idle conditions	4.9 PSIA
-14.7 PSIG (14.7 PSI vacuum)	Perfect vacuum (no fluid molecules present)	0 PSIA

Note that the only difference between each of the corresponding *gauge* and *absolute* pressures is an offset of 14.7 PSI, with absolute pressure being the larger (more positive) value.

This offset of 14.7 PSI between *absolute* and *gauge* pressures can be confusing if we must convert between different pressure units. Suppose we wished to express the tire pressure of 35 PSIG in units of inches of water column ("W.C.). If we stay in the gauge-pressure scale, all we have to do is multiply by 27.68:

$$\frac{35 \text{ PSI}}{1} \times \frac{27.68 \text{ "W.C.}}{1 \text{ PSI}} = 968.8 \text{ "W.C.}$$

Note how the fractions have been arranged to facilitate cancellation of units. The "PSI" unit in the numerator of the first fraction cancels with the "PSI" unit in the denominator of the second fraction, leaving inches of water column ("W.C.) as the only unit standing. Multiplying the first fraction (35 PSI over 1) by the second fraction (27.68 "W.C. over 1 PSI) is "legal" to do since the second fraction has a *physical* value of unity (1): being that 27.68 inches of water column is the same physical pressure as 1 PSI, the second fraction is really the number "1" in disguise. As we know, multiplying any quantity by unity does not change its value, so the result of 968.8 "W.C. we get has the exact same physical meaning as the original figure of 35 PSI. This technique of unit conversion is sometimes known as *unity fractions*, and it is discussed in more general terms in another section of this book (refer to section 2.4 beginning on page 60).

If, however, we wished to express the car's tire pressure in terms of inches of water column *absolute* (in reference to a perfect vacuum), we would have to include the 14.7 PSI offset in our calculation, and do the conversion in two steps:

$$35 \text{ PSIG} + 14.7 \text{ PSI} = 49.7 \text{ PSIA}$$

$$\frac{49.7 \text{ PSIA}}{1} \times \frac{27.68 \text{ "W.C.}}{1 \text{ PSIA}} = 1375.7 \text{ "W.C.}$$

The ratio between inches of water column and pounds per square inch is still the same (27.68:1) in the absolute scale as it is in the gauge scale. The only difference is that we included the 14.7 PSI offset in the very beginning to express the tire's pressure on the absolute scale rather than on the gauge scale. From then on, all conversions were performed in absolute units.

This two-step conversion process is not unlike converting between different units of temperature (degrees Celsius versus degrees Fahrenheit), and for the exact same reason. To convert from °F to

$^{\circ}\text{C}$, we must first *subtract* an offset of 32 degrees, then *multiply* by $\frac{5}{9}$. The reason an offset is involved in this temperature conversion is because the two temperature scales do not share the same “zero” point: 0°C is *not* the same temperature as 0°F . Likewise, 0 PSIG is *not* the same pressure as 0 PSIA, and so an offset is always necessary to convert between gauge and absolute pressure units.

As seen with the unit of pounds per square inch (PSI), the distinction between gauge and absolute pressure is typically shown by a lettered suffix “G” or “A” following the unit, respectively. Following this convention, we may encounter other units of pressure measurement qualified as either gauge or absolute by these letters: kPaA (kilopascals absolute), inches HgG (inches of mercury gauge), inches W.C.A (inches of water column absolute), etc.

There are some pressure units that are *always* in absolute terms, and as such require no letter “A” to specify. One is the unit of *atmospheres*, 1 atmosphere being 14.7 PSIA. There is no such thing as “atmospheres gauge” pressure. For example, if we were given a pressure as being 4.5 atmospheres and we wanted to convert that into pounds per square inch gauge (PSIG), the conversion would be a two-step process:

$$\frac{4.5 \text{ atm}}{1} \times \frac{14.7 \text{ PSIA}}{1 \text{ atm}} = 66.15 \text{ PSIA}$$

$$66.15 \text{ PSIA} - 14.7 \text{ PSI} = 51.45 \text{ PSIG}$$

Another unit of pressure measurement that is always absolute is the *torr*, equal to 1 millimeter of mercury column absolute (mmHgA). 0 torr is absolute zero, equal to 0 atmospheres, 0 PSIA, or -14.7 PSIG. Atmospheric pressure at sea level is 760 torr, equal to 1 atmosphere, 14.7 PSIA, or 0 PSIG.

If we wished to convert the car tire’s pressure of 35 PSIG into torr, we would once again have to offset the initial value to get everything into absolute terms.

$$35 \text{ PSIG} + 14.7 \text{ PSI} = 49.7 \text{ PSIA}$$

$$\frac{49.7 \text{ PSIA}}{1} \times \frac{760 \text{ torr}}{14.7 \text{ PSIA}} = 2569.5 \text{ torr}$$

One last unit of pressure measurement deserves special comment, for it may be used to express either gauge or absolute pressure, yet it is *not* customary to append a “G” or an “A” to the unit. This unit is the *bar*, exactly equal to 100 kPa, and approximately equal⁷⁰ to 14.5 PSI. Some technical references append a lower-case letter “g” or “a” to the word “bar” to show either gauge pressure (*barg*) or absolute pressure (*bara*), but this notation seems no longer favored. Modern usage typically omits the “g” or “a” suffix in favor of context: the word “gauge” or “absolute” may be included in the expression to clarify the meaning of “bar.” Sadly, many references fail to explicitly declare either “gauge” or “absolute” when using units of *bar*, leaving the reader to interpret the intended context. Despite this ambiguity, the *bar* is frequently used in European literature as a unit of pressure measurement.


⁷⁰The origin of this unit for pressure is the atmospheric pressure at sea level: 1 atmosphere, or 14.7 PSIA. The word “bar” is short for *barometric*, in reference to Earth’s ambient atmospheric pressure.

2.11.6 Negative pressure


If a chamber is completely evacuated of any and all fluid molecules such that it contains nothing but empty space, we say that it contains a perfect *vacuum*. With no fluid molecules inside the chamber whatsoever, there will be no pressure exerted on the chamber walls by any fluid. This is the defining condition of zero absolute pressure (e.g. 0 PSIA, 0 torr, 0 atmospheres, etc.). Referencing atmospheric air pressure⁷¹ outside of this vessel, we could say that the “gauge” pressure of a perfect vacuum is -14.7 PSIG.

A commonly-taught principle is that a perfect vacuum is the lowest pressure possible in any physical system. However, this is not strictly true. It is, in fact, possible to generate pressures *below* 0 PSIA – pressures that are actually *less* than that of a perfect vacuum. The key to understanding this is to consider non-gaseous systems, where the pressure in question exists within a solid or a liquid substance.

Let us begin our exploration of this concept by considering the case of weight applied to a solid metal bar:




Recall that pressure is defined as force exerted over area. This metal bar certainly has a cross-sectional area, and if a compressive force is applied to the bar then the molecules of metal inside the bar will experience a pressure attempting to force them closer together. Supposing the bar in question measured 1.25 inches wide and thick, its cross-sectional area would be $(1.25 \text{ in})^2$, or 1.5625 in^2 . Applying a force of 80 pounds along the bar's length would set up an internal pressure within the bar of 51.2 pounds per square inch, or 51.2 PSI:



⁷¹At sea level, where the absolute pressure is 14.7 PSIA. Atmospheric pressure will be different at different elevations above (or below) sea level.

Now suppose we reverse the direction of the applied force to the bar, applying *tension* to the bar rather than *compression*. If the force is still 80 pounds and the cross-sectional area is still 1.5625 square inches, then the internal pressure inside the bar must be -51.2 PSI:



The negative pressure value describes the tensile force experienced by the molecules of metal inside the bar: a degree of force per unit area attempting to pull those molecules apart from each other rather than push them closer together as was the case with a compressive force.

If you believe that the lowest possible pressure is a perfect vacuum (0 PSIA, or -14.7 PSIG), then this figure of -51.2 PSI seems impossible. However, it is indeed possible because we are dealing with a solid rather than with a gas. Gas molecules exert pressure on a surface by striking that surface and exerting a force by the momentum of their impact. Since gas molecules can only strike (i.e. *push*) against a surface, and cannot *pull* against a surface, one cannot generate a negative absolute pressure using a gas. In solids, however, the molecules comprising the sample exhibit *cohesion*, allowing us to set up a tension within that material impossible in a gaseous sample where there is no cohesion between the molecules. Thus, negative pressures are possible within samples of solid material even though they are impossible within gases.

Negative pressures are also possible within *liquid* samples, provided there are no bubbles of gas or vapor anywhere within the sample. Like solids, the molecules within a liquid also exhibit cohesion (i.e. they tend to “stick” together rather than drift apart from each other). If a piston-and-cylinder arrangement is completely filled with liquid, and a tension applied to the movable piston, the molecules within that liquid will experience tension as well. Thus, it is possible to generate negative pressures (below 0 PSIA) within liquids that are impossible with gases.

Even vertical columns of liquid may generate negative pressure. The famous British scientists Hooke and Boyle demonstrated a negative pressure of -0.2 MPa (-29 PSI) using a column of liquid mercury. Trees naturally generate huge negative pressures in order to draw water to their full height, up from the ground. Two scientists, H.H. Dixon and J. Joly, presented a scientific paper entitled *On the Ascent of Sap* in 1895 proposing liquid tension as the mechanism by which trees could draw water up tremendous heights.


If even the smallest bubble of gas exists within a liquid sample, however, negative pressures become impossible. Since gases can only exert positive pressures, and Pascal’s Principle tells us that pressure will be equally distributed throughout a fluid sample, the low-limit of 0 PSIA for gases establishes a low pressure limit for the entire liquid/gas sample. In other words, the presence of any gas within an otherwise liquid sample prevents the entire sample from experiencing tension.

One limitation to the generation of negative pressures within liquids is that disturbances and/or impurities within the liquid may cause that liquid to spontaneously boil (changing phase from liquid to vapor), at which point a sustained negative pressure becomes impossible.


2.11.7 Buoyancy

When a solid body is immersed in a fluid, it *displaces* an equal volume of that fluid. This displacement of fluid generates an upward force on the object called the *buoyant force*. The magnitude of this force is equal to the weight of the fluid displaced by the solid body, and it is always directed exactly opposite the line of gravitational attraction. This is known as *Archimedes' Principle*.

Buoyant force is what makes ships float. A ship sinks into the water just enough so the weight of the water displaced is equal to the total weight of the ship and all it holds (cargo, crew, food, fuel, etc.):



If we could somehow measure the weight of that water displaced, we would find it exactly equals the dry weight of the ship:



Expressed mathematically, Archimedes' Principle states that the buoyant force is the product of the liquid volume and liquid density:

$$F_{\text{buoyant}} = \gamma V$$

Where,

F_b = Buoyant force exerted on object, opposite in direction from gravity

γ = Weight density of liquid

V = Volume of liquid displaced by the submerged object

We may use dimensional analysis to confirm correct cancellation of British units in the Archimedes' Principle formula:

$$F_{buoyant} = \gamma V$$

$$[\text{lb}] = \frac{[\text{lb}]}{[\text{ft}^3]} [\text{ft}^3]$$

Notice how the units of measurement for weight density (pounds per cubic foot) combine with the unit of measurement for volume (cubic feet) to cancel the unit of cubic feet and leave us with force measured in pounds.

Archimedes' Principle also explains why hot-air balloons and helium aircraft float. By filling a large enclosure with a gas that is less dense than the surrounding air, that enclosure experiences an upward (buoyant) force equal to the difference between the weight of the air displaced and the weight of the gas enclosed. If this buoyant force equals the weight of the craft and all it holds (cargo, crew, food, fuel, etc.), it will exhibit an apparent weight of zero, which means it will float. If the buoyant force exceeds the weight of the craft, the resultant force will cause an upward acceleration according to Newton's Second Law of motion ($F = ma$).


Submarines also make use of Archimedes' Principle, adjusting their buoyancy by adjusting the amount of water held by *ballast tanks* on the hull. Positive buoyancy is achieved by "blowing" water out of the ballast tanks with high-pressure compressed air, so the submarine weighs less (but still occupies the same hull volume and therefore displaces the same amount of water). Negative buoyancy is achieved by "flooding" the ballast tanks so the submarine weighs more. Neutral buoyancy is when the buoyant force exactly equals the weight of the submarine and the remaining water stored in the ballast tanks, so the submarine is able to "hover" in the water with no vertical acceleration or deceleration.

An interesting application of Archimedes' Principle is the quantitative determination of an object's density by submersion in a liquid. For instance, copper is 8.96 times as dense as water, with a mass of 8.96 grams per cubic centimeter (8.96 g/cm^3) as opposed to water at 1.00 gram per cubic centimeter (1.00 g/cm^3). If we had a sample of pure, solid copper exactly 1 cubic centimeter in volume, it would have a mass of 8.96 grams. Completely submerged in pure water, this same sample of solid copper would appear to have a mass of only 7.96 grams, because it would experience a buoyant force equivalent to the mass of water it displaces (1 cubic centimeter = 1 gram of water). Thus, we see that the difference between the dry mass (mass measured in air) and the wet mass (mass measured when completely submerged in water) is the mass of the water displaced. Dividing the sample's dry mass by this mass difference (dry – wet mass) yields the ratio between the sample's mass and the mass of an equivalent volume of water, which is the very definition of specific gravity. The same calculation yields a quantity for specific gravity if *weights* instead of *masses* are used, since weight is nothing more than mass multiplied by the acceleration of gravity ($F_{weight} = mg$), and the constant g cancels out of both numerator and denominator:

$$\text{Specific Gravity} = \frac{m_{dry}}{m_{dry} - m_{wet}} = \frac{m_{dry}g}{m_{dry}g - m_{wet}g} = \frac{\text{Dry weight}}{\text{Dry weight} - \text{Wet weight}}$$


Another application of Archimedes' Principle is the use of a *hydrometer* for measuring liquid density. If a narrow cylinder of precisely known volume and weight (most of the weight concentrated at one end) is immersed in liquid, that cylinder will sink to a level dependent on the liquid's density. In other words, it will sink to a level sufficient to displace its own weight in fluid. Calibrated marks made along the cylinder's length may then serve to register liquid density in any unit desired.

A simple style of hydrometer used to measure the density of lead-acid battery electrolyte is shown in this illustration:



To use this hydrometer, you must squeeze the rubber bulb at the top and dip the open end of the tube into the liquid to be sampled. Relaxing the rubber bulb will draw a sample of liquid up into the tube where it immerses the float. When enough liquid has been drawn into the tube to suspend the float so that it neither rests on the bottom of the tapered glass tube or "tops out" near the bulb, the liquid's density may be read at the air/liquid interface.


A denser electrolyte liquid results in the float rising to a higher level inside the hydrometer tube:



Like all floating objects, the hydrometer float naturally seeks a condition of neutral buoyancy where the weight of the displaced liquid exactly equals the dry weight of the float. If the liquid happens to be very dense, the float will not have to sink very far in order to achieve neutral buoyancy; the less dense the liquid, the deeper the float must sink in order to achieve neutral buoyancy.


This means the float's graduated density scale will read less density toward the top and greater density toward the bottom.

The following photograph shows a set of antique hydrometers used to measure the density of beer. The middle hydrometer bears a label showing its calibration to be in degrees Baumé (heavy):



Liquid density measurement is useful in the alcoholic beverage industry to infer alcohol content. Since alcohol is less dense than water, a sample containing a greater concentration of alcohol (a greater *proof* rating) will be less dense than a “weaker” sample, all other factors being equal.

A less sophisticated version of hydrometer uses multiple balls of differing density. A common application for such a hydrometer is measuring the concentration of “antifreeze” coolant for automobile engines, comprised of a mixture of ethylene glycol and water. Ethylene glycol is a denser compound than water, and so a “stronger” mixture of antifreeze will have a greater bulk density than a “weaker” density of antifreeze. This style of hydrometer yields a crude measurement of ethylene glycol concentration based on the number of balls that float:



A greater number of floating balls represents a “stronger” concentration of glycol in the coolant. “Weak” glycol concentrations represent a greater percentage of water in the coolant, with a correspondingly higher freezing temperature.

Similar hydrometers are used to measure the concentration of sulfuric acid in lead-acid battery electrolyte, comprised of acid and water. The more fully charged a lead-acid battery is, the higher the concentration of sulfuric acid in the electrolyte fluid. The more discharged a lead-acid battery becomes, the less sulfuric acid (and the more water) is present in the electrolyte. Since sulfuric acid is a denser compound than water, measuring electrolyte density with a hydrometer yields a crude measurement of battery charge state.

2.11.8 Gas Laws

The *Ideal Gas Law* relates pressure, volume, molecular quantity, and temperature of an ideal gas together in one concise mathematical expression:

$$PV = nRT$$

Where,

P = Absolute pressure (atmospheres)

V = Volume (liters)

n = Gas quantity (moles)

R = Universal⁷² gas constant (0.0821 L · atm / mol · K)

T = Absolute temperature (K)

For example, the Ideal Gas Law predicts five moles of helium gas (20 grams worth) at a pressure of 1.4 atmospheres and a temperature of 310 Kelvin will occupy 90.9 liters of volume.

An alternative form of the Ideal Gas Law uses the number of actual gas molecules (N) instead of the number of moles of molecules (n):

$$PV = NkT$$

Where,

P = Absolute pressure (Pascals)

V = Volume (cubic meters)

N = Gas quantity (molecules)

k = Boltzmann's constant (1.38×10^{-23} J / K)

T = Absolute temperature (K)

Interestingly, the Ideal Gas Law holds true for *any* gas. The theory behind this assumption is that gases are mostly empty space: there is far more volume of empty space *separating* individual gas molecules in a sample than there is space *occupied* by the gas molecules themselves. This means variations in the sizes of individual gas molecules within any sample is negligible, and therefore the type of gas molecules contained within the sample is irrelevant. Thus, we may apply either form of the Ideal Gas Law to situations regardless of the type of gas involved. This is also why the Ideal Gas Law does *not* apply to liquids or to phase changes (e.g. liquids boiling into gas): only in the gaseous phase will you find individual molecules separated by relatively large distances.

To modify the previous example, where 5 moles of helium gas occupied 90.9 liters at 1.4 atmospheres and 310 Kelvin, it is also true that 5 moles of *nitrogen* gas will occupy the same volume (90.9 liters) at 1.4 atmospheres and 310 Kelvin. The only difference will be the *mass* of each gas sample. 5 moles of helium gas (${}^4\text{He}$) will have a mass of 20 grams, whereas 5 moles of nitrogen gas (${}^{14}\text{N}_2$) will have a mass of 140 grams.

Although no gas in real life is ideal, the Ideal Gas Law is a close approximation for conditions of modest gas density, and no phase changes (gas turning into liquid or vice-versa). You will find this

⁷²It should be noted that many different values exist for R , depending on the units of measurement. For liters of volume, atmospheres of pressure, moles of substance, and Kelvin for temperature, $R = 0.0821$. If one prefers to work with different units of measurement for volume, pressure, molecular quantity, and/or temperature, different values of R are available.

Law appearing again and again in calculations of gas volume and gas flow rates, where engineers and technicians must know the relationship between gas volume, pressure, and temperature.

Since the molecular quantity of an enclosed gas is constant, and the universal gas constant *must* be constant, the Ideal Gas Law may be written as a proportionality instead of an equation:

$$PV \propto T$$

Several “gas laws” are derived from this proportionality. They are as follows:

$$PV = \text{Constant} \quad \text{Boyle's Law (assuming constant temperature } T)$$

$$V \propto T \quad \text{Charles's Law (assuming constant pressure } P)$$

$$P \propto T \quad \text{Gay-Lussac's Law (assuming constant volume } V)$$

You will see these laws referenced in explanations where the specified quantity is constant (or very nearly constant).

For non-ideal conditions, the “Real” Gas Law formula incorporates a corrected term for the *compressibility* of the gas:

$$PV = ZnRT$$

Where,

P = Absolute pressure (atmospheres)

V = Volume (liters)

Z = Gas compressibility factor (unitless)

n = Gas quantity (moles)

R = Universal gas constant (0.0821 L · atm / mol · K)


T = Absolute temperature (K)

The compressibility factor for an ideal gas is unity ($Z = 1$), making the Ideal Gas Law a limiting case of the Real Gas Law. Real gases have compressibility factors less than unity (< 1). What this means is real gases tend to compress more than the Ideal Gas Law would predict (i.e. occupies less volume for a given amount of pressure than predicted, and/or exerts less pressure for a given volume than predicted).

2.11.9 Fluid viscosity

Viscosity is a measure of a fluid’s resistance to shear. It may be visualized as a sort of internal friction, where individual fluid molecules experience either cohesion or collision while flowing past one another. The more “viscous” a fluid is, the “thicker” it is when stirred. Clean water is an example of a low-viscosity liquid, while liquid honey at room temperature is an example of a high-viscosity liquid.

There are two different ways to quantify the viscosity of a fluid: *absolute viscosity* and *kinematic viscosity*. Absolute viscosity (symbolized by the Greek symbol “eta” η , or sometimes by the Greek symbol “mu” μ), also known as *dynamic viscosity*, is a direct relation between stress placed on a fluid and its rate of deformation (or shear). The textbook definition of absolute viscosity is based on a model of two flat plates moving past each other with a film of fluid separating them. The relationship between the shear stress applied to this fluid film (force divided by area) and the velocity/film thickness ratio is viscosity:



$$\eta = \frac{FL}{Av}$$

Where,

η = Absolute viscosity (pascal-seconds), also symbolized as μ

F = Force (newtons)

L = Film thickness (meters) – typically *much* less than 1 meter for any realistic demonstration!

A = Plate area (square meters)

v = Relative velocity (meters per second)

Another common unit of measurement for absolute viscosity is the *poise*, with 1 poise being equal to 0.1 pascal-seconds. Both units are too large for common use, and so absolute viscosity is often expressed in *centipoise*. Water has an absolute viscosity of very nearly 1.000 centipoise.

Kinematic viscosity (symbolized by the Greek letter “nu” ν) includes an assessment of the fluid’s density in addition to all the above factors. It is calculated as the quotient of absolute viscosity and mass density:

$$\nu = \frac{\eta}{\rho}$$

Where,

ν = Kinematic viscosity (stokes)

η = Absolute viscosity (poise)

ρ = Mass density (grams per cubic centimeter)

As with the unit of poise, the unit of stokes is too large for convenient use, so kinematic viscosities are often expressed in units of *centistokes*. Water has a kinematic viscosity of very nearly 1.000 centistokes.

The mechanism of viscosity in liquids is inter-molecular *cohesion*. Since this cohesive force is overcome with increasing temperature, most liquids tend to become “thinner” (less viscous) as they heat up. The mechanism of viscosity in gases, however, is inter-molecular *collisions*. Since these collisions increase in frequency and intensity with increasing temperature, gases tend to become “thicker” (more viscous) as they heat up.

As a ratio of stress to strain (applied force to yielding velocity), viscosity is often constant for a given fluid at a given temperature. Interesting exceptions exist, though. Fluids whose viscosities change with applied stress, and/or over time with all other factors constant, are referred to as *non-Newtonian fluids*. A simple example of a non-Newtonian fluid is cornstarch mixed with water, which “solidifies” under increasing stress and then returns to a liquid state when the stress is removed.

2.11.10 Reynolds number

Viscous flow is a condition where friction forces dominate the behavior of a moving fluid, typically in cases where viscosity (internal fluid friction) is great. *Inviscid flow*, by contrast, is a condition where friction within a moving fluid is negligible and the fluid moves freely. The *Reynolds number* of a fluid is a dimensionless quantity expressing the ratio between a moving fluid's momentum and its viscosity, and is a helpful gauge in predicting how a fluid stream will move.

A couple of formulae for calculating Reynolds number of a flow are shown here:

$$\text{Re} = \frac{D\bar{v}\rho}{\mu}$$

Where,

Re = Reynolds number (unitless)

D = Diameter of pipe, (meters)

\bar{v} = Average velocity of fluid (meters per second)

ρ = Mass density of fluid (kilograms per cubic meter)

μ = Absolute viscosity of fluid (pascal-seconds)

$$\text{Re} = \frac{(3160)G_f Q}{D\mu}$$

Where,

Re = Reynolds number (unitless)

G_f = Specific gravity of liquid (unitless)

Q = Flow rate (gallons per minute)

D = Diameter of pipe (inches)

μ = Absolute viscosity of fluid (centipoise)

3160 = Conversion factor for British units

The first formula, with all metric units, is the textbook “definition” for Reynolds number. If you take the time to dimensionally analyze this formula, you will find that all units do indeed cancel to leave the Reynolds number unitless:

$$\text{Re} = \frac{D\bar{v}\rho}{\mu}$$

$$\text{Re} = \frac{[\text{m}] \left[\frac{\text{m}}{\text{s}}\right] \left[\frac{\text{kg}}{\text{m}^3}\right]}{[\text{Pa} \cdot \text{s}]}$$

Recalling that the definition of a “pascal” is one Newton of force per square meter:

$$\text{Re} = \frac{\left[\frac{\text{kg}}{\text{m} \cdot \text{s}}\right]}{\left[\frac{\text{N} \cdot \text{s}}{\text{m}^2}\right]}$$

$$\text{Re} = \left[\frac{\text{kg}}{\text{m} \cdot \text{s}}\right] \cdot \left[\frac{\text{m}^2}{\text{N} \cdot \text{s}}\right]$$

$$\text{Re} = \left[\frac{\text{kg} \cdot \text{m}}{\text{N} \cdot \text{s}^2}\right]$$

Recalling that the definition of a “newton” is one kilogram times meters per second squared (from Newton’s Second Law equation $F = ma$):

$$\text{Re} = \left[\frac{\text{kg} \cdot \text{m} \cdot \text{s}^2}{\text{kg} \cdot \text{m} \cdot \text{s}^2}\right]$$

$$\text{Re} = \textit{unitless}$$

The second formula given for calculating Reynolds number includes a conversion constant of 3160, which bears the unwieldy unit of “inches-centipoise-minutes per gallon” in order that the units of all variables (flow in gallons per minute, pipe diameter in inches, and viscosity in centipoise) may cancel. Note that specific gravity (G_f) is unitless and therefore does not appear in this dimensional analysis:

$$\text{Re} = \frac{(3160)G_f Q}{D\mu}$$


$$\text{Re} = \frac{\left[\frac{\text{in} \cdot \text{cp} \cdot \text{min}}{\text{gal}}\right] \left[\frac{\text{gal}}{\text{min}}\right]}{[\text{in} \cdot \text{cp}]}$$

$$\text{Re} = \textit{unitless}$$

You will often find this formula, and the conversion constant of 3160, shown without units at all. Its sole purpose is to make the calculation of Reynolds number easy when working with British units customary in the United States.


The Reynolds number of a fluid stream may be used to qualitatively predict whether the flow regime will be *laminar* or *turbulent*. Low Reynolds number values predict laminar (viscous) flow, where fluid molecules move in straight “stream-line” paths, and fluid velocity near the center of the pipe is substantially greater than near the pipe walls:

Laminar flow (low Re)



High Reynolds number values predict turbulent (inviscid) flow, where individual molecule motion is chaotic on a microscopic scale, and fluid velocities across the face of the flow profile are similar:

Turbulent flow (high Re)



It should be emphasized that this turbulence is microscopic in nature, and occurs even when the fluid flows through a piping system free of obstructions, rough surfaces, and/or sudden directional changes. At high Reynolds number values, turbulence simply *happens*.

Other forms of turbulence, such as *eddies* and *swirl* are possible at high Reynolds numbers, but are caused by disturbances in the flow stream such as pipe elbows, tees, control valves, thermowells, and other irregular surfaces. The “micro-turbulence” naturally occurring at high Reynolds numbers will actually randomize such macroscopic (large-scale) motions if the fluid subsequently passes through a long enough length of straight pipe.

Turbulent flow is actually the desired condition for many industrial processes. When different fluids must be mixed together, for example, laminar flow is a bad thing: only turbulent flow will guarantee thorough mixing. The same is true for convective heat exchange: in order for two fluids to effectively exchange heat energy within a heat exchanger, the flow must be turbulent so that molecules from all portions of the flow stream will come into contact with the exchanger walls. Many types of flowmeters require a condition called *fully-developed turbulent flow*, where the flow profile is relatively flat and the only turbulence is that existing on a microscopic scale. Large-scale disturbances in the flow profile such as eddies and swirl tend to negatively affect the measurement performance of many flowmeter designs. This is why such flowmeters usually require long lengths of “straight-run” piping both upstream and downstream: to give micro-turbulence the opportunity to randomize any large-scale motions and homogenize the velocity profile.

A generally accepted rule-of-thumb is that Reynolds number values less than 2000 will probably be laminar, while values in excess of 10000 will probably be turbulent. There is no definite threshold value for all fluids and piping configurations, though. To illustrate, I will share with you some examples of Reynolds number thresholds for laminar versus turbulent flows given by various technical sources:

Chapter 2.8: Laminar Flowmeters of the *Instrument Engineers' Handbook, Process Measurement and Analysis, Third Edition* (pg. 105 – authors: R. Siev, J.B. Arant, B.G. Lipták) define $Re < 2000$ as “laminar” flow, $Re > 10000$ as “fully developed turbulent” flow, and any Reynolds number values between 2000 and 10000 as “transitional” flow.

Chapter 2: Fluid Properties – Part II of the *ISA Industrial Measurement Series – Flow* (pg. 11) define “laminar” flow as $Re < 2000$, “turbulent” flow as $Re > 4000$, and any Reynolds values in between 2000 and 4000 as “transitional” flow.

The Laminar Flow in a Pipe section in the *Standard Handbook of Engineering Calculations* (pg. 1-202) defines “laminar” flow as $Re < 2100$, and “turbulent” flow as $Re > 3000$. In a later section of that *same book* (Piping and Fluid Flow – page 3-384), “laminar” flow is defined as $Re < 1200$ and “turbulent” flow as $Re > 2500$.

Douglas Giancoli, in his physics textbook *Physics* (third edition, pg. 11), defines “laminar” flow as $Re < 2000$ and “turbulent” flow as $Re > 2000$.


Finally, a source on the Internet (<http://flow.netfirms.com/reynolds/theory.htm>) attempts to define the threshold separating laminar from turbulent flow to an unprecedented degree of precision: $Re < 2320$ is supposedly the defining point of “laminar” flow, while $Re > 2320$ is supposedly marks the onset of “turbulent” flow.

Clearly, Reynolds number alone is insufficient for consistent prediction of laminar or turbulent flow, otherwise we would find far greater consistency in the reported Reynolds number values for each regime. Pipe roughness, swirl, and other factors influence flow regime, making Reynolds number an approximate indicator only. It should be noted that laminar flow may be sustained at Reynolds numbers significantly in excess of 10000 under very special circumstances. For example, in certain coiled capillary tubes, laminar flow may be sustained all the way up to $Re = 15000$, due to a phenomenon known as the *Dean effect*!

2.11.11 Law of Continuity

Any fluid moving through a pipe obeys the Law of Continuity, which states that the product of average velocity (\bar{v}), pipe cross-sectional area (A), and fluid density (ρ) for a given flow stream must remain constant:

$$\rho_1 A_1 \bar{v}_1 = \rho_2 A_2 \bar{v}_2 = \cdots \rho_n A_n \bar{v}_n$$



Fluid continuity is an expression of a more fundamental law of physics: the *Conservation of Mass*. If we assign appropriate units of measurement to the variables in the continuity equation, we see that the units cancel in such a way that only units of mass per unit time remain:

$$\rho A \bar{v} = \left[\frac{\text{kg}}{\text{m}^3} \right] \left[\frac{\text{m}^2}{1} \right] \left[\frac{\text{m}}{\text{s}} \right] = \left[\frac{\text{kg}}{\text{s}} \right]$$

This means we may define the product $\rho A \bar{v}$ as an expression of *mass flow rate*, or W :

$$W = \rho A \bar{v}$$

In order for the product $\rho A \bar{v}$ to differ between any two points in a pipe, mass would have to mysteriously appear and disappear. So long as the flow is continuous (not pulsing), and the pipe does not leak, it is impossible to have different rates of mass flow at different points along the flow path without violating the Law of Mass Conservation. The continuity principle for fluid through a pipe is analogous to the principle of current being the same everywhere in a series-connected electric circuit, and for equivalently the same reason⁷³.

⁷³The conservation law necessitating equal current at all points in a series electric circuit is the *Law of Charge Conservation*, which states that electric charges cannot be created or destroyed.

We refer to a flowing fluid as *incompressible* if its density does not substantially change with modest changes in pressure⁷⁴. For this limiting case, ρ is constant and the continuity equation simplifies to the following form:

$$A_1 \bar{v}_1 = A_2 \bar{v}_2$$

Examining this equation in light of dimensional analysis, we see that the product $A\bar{v}$ is also an expression of flow rate:

$$A\bar{v} = \left[\frac{\text{m}^2}{1} \right] \left[\frac{\text{m}}{\text{s}} \right] = \left[\frac{\text{m}^3}{\text{s}} \right]$$

Cubic meters per second is an expression of *volumetric flow rate*, often symbolized by the variable Q :

$$Q = A\bar{v}$$

The practical implication of this principle is that fluid velocity is inversely proportional to the cross-sectional area of a pipe. That is, fluid slows down when the pipe's diameter expands, and vice-versa. We readily see this principle manifest in the natural world: rivers run slowest where they are deep and wide, and run fastest where they are shallow and narrow.

More specifically, we may say that the average velocity of a fluid through a pipe varies inversely with the square of the diameter, since cross-sectional area is proportional to the square of the pipe diameter. For example, if fluid flows at a velocity of 2 feet per second through a 12-inch pipe, and that pipe extends to a narrower section only 6 inches (half the diameter of the wide section), the velocity at the narrower section will be *four times* as great (8 feet per second), since the area of that skinnier section is one-quarter the area of the wider section.

⁷⁴Although not grammatically correct, this is a common use of the word in discussions of fluid dynamics. By definition, something that is “incompressible” *cannot* be compressed, but that is not how we are using the term here. We commonly use the term “incompressible” to refer to either a moving liquid (in which case the actual compressibility of the liquid is inconsequential) or a gas/vapor that *does not happen to undergo substantial compression or expansion as it flows through a pipe*. In other words, an “incompressible” flow is a moving fluid whose ρ does not substantially change, whether by actual impossibility or by circumstance.

For example, consider a pipe with an inside diameter of 8 inches ($2/3$ of a foot), passing a liquid flow of 5 cubic feet per minute. The average velocity (v) of this fluid may be calculated as follows:

$$Q = A\bar{v}$$

$$\bar{v} = \frac{Q}{A}$$

Solving for A in units of square feet:

$$A = \pi r^2$$

$$A = \pi \left(\frac{1}{3} \text{ ft} \right)^2 = \frac{\pi}{9} \text{ ft}^2$$

Now, solving for average velocity \bar{v} :

$$\bar{v} = \frac{Q}{A} = \frac{5 \frac{\text{ft}^3}{\text{min}}}{\frac{\pi}{9} \text{ ft}^2}$$

$$\bar{v} = \left(\frac{5 \text{ ft}^3}{\text{min}} \right) \left(\frac{9}{\pi \text{ ft}^2} \right)$$

$$\bar{v} = \frac{45 \text{ ft}}{\pi \text{ min}} = 14.32 \frac{\text{ft}}{\text{min}}$$

Thus, the average fluid velocity inside an 8-inch pipe passing a volumetric flow rate of 5 cubic feet per minute is 14.32 feet per minute.

2.11.12 Viscous flow

The pressure dropped by a slow-moving, viscous fluid through a pipe is described by the *Hagen-Poiseuille equation*. This equation applies only for conditions of low Reynolds number; i.e. when viscous forces are the dominant restraint to fluid motion through the pipe, and turbulence is nonexistent:

$$Q = k \left(\frac{\Delta P D^4}{\mu L} \right)$$

Where,

Q = Flow rate (gallons per minute)

k = Unit conversion factor = 7.86×10^5

ΔP = Pressure drop (inches of water column)

D = Pipe diameter (inches)

μ = Liquid viscosity (centipoise) – this is a temperature-dependent variable!

L = Length of pipe section (inches)

2.11.13 Bernoulli's equation

Bernoulli's equation is an expression of the *Law of Energy Conservation* for an inviscid (frictionless) fluid stream, named after Daniel Bernoulli⁷⁵. It states that the sum total energy at any point in a passive fluid stream (i.e. no pumps or other energy-imparting machines in the flow path, nor any energy-dissipating elements) must be constant. Two versions of the equation are shown here:

$$z_1\rho g + \frac{v_1^2\rho}{2} + P_1 = z_2\rho g + \frac{v_2^2\rho}{2} + P_2$$

$$z_1 + \frac{v_1^2}{2g} + \frac{P_1}{\gamma} = z_2 + \frac{v_2^2}{2g} + \frac{P_2}{\gamma}$$

Where,

z = Height of fluid (from a common reference point, usually ground level)

ρ = Mass density of fluid

γ = Weight density of fluid ($\gamma = \rho g$)

g = Acceleration of gravity

v = Velocity of fluid

P = Pressure of fluid

Each of the three terms in Bernoulli's equation is an expression of a different kind of energy, commonly referred to as *head*:

$z\rho g$ Elevation head

$\frac{v^2\rho}{2}$ Velocity head

P Pressure head

Elevation and Pressure heads are potential forms of energy, while Velocity head is a kinetic form of energy. Note how the elevation and velocity head terms so closely resemble the formulae for potential and kinetic energy of solid objects:

$$E_p = mgh \quad \text{Potential energy formula}$$

$$E_k = \frac{1}{2}mv^2 \quad \text{Kinetic energy formula}$$

The only real differences between the solid-object and fluid formulae for energies is the use of mass *density* (ρ) for fluids instead of mass (m) for solids, and the arbitrary use of the variable z for height instead of h . In essence, the elevation and velocity head terms within Bernoulli's equation come from the assumption of individual fluid molecules behaving as miniscule solid masses.

⁷⁵According to Ven Te Chow in *Open Channel Hydraulics*, who quotes from Hunter Rouse and Simon Ince's work *History of Hydraulics*, Bernoulli's equation was first formulated by the great mathematician Leonhard Euler and made popular by Julius Weisbach, not by Daniel Bernoulli himself.

It is very important to maintain consistent units of measurement when using Bernoulli's equation! Each of the three energy terms (elevation, velocity, and pressure) *must* possess the exact same units if they are to add appropriately⁷⁶. Here is an example of dimensional analysis applied to the first version of Bernoulli's equation (using British units):

$$z\rho g + \frac{v^2\rho}{2} + P$$

$$[\text{ft}] \left[\frac{\text{slug}}{\text{ft}^3} \right] \left[\frac{\text{ft}}{\text{s}^2} \right] + \left[\frac{\text{ft}}{\text{s}} \right]^2 \left[\frac{\text{slug}}{\text{ft}^3} \right] + \left[\frac{\text{lb}}{\text{ft}^2} \right] = \left[\frac{\text{slug}}{\text{ft} \cdot \text{s}^2} \right]$$

As you can see, both the first and second terms of the equation (elevation and velocity heads) bear the same unit of slugs per foot-second squared after all the “feet” are canceled. The third term (pressure head) does not appear as though its units agree with the other two terms, until you realize that the unit definition of a “pound” is a slug of mass multiplied by the acceleration of gravity in feet per second squared, following Newton's Second Law of motion ($F = ma$):

$$[\text{lb}] = [\text{slug}] \left[\frac{\text{ft}}{\text{s}^2} \right]$$

Once we make this substitution into the pressure head term, the units are revealed to be the same as the other two terms, slugs per foot-second squared:

$$\left[\frac{\text{lb}}{\text{ft}^2} \right] = \left[\frac{\text{slug} \left[\frac{\text{ft}}{\text{s}^2} \right]}{\text{ft}^2} \right] = \left[\frac{\text{slug}}{\text{ft} \cdot \text{s}^2} \right]$$

In order for our British units to be consistent here, we must use *feet* for elevation, *slugs* per cubic *foot* for mass density, *feet* per *second* squared for acceleration, *feet* per *second* for velocity, and *pounds* per square *foot* for pressure. If one wished to use the more common pressure unit of PSI (pounds per square inch) with Bernoulli's equation instead of PSF (pounds per square foot), all the other units would have to change accordingly: elevation in *inches*, mass density in slugs per cubic *inch*, acceleration in *inches* per second squared, and velocity in *inches* per second.

Just for fun, we can try dimensional analysis on the second version of Bernoulli's equation, this time using metric units:

$$z + \frac{v^2}{2g} + \frac{P}{\gamma}$$


$$[\text{m}] + \left[\frac{\left[\frac{\text{m}}{\text{s}} \right]^2}{\left[\frac{\text{m}}{\text{s}^2} \right]} \right] + \left[\frac{\left[\frac{\text{N}}{\text{m}^2} \right]}{\left[\frac{\text{N}}{\text{m}^3} \right]} \right] = [\text{m}]$$

Here, we see that all three terms end up being cast in simple units of meters. That is, the fluid's elevation, velocity, and pressure heads are all expressed as simple elevations. In order for our metric

⁷⁶Surely you've heard the expression, “Apples and Oranges don't add up.” Well, pounds per square inch and pounds per square foot don't add up either! A general mathematical rule in physics is that any quantities added to or subtracted from each other *must* bear the exact same units. This rule does not hold for multiplication or division, which is why we see units canceling in those operations. With addition and subtraction, no unit cancellation occurs.

units to be consistent here, we must use *meters* for elevation, *meters per second* for velocity, *meters per second squared* for acceleration, *pascals (newtons per square meter)* for pressure, and *newtons per cubic meter* for weight density.

Applying Bernoulli's equation to real-life applications can be a bit daunting, as there are so many different units of measurement to contend with, and so many calculations which must be precise in order to arrive at a correct final answer. The following example serves to illustrate how Bernoulli's equation may be applied to the solution of pressure at a point in a water piping system, assuming no frictional losses anywhere in the system:



We know without a doubt that Bernoulli's equation will be what we need to evaluate in order to solve for the unknown pressure P_2 , but where do we begin? A good place to start is by writing the equation we know we will need, then identifying all known values and all unknown values:

$$z_1 \rho g + \frac{v_1^2 \rho}{2} + P_1 = z_2 \rho g + \frac{v_2^2 \rho}{2} + P_2$$

Here is a list of known values, given to us already:

Known quantity	Comments
z_1	0 ft (arbitrarily assigned as 0 height)
z_2	3 ft (if z_1 is 0 feet, then z_2 is 3 ft above it)
v_1	11 ft/s
P_1	46 PSI (<i>need to convert into PSF so all units match</i>)
g	32.2. ft/s ²

The conversion for P_1 from units of PSI into units of PSF is quite simple: multiply 46 PSI by 144 to get 6624 PSF.

Here is a list of values unknown to us at this time:

Unknown quantity	Comments
ρ	(needs to be in units of slugs/ft ³)
v_2	(needs to be in units of ft/s just like v_1)
P_2	(the quantity we are ultimately solving for)

Now all we must do is solve for ρ and v_2 , and we will be ready to use Bernoulli's equation to solve for P_2 . The important of identifying all the known and unknown quantities *before* beginning any calculations cannot be overstated. Doing so allows us to *develop a plan* for solving the problem. Without a plan, one has no idea of where or how to proceed, which is a condition many students repeatedly find themselves in when solving physics-type problems.

We know that ρ is an expression of mass density for the fluid, and we were told the fluid in this example is water. Water has a maximum density of 62.4 pounds per cubic foot, but this figure is not usable in our chosen form of Bernoulli's equation because it is *weight* density (γ) and not *mass* density (ρ).

The relationship between weight density γ and mass density ρ is the exact same relationship between weight (F_W) and mass (m) in a gravitational field (g). Newton's Second Law equation relating force to mass and acceleration ($F = ma$) works well to relate weight to mass and gravitational acceleration:

$$F = ma$$

$$F_W = mg$$

Dividing both sides of this equation by volumetric units (V) (e.g. cubic feet) gives us our relationship between γ and ρ :

$$\frac{F_W}{V} = \frac{m}{V}g$$

$$\gamma = \rho g$$

Water has a weight density of 62.4 pounds per cubic foot in Earth gravity (32.2 feet per second squared), so:

$$\rho = \frac{\gamma}{g}$$

$$\rho = \frac{62.4 \text{ lb/ft}^3}{32.2 \text{ ft/s}^2} = 1.94 \text{ slugs/ft}^3$$

Now we may calculate the total value for the left-hand side of Bernoulli's equation, representing the sum total of potential and kinetic heads for the fluid within the 10-inch pipe:

$$z_1 \rho g + \frac{v_1^2 \rho}{2} + P_1 = \text{Total head at 10-inch pipe}$$

Head	Calculation at 10 inch pipe	Value
$z_1 \rho g$	(0 ft) (1.94 slugs/ft ³) (32.2 ft/s ²)	0 lb/ft ²
$v_1^2 \rho / 2$	(11 ft/s) ² (1.94 slugs/ft ³) / 2	117.4 lb/ft ²
P_1	(46 lb/in ²) (144 in ² /1 ft ²)	6624 lb/ft ²
Total	0 lb/ft ² + 117.4 lb/ft ² + 6624 lb/ft ²	6741.4 lb/ft²

Note the absolutely consistent use of units: all units of distance are *feet*. All units of mass as *slugs*. All units of time are *seconds*. Failure to maintain consistency of units will result in (often severely) incorrect results!⁷⁷

⁷⁷It is entirely possible to perform all our calculations using inches and/or minutes as the primary units instead of feet and seconds. The only caveat is that *all* units throughout all terms of Bernoulli's equation must be consistent. This means we would also have to express mass density in units of slugs per cubic *inch*, the acceleration of gravity in *inches* per second squared (or *inches* per *minute* squared), and velocity in units of *inches* per second (or *inches* per *minute*). The only real benefit of doing this is that pressure would remain in the more customary units of pounds per square *inch*. My personal preference is to do all calculations using units of feet and seconds, then convert pressures in units of PSF to units of PSI at the very end.

There is one more unknown quantity to solve for before we may calculate values at the 6-inch pipe, and that unknown quantity is v_2 . We know that the Continuity equation gives us a mathematical relationship between volumetric flow (Q), pipe area (A), and velocity (v):

$$Q = A_1 v_1 = A_2 v_2$$

Looking at this equation, the only variable we know the value of at this point is v_1 , and we need to find v_2 . However, if we could find the values of A_1 and A_2 , and/or Q , we would have the information we need to solve for v_2 , which in turn would give us the information we would need to solve for P_2 in Bernoulli's equation.

One way to approach this problem is to express the areas and velocities as ratios, eliminating Q entirely so all we need to find are A_1 and A_2 :

$$\frac{A_1}{A_2} = \frac{v_2}{v_1}$$

The area of a circular pipe is given by the basic equation $A = \pi r^2$. Since the problem gives us each pipe's diameter (10 inches and 6 inches), we know the radii (5 inches and 3 inches, respectively) which we may then plug into our ratio equation:

$$\frac{\pi(5 \text{ in})^2}{\pi(3 \text{ in})^2} = \frac{v_2}{v_1}$$

$$\frac{25}{9} = \frac{v_2}{v_1}$$

Knowing v_1 has a value of 11 feet per second, the solution for v_2 is now quite simple:

$$v_2 = 11 \text{ ft/s} \left(\frac{25}{9} \right)$$

$$v_2 = (11 \text{ ft/s})(2.778) = 30.56 \text{ ft/s}$$

Finally, we have all the pieces necessary to solve for P_2 in the right-hand side of Bernoulli's equation:

$$z_2 \rho g + \frac{v_2^2 \rho}{2} + P_2 = \text{Total head at 6-inch pipe}$$

Head	Calculation at 6 inch pipe	Value
$z_2 \rho g$	(3 ft) (1.94 slugs/ft ³) (32.2 ft/s ²)	187.4 lb/ft ²
$v_2^2 \rho / 2$	(30.56 ft/s) ² (1.94 slugs/ft ³) / 2	905.6 lb/ft ²
P_2		(unknown)
Total	187.4 lb/ft ² + 905.6 lb/ft ² + P_2	1093 lb/ft² + P_2

Knowing that the total head calculated at the first location was 6741.4 lb/ft², and the Conservation of Energy requires total heads at both locations be equal (assuming no energy lost to fluid friction along the way), P_2 must be equal to:

$$6741.4 \text{ lb/ft}^2 = 1093 \text{ lb/ft}^2 + P_2$$


$$P_2 = 6741.4 \text{ lb/ft}^2 - 1093 \text{ lb/ft}^2 = 5648.3 \text{ lb/ft}^2$$

Converting pounds per square foot into the more customary unit of pounds per square inch (PSI):

$$P_2 = (5648.3 \text{ lb/ft}^2) \left(\frac{1 \text{ ft}^2}{144 \text{ in}^2} \right)$$

$$P_2 = 39.2 \text{ lb/in}^2$$

Before discussing the larger meaning of our solution, it would be good to review the problem-solving plan we followed to calculate P_2 :




First, we identified Bernoulli's equation as being the central equation necessary for solving P_2 . Then, we identified all the known variables within Bernoulli's equation given to us in the problem, and also if there were any unit-conversion operations necessary. Next, we identified any unknown variables necessary to solve for P_2 in Bernoulli's equation. For each of those unknown variables, we found or developed equations to solve for them, based on variables known to us. The graphic shown above illustrates our plan of solution, with arrows showing the dependent relationships where equations supplied values for unknown quantities in other equations.

This is not just a problem-solving technique unique to Bernoulli's equation; it is a general strategy applicable to *any* type of problem where multiple equations must be used to solve for some quantity. The study of physics in general is filled with problems like this!

Note how our calculated value for P_2 at the second gauge is so much lower than the pressure at the first gauge: 39.2 PSI compared to 46 PSI. This represents nearly a 7 PSI decrease in pressure! Note also how little vertical distance separates the two gauges: only 3 feet. Clearly, the change in elevation between those two points is insufficient to account for the large loss in pressure⁷⁸. Given a 3 foot difference in elevation, one would expect a pressure reduction of about 1.3 PSI for a static column of water, but what we're seeing in this piping system is a pressure drop of nearly 7 PSI. The difference is due to an exchange of energy from potential to kinetic form, as the fluid enters a much narrower pipe (6 inches instead of 10) and must increase velocity.

Furthermore, if we were to increase the flow rate discharged from the pump, resulting in even more velocity through the narrow pipe, pressure at P_2 might even drop lower than atmospheric. In other words, Bernoulli's equation tells us we can actually produce a *vacuum* by accelerating a fluid through a constriction. This principle is widely used in industry with devices known as *eductors* or *ejectors*⁷⁹: tapered tubes through which fluid flows at high velocity to produce a vacuum at the throat.



This, in fact, is how a *carburetor* works in an older automobile engine to vaporize liquid gasoline fuel into a stream of air drawn into the engine: the engine's intake air passes through a venturi tube, where vacuum at the throat of the venturi produces enough negative pressure to draw liquid gasoline into the stream to produce a fine mist.

⁷⁸A simple approximation for pressure loss due to elevation gain is approximately 1 PSI for every 2 vertical feet of water (1 PSI for every 27.68 inches to be more exact).

⁷⁹Technically, an *eductor* uses a liquid such as water to generate the vacuum, while an *ejector* uses a gas or a vapor such as steam.


Ejectors use a high-velocity gas or vapor (e.g. superheated steam) to produce significant vacuums. Eductors use process liquid flow, such as the eductor shown in this next photograph where wastewater flow creates a vacuum to draw gaseous chlorine into the stream for biological disinfection:



Here, the eductor helps fulfill an important safety function. By creating a vacuum to draw toxic chlorine gas from the supply tank into the water stream, the chlorine gas piping may be continuously maintained at a slightly negative pressure throughout. If ever a leak were to develop in the chlorine system, this vacuum would cause ambient air to enter the chlorine pipe rather than toxic chlorine gas to exit the pipe, making a leak far less dangerous than if the chlorine gas piping were maintained in a pressurized state.

2.11.14 Torricelli's equation

The velocity of a liquid stream exiting from a nozzle, pressured solely by a vertical column of that same liquid, is equal to the free-fall velocity of a solid mass dropped from the same height as the top of the liquid column. In both cases, potential energy (in the form of vertical height) converts to kinetic energy (motion):



This was discovered by Evangelista Torricelli almost 100 years prior to Bernoulli's more comprehensive formulation. The velocity may be determined by solving for v after setting the potential and kinetic energy formulae equal to each other (since all potential energy at the upper height must translate into kinetic energy at the bottom, assuming no frictional losses):

$$mgh = \frac{1}{2}mv^2$$

$$gh = \frac{1}{2}v^2$$


$$2gh = v^2$$

$$v = \sqrt{2gh}$$

Note how mass (m) simply disappears from the equation, neatly canceling on both sides. This means the nozzle velocity depends only on height, not the mass density of the liquid. It also means the velocity of the falling object depends only on height, not the mass of the object.

2.11.15 Flow through a venturi tube


If an incompressible fluid moves through a *venturi tube* (i.e. a tube purposefully built to be narrow in the middle), the continuity principle tells us the fluid velocity must increase through the narrow portion. This increase in velocity causes kinetic energy to increase at that point. If the tube is level, there will be negligible difference in elevation (z) between different points of the tube's centerline, which means elevation head remains constant. According to the Law of Energy Conservation, some other form of energy must decrease to account for the increase in kinetic energy. This other form is the pressure head, which decreases at the throat of the venturi:



Ideally, the pressure downstream of the narrow throat should be the same as the pressure upstream, assuming equal pipe diameters upstream and down. However, in practice the downstream pressure gauge will show slightly less pressure than the upstream gauge due to some inevitable energy loss as the fluid passed through the venturi. Some of this loss is due to fluid friction against the walls of the tube, and some is due to viscous losses within the fluid driven by turbulent fluid motion at the high-velocity throat passage.

The difference between upstream and downstream pressure is called *permanent pressure loss*, while the difference in pressure between the narrow throat and downstream is called *pressure recovery*.

If we install vertical sight-tubes called *piezometers*⁸⁰ along a horizontal venturi tube, the differences in pressure will be shown by the heights of liquid columns within the tubes. Here, we assume an ideal (inviscid) liquid with no permanent pressure loss:




The height of liquid in each piezometer tube represents the amount of *potential energy*⁸¹ in the fluid at that point along the venturi tube.

⁸⁰A piezometer tube is nothing more than a manometer (minus the well or the other half of the U-tube).

⁸¹For a moving fluid, potential energy is the sum of fluid height and static pressure.

We may gain more insight into the nature of energy in this moving fluid stream if we add three more piezometers, each one equipped with its own *Pitot tube* facing upstream to “catch” the velocity of the fluid. Rather than represent potential energy by liquid height as the straight-tube piezometers do, the Pitot tube piezometers represent the *total energy* (potential plus kinetic) of the fluid. As such, the liquid heights in these new piezometers are all equal to each other, showing that total energy is indeed conserved at every point in the system:


$$z + \frac{v^2}{2g} + \frac{P}{\gamma} = (\text{constant})$$



Here, each of the “heads” represented⁸² in Bernoulli’s equation are shown in relation to the different piezometer heights. The difference in liquid column height between each Pitot tube piezometer (potential + kinetic energy) and its corresponding straight-tube piezometer (potential energy alone) reflects the amount of kinetic energy possessed by the fluid stream at that point in the venturi tube.

⁸²The form of Bernoulli’s equation with each term expressed in units of distance (e.g. $z = [\text{feet}]$; $\frac{v^2}{2g} = [\text{feet}]$; $\frac{P}{\gamma} = [\text{feet}]$) was chosen so that the piezometers’ liquid heights would directly correspond.

In a real venturi tube, there is some energy permanently lost in the moving fluid due to friction. Consequently the piezometer measurements in a real venturi tube would look something like this:



The “energy line” is seen to slope downhill from inlet to outlet on the venturi tube, showing a degradation in total energy content from beginning to end.

References

- Caupin, F. and Herbert, C., *Cavitation in Water: A Review*, C.R. Physique 7, pages 1000-1017, 2006.
- Chow, Ven Te., *Open-Channel Hydraulics*, McGraw-Hill Book Company, Inc., New York, NY, 1959.
- Considine, Douglas C., *Energy Technology Handbook*, McGraw-Hill Book Company, New York, NY, 1977.
- Control Valve Handbook*, Third Edition, Fisher Controls International, Inc., Marshalltown, IA, 1999.
- Coy, John J.; Townsend, Dennis P.; and Zaretsky, Erwin V., *Gearing*, NASA Reference Publication 1152, AVSCOM Technical Report 84-C-15, National Aeronautics and Space Administration, Scientific and Technical Information Branch, Cleveland, OH, 1985.
- Faydor, L. Litvin; Egelja, A.; Tan, J.; Chen, D.Y-D.; and Heath, G., *Handbook on Face Gear Drives With a Spur Involute Pinion*, University of Illinois at Chicago report E-12127, NASA report NASA CR-2000-209909, U.S. Army Research Laboratory report ARL-CR-447, National Aeronautics and Space Administration, Washington D.C., March 2000.
- Faydor, L. Litvin; Fuentes, Alfonso; Vecchiato, Daniele; and Gonzalez-Perez, Ignacio, *New Design and Improvement of Planetary Gear Trains*, University of Illinois at Chicago report E-14576, NASA report NASA CR-2004-213101, U.S. Army Research Laboratory report ARL-CR-0540, National Aeronautics and Space Administration, Washington D.C., July 2004.
- Giancoli, Douglas C., *Physics for Scientists & Engineers*, Third Edition, Prentice Hall, Upper Saddle River, NJ, 2000.
- Hicks, Tyler G., *Standard Handbook of Engineering Calculations*, McGraw-Hill, Inc., New York, NY, 1972.
- Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.
- Miller, Richard W., *Flow Measurement Engineering Handbook*, Second Edition, McGraw-Hill Publishing Company, New York, NY, 1989.
- Pauling, Linus, *General Chemistry*, Dover Publications, Inc., Mineola, NY, 1988.
- Rouse, Hunter, *Characteristics of Laminar and Turbulent Flow* (video), Iowa Institute of Hydraulic Research, University of Iowa.
- Shapiro, Ascher H., *Pressure Fields and Fluid Acceleration* (video), Massachusetts Institute of Technology, Educational Services Incorporated, 1962.
- Thompson, Ambler and Taylor, Barry N., *Guide for the Use of the International System of Units (SI)*, special publication 811 (second printing), National Institute of Standards and Technology,

Gaithersburg, MD, 2008.

Vennard, John K., *Elementary Fluid Mechanics*, 3rd Edition, John Wiley & Sons, Inc., New York, NY, 1954.

Wall, Göran, *Exergetics*, Bucaramanga, January 2009.

Weast, Robert C.; Astel, Melvin J.; and Beyer, William H., *CRC Handbook of Chemistry and Physics*, 64th Edition, CRC Press, Inc., Boca Raton, FL, 1984.

Chapter 3

Chemistry

Everything there is to know about chemistry – excerpted from “Everything you need to know about school” in the September 16, 2008 edition of the Seattle periodical *The Stranger*:

Stuff is made up of different arrangements of atoms; atoms are made up of nucleus surrounded by buzzing electrons. The outer shell always wants to be filled with eight electrons. So, any arrangement that gets you there – sodium with chloride, oxygen with two hydrogens, carbon with four chlorides – will work. This is why the periodic table has eight columns and helium (with eight outer electrons of its own) doesn’t explode. Some arrangements adding up to eight shared electrons are happier than others. Chemical reactions rearrange from less stable to more stable arrangements on their own, giving off energy in the process. To make a less stable arrangement, you have to put in energy as payment. Chemistry is simply accounting: You must not gain or lose atoms at any point. Ignore the nuclear physicists at this point.

As suggested by this quote, chemistry is the study of matter (stuff) on an atomic scale. It is relevant to industrial instrumentation because so many industrial processes rely on specific chemical reactions to achieve desired outcomes, and we must use instruments to monitor and regulate these chemical reactions. Chemistry can be a confounding subject of study, principally because it seems to defy any simple rule. Many of the “rules” learned by chemistry students, such as the rule of eight electrons referenced in the humorous quote, are not general and in fact only apply to certain elements in the Periodic Table. It should be noted that helium is actually an exception to this rule (an atom of helium only has two electrons, not eight – but at least the quote was correct in saying helium doesn’t explode!). It should also be noted that only a small portion of the Periodic Table has eight columns – most of the table in fact has *eighteen* columns.

Perhaps the most accurate portion of the quote is where it tells us atoms are never lost or gained in a chemical reaction: every atom entering a reaction must somewhere exit that reaction. This simple rule goes by the clumsy name of *stoichiometry* and it is inviolable for all practical purposes. Chemistry, therefore, is the shuffling of atoms between different arrangements which we call *molecules*.

Chemistry is the study of matter: in particular how and why atoms join with one another to form molecules, and the processes by which molecules may be formed and re-formed. Any process where atoms either join with one another to form molecules, or break apart to become individual atoms, is called a *chemical reaction*. Applications of chemistry abound, from the formation of rocks and minerals in the Earth to industrial processes to the processes of organic life itself. Chemistry plays a particularly important role in industrial instrumentation in the form of devices called *analyzers* which exist to measure concentrations of certain chemicals. Analytical instrumentation is essential for industrial processes such as wastewater treatment, combustion, and fermentation to proceed safely and efficiently. Analyzers are also essential for quantitatively tracking pollutants emitted by industrial processes.

Like so many other areas of physical science, the patterns and limits we see in chemical reactions are dominated by two fundamental laws of physics: the *Conservation of Mass* and the *Conservation of Energy*. The particles of matter comprising atoms have the ability to store energy in potential form, and their tendency is to “seek” states having the lowest available energy¹. The arrangement of electrons around the nucleus of an atom is largely dictated by the tendency of electrons to “prefer” stable energy states, and so is the formation of molecules (atoms bonded together): electrons seeking energy states least liable to disturbance. The rest, as they say, is mere detail.

We exploit this property of energy storage in the fuels we use. Atoms bound together to form molecules are in a lower energy state than when they exist as separate atoms. Therefore, an investment of energy is required to force molecules apart (into separate atoms), and energy is returned (released) when atoms join together to form molecules. The combustion of a *fuel*, for example, is nothing more than a process of the atoms in relatively unstable (high-energy) fuel molecules joining with oxygen atoms in air to form stable (low-energy) molecules such as water (H_2O) and carbon dioxide (CO_2).

Natural gas, for example, is a relatively stable combination of hydrogen (H) and carbon (C) atoms, mostly in the form of molecules with a 4:1 hydrogen-to-carbon ratio (CH_4). However, when placed in the vicinity of free oxygen (O) atoms, and given enough energy (a spark) to cause the hydrogen and carbon atoms to separate from each other, the hydrogen atoms strongly bond with oxygen atoms to form water molecules (H_2O), while the carbon atoms also strongly bond with oxygen atoms to form carbon dioxide molecules (CO_2). These strong bonds formed between hydrogen, carbon, and oxygen in the water and carbon dioxide molecules are the result of electrons within those atoms seeking lower energy states than they possessed while forming molecules of natural gas (CH_4). In other words, the electrons binding hydrogen and carbon atoms together to form natural gas are at higher energy states than the electrons binding hydrogen and carbon atoms to oxygen atoms to form water and carbon dioxide, respectively. As those electrons attain lower energy states, their difference of energy must go somewhere (since energy cannot be created or destroyed), and so the chemical reaction releases that energy in the forms of heat and light. This is what you see and feel in the presence of a natural gas flame: the heat and light emitted by hydrogen and carbon atoms joining with oxygen atoms.

¹This generally means to seek the *lowest* gross potential energy, but there are important exceptions where chemical reactions actually proceed in the opposite direction (with atoms seeking *higher* energy states and absorbing energy from the surrounding environment to achieve those higher states). A more general and consistent understanding of matter and energy interactions involves a more complex concept called *entropy*, and a related concept known as *Gibbs Free Energy*.

The Law of Mass Conservation plays an important role in chemistry as well. When atoms join to form molecules, their masses add. That is, the mass of a molecule is precisely equal² to the mass of its constituent atoms. Furthermore, the total mass is unaffected when atoms separate and then re-join to form different molecules. In our natural gas combustion example, the mass of the CH₄ molecules plus the mass of the oxygen atoms they combust with precisely equals the sum total mass of the water and carbon dioxide molecules produced by the combustion. Another way of saying this is that all mass entering a chemical reaction must equal the mass exiting that same reaction. Chemical engineers apply this principle when they calculate *mass balance* in a chemical process: accounting for all mass entering and exiting the process based on the safe assumption that no mass will be gained or lost.

Too many other practical applications of chemistry exist to summarize in these pages, but this chapter aims to give you a foundation to understand basic chemistry concepts necessary to comprehend the function of certain instruments (notably *analyzers*) and processes.

²This statement is not perfectly honest. When atoms join to form molecules, the subsequent release of energy is translated into an incredibly small loss of mass for the molecule, as described by Albert Einstein's famous mass-energy equation $E = mc^2$. However, this mass discrepancy is so small (typically less than one part per *billion* of the original mass!), we may safely ignore it for the purposes of understanding chemical reactions in industrial processes. This is what the humorous quote at the start of this chapter meant when it said "ignore the nuclear physicists at this point".

3.1 Terms and Definitions

- *Atom*: the smallest unit of matter that may be isolated by chemical means.
- *Particle*: a part of an atom, separable from the other portions only by levels of energy far in excess of chemical reactions.
- *Proton*: a type of “elementary” particle, found in the nucleus of an atom, possessing a positive electrical charge.
- *Neutron*: a type of “elementary” particle, found in the nucleus of an atom, possessing no electrical charge, and having nearly the same amount of mass as a proton.
- *Electron*: a type of “elementary” particle, found in regions surrounding the nucleus of an atom, possessing a negative electrical charge, and having just a small fraction of the mass of a proton or neutron.
- *Element*: a substance composed of atoms all sharing the same number of protons in their nuclei (e.g. hydrogen, helium, nitrogen, iron, cesium, fluorine).
- *Atomic number*: the number of protons in the nucleus of an atom – this quantity defines the chemical identity of an atom.
- *Atomic mass* or *Atomic weight*: the total number of elementary particles in the nucleus of an atom (protons + neutrons) – this quantity defines the vast majority of an atom’s mass, since the only other elementary particle (electrons) are so light-weight by comparison to protons and neutrons.
- *Ion*: an atom or molecule that is not electrically balanced (i.e. equal numbers of protons and electrons).
 - *Cation*: a positively-charged ion, called a “cation” because it is attracted toward the negative electrode (cathode) immersed in a fluid solution.
 - *Anion*: a negatively-charged ion, called an “anion” because it is attracted toward the positive electrode (anode) immersed in a fluid solution.
- *Isotope*: a variation on the theme of an element – atoms sharing the same number of protons in their nuclei, but having different numbers of neutrons, are called “isotopes” (e.g. uranium-235 versus uranium-238).
- *Molecule*: the smallest unit of matter composed of two or more atoms joined by electron interaction in a fixed ratio (e.g. water: H₂O). The smallest unit of a *compound*.
- *Compound*: a substance composed of identical molecules (e.g. pure water).
- *Isomer*: a variation on the theme of a compound – molecules sharing the same numbers and types of atoms, but having different structural forms, are called “isomers”. For example, the sugars glucose and fructose are isomers, both having the same formula C₆H₁₂O₆ but having different molecular structures. An isomer is to a molecule as an isotope is to an atomic nucleus.

- *Mixture*: a substance composed of different atoms or molecules not electronically bonded to each other.
- *Solution*: an homogeneous mixture at the molecular level (different atoms/molecules thoroughly mixed together). A solution may be a gas, a liquid, or a solid (e.g. air, saltwater, steel).
 - *Solvent*: the majority element or compound in a solution. Chemists usually consider water to be the *universal solvent*.
 - *Solute*: the minority element or compound in a solution (may be more than one).
 - *Precipitate*: (noun) solute that has “fallen out of solution” due to the solution being saturated with that element or compound; (verb) the process of solute separating from the rest of the solution. (e.g. If you mix too much salt with water, some of the salt will *precipitate* out of the water to form a solid pile at the bottom.)
 - *Supernatant*: the solution remaining above the precipitate.
- *Suspension*: an heterogeneous mixture where separation occurs due to gravity (e.g. mud).
- *Colloid* or *Colloidal suspension*: an heterogeneous mixture where separation either does not occur or occurs at a negligible pace under the influence of gravity (e.g. milk).
 - *Aerosol*: A colloid formed of a solid or liquid substance dispersed in a gas medium.
 - *Foam*: A colloid formed of a gas dispersed in either a liquid or a solid medium.
 - *Emulsion*: A colloid formed of a liquid dispersed in either a liquid or a solid medium.
 - *Sol*: A colloid formed of a solid dispersed in either a liquid or a solid medium.

3.2 Atomic theory and chemical symbols

The three “elementary” particles of matter comprising all atoms are *electrons*, *protons*, and *neutrons*. Combinations of these three particle types in various whole-number quantities constitute every type of atom. These fundamental particles are absolutely miniscule in comparison to the macroscopic existence of human beings. Just to illustrate, the mass of a single proton is approximately 1.67×10^{-27} kilograms: written without scientific notation, it would be 0.0000000000000000000000000000000167 kg. An electron is even smaller: weighing in at 9.11×10^{-31} kg (about *1800 times* less mass than a proton!). Being far smaller in size than a wavelength of visible light³, we cannot see these particles even with the most powerful optical microscope.

Protons and neutrons are very tightly bound together in the nucleus (center) of an atom. This bond is so secure that only extraordinary forces are able to pry an atom’s nucleus apart. Suffice it to say, one cannot disturb the stability of an atomic nucleus by rubbing, cutting, grinding, heating, smashing, or any other macroscopic physical process. The force binding protons and neutrons together in the nucleus is known as the *strong nuclear force*.


Electrons “orbit” the nucleus of atoms, and are held in proximity to those nuclei by electrostatic attraction (the so-called *electromagnetic force*), which is many orders of magnitude weaker than the strong nuclear force. Thus, electrons *can* be dislodged from or added to atoms through the action of macroscopic forces such as rubbing, cutting, grinding, heating, etc. It is the changeable configurations of electrons that accounts for different atoms joining together to form *molecules*.

The chemical identity of any atom is a simple and direct function of how many protons that atom has in its nucleus. Each nitrogen atom, for example, has seven (7) protons in its nucleus. This quantity is called the *atomic number* of an atom. In order for an atom to have a net neutral electric charge, there must be as many electrons orbiting the nucleus as there are protons in the nucleus, since protons are positively charged and electrons are negatively charged (equal and opposite electric charges, each). Therefore, a neutral atom of nitrogen will have seven electrons orbiting around the nucleus, electrically balancing the seven protons within the nucleus.


The number of neutrons within the nucleus of an atom does not affect the atom’s chemical identity, but it may affect its nuclear properties (e.g. whether or not it is radioactive; to what degree it captures certain forms of radiation, etc.). For example, most nitrogen atoms have seven neutrons along with seven protons in their nuclei, giving a total nuclear particle count of fourteen – the *atomic mass* of the atom, sometimes called the *atomic weight*. However, it is possible for a nitrogen atom to have eight neutrons (an atomic mass of fifteen) and still be “nitrogen,” with all the same chemical properties.

³In order for a wave of light to be influenced at all by an object, that object must be at least the size of the wave’s length. To use an analogy with water waves, it would be comparing the interaction of a water wave on a beach against a large rock (a disturbance in the wave pattern) versus the non-disturbance of that same wave as it encounters a small buoy.

A tremendously simplified model of a common nitrogen atom is shown here, with 7 protons and 7 neutrons in the nucleus, and 7 electrons in “orbit” around the nucleus:




The atomic number of this atom (the number of protons in the nucleus) is seven, which is what defines it as nitrogen. The *atomic mass* of this atom (the sum of protons and neutrons in the nucleus) is fourteen. The chemical symbol for this atom is shown here:




The atomic number is redundant to the letter “N” for nitrogen, since only the element nitrogen can have an atomic number of seven. The atomic mass is only relevant when we need to distinguish one *isotope* of nitrogen from another (variations of elements having the same number of protons but different numbers of neutrons), and this is seldom because the chemical properties of isotopes are identical – only their masses differ. For these reasons, you will usually find no left-hand subscripts or superscripts placed near chemical symbols of elements in chemical expressions.

By contrast, subscripts and superscripts placed to the right of a chemical symbol have very important meanings in chemistry. A right-hand subscript refers to the number of atoms bound together to form a molecule. A right-hand superscript refers to the electrical charge possessed by an atom (or by a molecule) by virtue of the number of electrons not matching the number of protons:

An N_2 molecule may be represented simplistically as follows, the two nitrogen atoms joined by a mutual sharing of three of its highest-energy (valence) electrons, shown in this illustration as those electrons residing in the largest-diameter “orbits”. Incidentally, this “triple bond” characterizing the nitrogen molecule is very strong (the more electrons participating in the joining of two atoms, the stronger the bond, all other factors being equal):



An N^{3-} ion is an atom of nitrogen having three more electrons than it normally would when electrically balanced:



A chemical *formula* is a written description of a molecule's constituent atoms. Ethanol (ethyl alcohol), for example, is a conglomerate of two carbon atoms, six hydrogen atoms, and one oxygen atom. One way to express this structure is to write the following formula for ethanol, the right-hand subscripts showing the relative quantities of atoms in each ethanol molecule:




This is called a *molecular formula*, because it shows the proportions of atom types comprising each molecule.

A more common way to write the formula for ethanol, though, is this:




Here, an attempt is made to show the physical structure of the ethanol molecule, where one of the hydrogen atoms is located further away from the others. This is called a *structural formula*. If more detail of the bonds between atoms in a molecule is needed, a semi-graphic representation called a *displayed formula* (also known as an *expanded structural formula*) may be used in lieu of a structural formula:

Displayed formula for ethanol
($\text{C}_2\text{H}_5\text{OH}$)



Each letter in a displayed formula represents a single atom within that molecule, and each line segment in a displayed formula represents a bond between two.

In organic chemistry – the study of molecules principally centered around *carbon* atoms – a special type of notation is used to show the structural detail of the molecule with fewer lines and letters than a displayed formula. This notation is called a *line drawing*, where each line segment represents a single electron bond⁴ to a carbon atom, each vertex and line-end represents the location of a carbon atom, and any hydrogen atoms directly bound to a carbon atom are simply omitted for simplicity. Compare and contrast these displayed formulae and line drawings for a few different organic compounds:



An important principle in organic chemistry is that carbon atoms prefer to form exactly *four* bonds with surrounding atoms⁵. This fact is exploited in line-drawing notation where any bonds not explicitly shown at a vertex are assumed to be single-bonds with hydrogen atoms, enough of them to bring the total number of bonds with that carbon atom to four. Since a great many organic compounds are principally comprised of carbon and hydrogen, the line-drawing symbols for these molecules tend to be more lines than letters.

⁴One line represents a single bond, which is one electron shared per bound atom. Two parallel lines represent a *double bond*, where each carbon atom shares two of its valence electrons with the neighboring atom. Three parallel lines represent a *triple bond*, where each atom shares three of its outer electrons with the neighboring atom.

⁵Incidentally, nitrogen atoms preferentially form exactly *three* bonds, and oxygen atoms exactly *two* bonds. The reason for this pattern is the particular patterns of electrons orbiting each of these atoms, and their respective energy levels. For more information on this, see section 3.4 beginning on page 241.

Chemical engineers often perform mass and energy balance calculations for processes where *mixtures* of similar compounds exist. Wastewater treatment is one example, where an array of organic compounds must all be treated through oxidation (chemical reaction with oxygen). In such cases, it is common for chemical engineers to write formulae expressing the *average* ratios of elements, so that they may calculate the quantity of reactant(s) needed to complete the desired chemical reaction with compounds in the mixture. Primary sludge clarified from municipal wastewater, for example, may be represented by the *compositional formula* $C_{22}H_{39}O_{10}N$. This does not suggest the existence of some monstrous molecule consisting of twenty-two carbon atoms, thirty-nine hydrogen atoms, ten oxygen atoms, and a lone nitrogen atom somewhere in a sample of sludge, but rather that the combined *average* carbon, hydrogen, oxygen, and nitrogen quantities in that sludge exist in a variety of molecular forms in these approximate proportions. This aggregate formula expression helps the engineer quantify the gross chemical characteristics of the sludge, and from that determine how much oxygen will be necessary to completely oxidize it.

Sometimes, compositional formulae are written with non-integer subscripts. An example of this would be the compositional formula $C_{4.8}H_{8.4}O_{2.2}$, which also happens to be an average composition for municipal wastewater sludge (ignoring nitrogen content). The same formula could just as well have been written $C_{48}H_{84}O_{22}$, or even $C_{24}H_{42}O_{11}$, because these subscript values all express the exact same proportions.

3.3 Periodic table of the elements

All substances are comprised of various elements in various combinations and proportions. Elements may thus be thought of as the building-blocks of matter. A Periodic Table of the Elements is a table listing the known elements in order of their atomic numbers.

Periodic Table of the Elements

H Hydrogen 1.00794 1s ¹																	He Helium 4.00260 1s ²															
Li Lithium 6.941 2s ¹	Be Beryllium 9.012182 2s ²													B Boron 10.81 2p ¹	C Carbon 12.011 2p ²	N Nitrogen 14.0067 2p ³	O Oxygen 15.9994 2p ⁴	F Fluorine 18.9984 2p ⁵	Ne Neon 20.179 2p ⁶													
Na Sodium 22.989768 3s ¹	Mg Magnesium 24.3050 3s ²													Al Aluminum 26.9815 3p ¹	Si Silicon 28.0855 3p ²	P Phosphorus 30.9738 3p ³	S Sulfur 32.06 3p ⁴	Cl Chlorine 35.453 3p ⁵	Ar Argon 39.948 3p ⁶													
K Potassium 39.0983 4s ¹	Ca Calcium 40.078 4s ²	Sc Scandium 44.955910 3d ¹ 4s ²	Ti Titanium 47.88 3d ² 4s ²	V Vanadium 50.9415 3d ³ 4s ²	Cr Chromium 51.9961 3d ⁵ 4s ¹	Mn Manganese 54.93805 3d ⁵ 4s ²	Fe Iron 55.847 3d ⁶ 4s ²	Co Cobalt 58.93320 3d ⁷ 4s ²	Ni Nickel 58.69 3d ⁸ 4s ²	Cu Copper 63.546 3d ¹⁰ 4s ¹	Zn Zinc 65.39 3d ¹⁰ 4s ²	Ga Gallium 69.723 4p ¹	Ge Germanium 72.61 4p ²	As Arsenic 74.92159 4p ³	Se Selenium 78.96 4p ⁴	Br Bromine 79.904 4p ⁵	Kr Krypton 83.80 4p ⁶															
Rb Rubidium 85.4678 5s ¹	Sr Strontium 87.62 5s ²	Y Yttrium 88.90585 4d ¹ 5s ²	Zr Zirconium 91.224 4d ² 5s ²	Nb Niobium 92.90638 4d ⁴ 5s ¹	Mo Molybdenum 95.94 4d ⁵ 5s ¹	Tc Technetium (98) 4d ⁵ 5s ²	Ru Ruthenium 101.07 4d ⁷ 5s ¹	Rh Rhodium 102.90550 4d ⁸ 5s ¹	Pd Palladium 106.42 4d ¹⁰ 5s ⁰	Ag Silver 107.8682 4d ¹⁰ 5s ¹	Cd Cadmium 112.411 4d ¹⁰ 5s ²	In Indium 114.82 5p ¹	Sn Tin 118.710 5p ²	Sb Antimony 121.75 5p ³	Te Tellurium 127.60 5p ⁴	I Iodine 126.905 5p ⁵	Xe Xenon 131.30 5p ⁶															
Cs Cesium 132.90543 6s ¹	Ba Barium 137.327 6s ²	57 - 71 Lanthanide series		Hf Hafnium 178.49 5d ² 6s ²	Ta Tantalum 180.9479 5d ³ 6s ²	W Tungsten 183.85 5d ⁴ 6s ²	Re Rhenium 186.207 5d ⁵ 6s ²	Os Osmium 190.2 5d ⁶ 6s ²	Ir Iridium 192.22 5d ⁷ 6s ²	Pt Platinum 195.08 5d ⁹ 6s ¹	Au Gold 196.96654 5d ¹⁰ 6s ¹	Hg Mercury 200.59 5d ¹⁰ 6s ²	Tl Thallium 204.3833 6p ¹	Pb Lead 207.2 6p ²	Bi Bismuth 208.98037 6p ³	Po Polonium (209) 6p ⁴	At Astatine (210) 6p ⁵	Rn Radon (222) 6p ⁶														
Fr Francium (223) 7s ¹	Ra Radium (226) 7s ²	89 - 103 Actinide series		Unq Unnilquadium (261) 6d ⁷ 7s ²	Unp Unnilpentium (262) 6d ⁷ 7s ²	Unh Unnilhexium (263) 6d ⁷ 7s ²	Uns Unnilseptium (262) 6d ⁷ 7s ²	108		109																						
																		La Lanthanum 138.9055 5d ¹ 6s ²	Ce Cerium 140.115 4f ¹ 5d ¹ 6s ²	Pr Praseodymium 140.90765 4f ² 6s ²	Nd Neodymium 144.24 4f ⁴ 6s ²	Pm Promethium (145) 4f ⁵ 6s ²	Sm Samarium 150.36 4f ⁶ 6s ²	Eu Europium 151.965 4f ⁷ 6s ²	Gd Gadolinium 157.25 4f ⁷ 5d ¹ 6s ²	Tb Terbium 158.92534 4f ⁹ 6s ²	Dy Dysprosium 162.50 4f ⁹ 6s ²	Ho Holmium 164.93032 4f ¹⁰ 6s ²	Er Erbium 167.26 4f ¹⁰ 6s ²	Tm Thulium 168.93421 4f ¹¹ 6s ²	Yb Ytterbium 173.04 4f ¹⁴ 6s ²	Lu Lutetium 174.967 4f ¹⁴ 5d ¹ 6s ²
																		Ac Actinium (227) 6d ⁷ 7s ²	Th Thorium 232.0381 6d ² 7s ²	Pa Protactinium 231.03588 5f ² 6d ¹ 7s ²	U Uranium 238.0289 5f ³ 6d ¹ 7s ²	Np Neptunium (237) 5f ⁴ 6d ¹ 7s ²	Pu Plutonium (244) 5f ⁶ 6d ¹ 7s ²	Am Americium (243) 5f ⁷ 6d ² 7s ²	Cm Curium (247) 5f ⁷ 6d ¹ 7s ²	Bk Berkelium (247) 5f ⁹ 6d ² 7s ²	Cf Californium (251) 5f ¹⁰ 6d ² 7s ²	Es Einsteinium (252) 5f ¹¹ 6d ² 7s ²	Fm Fermium (257) 5f ¹² 6d ² 7s ²	Md Mendelevium (258) 5f ¹³ 6d ² 7s ²	No Nobelium (259) 6d ² 7s ²	Lr Lawrencium (260) 6d ¹ 7s ²

Multiple attributes appear for each element in the table. Two of these attributes – atomic number and atomic mass – are directly related to the number of particles in the nucleus of each atom. We will examine the table’s entry for the element *potassium* (K) to explore these concepts.

Potassium has an *atomic number* (number of protons in the nucleus of each potassium atom) of 19. This number defines the element. If we were somehow to add or subtract protons from the nucleus of a potassium atom⁶, it would cease being potassium and *transmute* into a different element. Note how *every* element in the table has its own unique atomic number, and how each of these numbers is whole (no fractions or decimals).

The *atomic mass* or *atomic weight* shown for potassium is 39.0983. This quantity is the sum of protons and neutrons found in the nucleus of each potassium atom. Like the atomic number (19), we would logically expect the atomic mass to be a whole number as well, since protons and neutrons

⁶The amount of energy required to rearrange particles in the nucleus for even just a single atom is *tremendous*, lying well outside the energy ranges of chemical reactions. Such energy levels are the exclusive domain of *nuclear* reactions and high-energy radiation (subatomic particles traveling at high velocity). The extremely large energy “investment” required to alter an atom’s nucleus is why atomic identities are so stable. This is precisely why alchemists of antiquity utterly failed to turn lead into gold: no materials, processes, or techniques they had at their disposal were capable of the targeted energy necessary to dislodge three protons from a nucleus of lead ($_{82}\text{Pb}$) to that it would turn into a nucleus of gold ($_{79}\text{Au}$). That, and the fact the alchemists had no clue about atomic structure to begin with, made their endeavor fruitless.

only come in whole quantities. The primary reason we see a non-whole number for potassium's atomic mass is that this table reflects the *average* atomic mass of potassium atoms as found in nature. Some potassium atoms have atomic masses greater than 39, and some have atomic masses less than 39. We know that the number of protons in every potassium atom is fixed (which is what gives potassium its elemental identity), which means the only quantity that may cause the atomic mass to vary is the number of *neutrons* in the nucleus. The most common form of potassium (^{39}K) atom possesses 19 protons and 20 neutrons in its nucleus, giving it an atomic mass of 39 (19 + 20). The next most common form of potassium found on Earth is (^{41}K), possessing 19 protons and 22 neutrons.

We refer to atoms of the same element with differing atomic masses as *isotopes*. From a chemical perspective, isotopes are identical. That is to say, they engage in the exact same chemical reactions in the exact same manner. To use potassium as an example, an atom of ^{39}K will join with a chlorine atom (Cl) to form the compound *potassium chloride* (KCl) just as readily as an atom of ^{41}K will join with a chlorine atom to form the same compound. The three isotopes of hydrogen (^1H , ^2H , and ^3H : hydrogen, deuterium, and tritium, respectively) are all chemically identical: all are highly flammable, combining with oxygen to create water (H_2O). However, deuterium (^2H) has twice the density of normal hydrogen (^1H), while tritium (^3H) has three times the density of normal hydrogen and is highly radioactive! Isotopes only differ in their mass and in their nuclear properties (such as *radioactivity*: the tendency for a nucleus to spontaneously decay, usually resulting in a loss or gain of protons that subsequently alters the identity of the decayed atom).


The Periodic Table is called “periodic” because its configuration reveals a repeating pattern of chemical behaviors approximately following atomic number. Horizontal rows in the table are called *periods*, while vertical columns are called *groups*. Elements in the same group (vertical column) share similar chemical reactivities – that is, they tend to engage in the same types of chemical reactions – despite having different masses and physical properties such as melting point, boiling point, etc. This *periodicity* is a function of how electrons are arranged around the nucleus of each atom, a subject we will explore in more detail later in this chapter. As mentioned previously, chemistry is the study of how atoms bond together to form molecules, and this bonding takes place through the interaction of the electrons surrounding the atoms' nuclei. It makes perfect sense, then, that the configuration of those electrons determine the chemical (bonding) properties of atoms.

Some periodic tables show the *first ionization energy* value for each element – the amount of energy required to force the first electron of an electrically balanced atom to separate from that atom – in addition to other attributes such as atomic number and atomic mass. If we note the ionization energies of the elements, reading each element in turn from left-to-right, starting with period 1 (hydrogen and helium) and progressing to subsequent periods, we see an interesting pattern:

Element	Period	First ionization energy (measured in “electron-volts”)
Hydrogen (H)	1	13.5984
Helium (He)	1	24.5874
Lithium (Li)	2	5.3917
Beryllium (Be)	2	9.3227
Boron (B)	2	8.2980
Carbon (C)	2	11.2603
Nitrogen (N)	2	14.5341
Oxygen (O)	2	13.6181
Fluorine (F)	2	17.4228
Neon (Ne)	2	21.5645
Sodium (Na)	3	5.1391
Magnesium (Mg)	3	7.6462
Aluminum (Al)	3	5.9858
Silicon (Si)	3	8.1517
Phosphorus (P)	3	10.4867
Sulfur (S)	3	10.3600
Chlorine (Cl)	3	12.9676
Argon (Ar)	3	15.7596
Potassium (K)	4	4.3407

First ionization energy represents the relative stability of the last electron balancing the electrical charge of an atom. We see from this table that 24.5874 electron-volts of energy is needed to remove one electron from an electrically-balanced atom of helium (changing He into He^{1+}), while only 13.5984 electron-volts of energy is required to do the same to an atom of hydrogen. This tells us the electron configuration of helium is at a lower energy (and therefore more stable) than that of hydrogen.

The ionization energies increase with increasing atomic number (with an occasional down-step) until the last column of the period is reached, and then there is a comparatively enormous down-step in energy at the first column of a new period. This pattern is clearly evident when the first ionization energies are plotted against atomic number:



This periodicity suggests that as atoms grow in atomic number, the additional electrons do not simply pile on in random fashion or in a plain and simple progression from inner orbits to outer orbits. Rather, they “fill in” a structured energy pattern, with major changes in structure at the start of each new period. More details of this structured pattern will be explored later in this chapter.

The low ionization energy values for all the “Group 1” elements (far left-hand column) suggest they are relatively easy to positively ionize, and indeed we find this to be the case through experimentation. Hydrogen, lithium, sodium, potassium, and the rest all readily become positively-charged ions upon interaction with other atoms, since their low ionization energy values means they may easily lose an electron.

The high ionization energy values for all the “Group 18” elements (far right-hand column) suggest they possess a very stable electron structure, which is also verified by experiment. These are the *noble* elements, possessing very little reactive potential⁷.

Looking at the “Group 17” column, just to the left of the noble elements, we notice that they are all just one electron shy of the stable electron structure enjoyed by the noble atoms when in their electrically-balanced states. This suggests it might be easy to *add* one more electron to atoms of these elements, which (once again!) is a principle validated by experiment. Fluorine, chlorine, bromine, iodine, and even astatine⁸ all readily ionize negatively, readily accepting an extra electron from surrounding atoms. As one might expect from this tendency, these elements readily bond through electrostatic attraction with the “Group 1” elements (hydrogen, lithium, sodium, potassium, etc.), each “Group 17” atom accepting an extra electron from each “Group 1” atom which readily provides it. Ordinary table salt (sodium chloride, or NaCl) is an example of a compound formed by this sort of bond.

⁷It used to be believed that these elements were completely *inert*: incapable of forming molecular bonds with other atoms. However, this is not precisely true, as some compounds are now known to integrate noble elements.


⁸All isotopes of astatine (At) are radioactive with very short half-lives, making this element difficult to isolate and study.

Thus, Group 1 and Group 17 elements are both highly reactive in a chemical sense, but in different ways. Group 1 elements easily form bonds with Group 17 elements, but Group 1 elements do not generally bond (solely) with other Group 1 elements, and likewise Group 17 elements do not generally bond (solely) with other Group 17 elements. It is the structure of the electrons surrounding each atom's nucleus that determines how those atoms will bond with other atoms.

3.4 Electronic structure

Earlier in this chapter you were shown a model of a nitrogen atom with a dense nucleus (comprised of protons and neutrons) surrounded by electrons whirling around like satellites around a planet. While there are some useful features of this model, it is largely in error. A more realistic view of atomic structure begins with the realization that electrons do not exist as discrete particles, but rather as wave packets. In a sense, they orbit the nucleus within certain *areas of probability*, as described by the principles of quantum mechanics. One way to envision this is to think of an electron's placement around the nucleus in the same way you might picture a city shrouded by a layer of fog. The electron does not have a discrete location (even though there *is* a discrete number of them found in every atom), but rather may be found anywhere within a certain region to varying degrees of probability.

Things get even stranger the more electrons there are in an atom. No two electrons may share the same quantum states in the same atom – a principle called the *Pauli Exclusion Principle*. This means the electrons surrounding a nucleus must exist in distinct patterns. Just a few of these patterns are shown here as *orbitals* (regions of high probability where up to two electrons may be found surrounding a nucleus):⁹



Electrons situate themselves around the nucleus of any atom according to one basic rule: the minimization of potential energy. That is, the electrons seek the lowest-energy positions available around the nucleus. Given the electrostatic attraction between negative electrons and the positive nucleus of an atom, there is potential energy stored in the “elevation” between an orbiting electron and the nucleus, just as there is gravitational potential energy in any object orbiting a planet. Electrons lose energy as they situate themselves closer to the nucleus, and it requires an external input of energy to move an electron farther away from its parent nucleus.

In a sense, most of chemistry may be explained by this principle of minimized potential energy. Electrons “want” to “fall” as close as they can to the positively-charged nucleus. However, there is limited “seating” around the nucleus. As described by Pauli’s Exclusion Principle, electrons cannot simply pile on top of each other in their quest for minimum energy, but rather must occupy certain regions of space allowed by their unique quantum states.

An analogy¹⁰ for visualizing this is to picture an atom as if it were an amphitheater, with the stage being the nucleus and the concentric array of seats being places where electrons may reside. All spectators (electrons) desire to be as close to the stage (nucleus) in an amphitheater (atom) as possible, but since everyone cannot occupy the best seat, people are forced to choose seats at

⁹These orbitals just happen to be the 1s, 2p, 3d, and 4f orbitals, as viewed from left to right. In each case, the nucleus lies at the geometric center of each shape. In a real atom, all orbitals share the same center, which means any atom having more than two electrons (that’s all elements except for hydrogen and helium!) will have multiple orbitals around one nucleus. This four-set of orbital visualizations shows what some orbitals would look like if viewed in isolation.

¹⁰Please understand that like all analogies, this one merely illustrates a complex concept in terms that are easier to recognize. Analogies do not *explain* why things work, but merely liken an abstract phenomenon to something more accessible to common experience.

different positions around the stage. As a result, the inner seats fill first, with most empty seats being furthest away from the stage. The concept of energy fits neatly into this analogy as well: just as electrons give up *energy* to “fall into” lower-energy regions around the nucleus, people must give up *money* to purchase seats closest to the action on stage.

The energy levels available for orbiting electrons are divided into categories of *shells* and *subshells*. A “shell” (or, *principal quantum number, n*) describes the main energy level of an electron. In our amphitheater analogy, this is equivalent to a *tier* or *seating level*. A “subshell” (or, *subsidiary quantum number, l*) further divides the energy levels within each electron shell, and assigns different shapes to the electrons’ probability “clouds.” In the amphitheater analogy, a subshell would be a row of seats within a particular tier. To make the analogy accurate, we would have to imagine each row of seats in a tier having a different shape (not all arcs or straight lines), with varying degrees of viewing comfort afforded by each shape. The first row in each tier faces uniformly toward the stage, allowing easy viewing. Successive rows (subshells) in each tier (shell) contain more seats, but are bent in such a way that the stage is not as easy to view, making these rows less desirable to occupy. Electron subshells always have an even-numbered electron capacity, analogous to theater rows containing even-numbered quantities of seats, because atomic electrons tend to gather in pairs called *orbitals*.

Chemists identify electron shells both by number (the value of the quantum number n) and/or by capital letters: the first shell by the letter K, the second by L, the third by M, and the fourth by N. Higher-order shells exist for atoms requiring¹¹ many electrons (high atomic number), and the lettering pattern is alphabetic (fifth shell is O, sixth is P, etc.). Each successive shell has a greater number of subshells available, like successive amphitheater tiers having more rows: the low-level tiers closest to the stage having the fewest rows, and the high-level tiers furthest from the stage having the most rows.

A numbering and lettering system is also used by chemists to identify subshells within each shell (the l quantum number value starting with zero, and lower-case letters beginning with “s”): the first subshell ($l = 0$) in any shell represented by the letter s, the second ($l = 1$) by p, the third ($l = 2$) by d, the fourth ($l = 3$) by f, and all others by successive lower-case letters of the alphabet¹². Each subshell of each shell has an even-numbered capacity for electrons, since the electrons in each subshell are organized in “orbital” regions, each orbital handling a maximum of two¹³ electrons. The number of orbitals per shell is equal to twice the l value plus one. An “s” subshell has one orbital holding up to two electrons. A “p” subshell has three orbitals holding up to six electrons total. A “d” subshell has five orbitals holding up to ten electrons total. An “f” subshell has seven orbitals holding up to 14 electrons total. A “g” subshell has nine orbitals holding up to 18 electrons total.

The number of subshells in any shell is the same as that shell’s n value. Thus, the first (K) shell has only one subshell, “s”. The second (L) shell has two subshells, an “s” and a “p”. The third (M) shell has three subshells available, an “s”, a “p”, and a “d”; and so on.

¹¹Truth be told, higher-order shells exist even in simple atoms like hydrogen, but are simply not occupied by that atom’s electron(s) unless they are “excited” into a higher energy state by an external input of energy.


¹²The letters *s*, *p*, *d*, and *f* refer to the words *sharp*, *principal*, *diffuse*, and *fundamental*, used to describe the appearance of spectral lines in the early days of atomic spectroscopy research. Higher-order subshells are labeled alphabetically after *f*: *g*, *h*, and *i*.

¹³The two electrons of any orbital have opposite *spin* values.

This table shows the first few shells, their subshells, and electron capacities of each:

Shell (n value)	Subshell (l value)	Subshell electron capacity $= 2(2l + 1)$
$n = 1$; K	$l = 0$; s	$2 \times (2 \times 0 + 1) = 2$
$n = 2$; L	$l = 0$; s	$2 \times (2 \times 0 + 1) = 2$
	$l = 1$; p	$2 \times (2 \times 1 + 1) = 6$
$n = 3$; M	$l = 0$; s	$2 \times (2 \times 0 + 1) = 2$
	$l = 1$; p	$2 \times (2 \times 1 + 1) = 6$
	$l = 2$; d	$2 \times (2 \times 2 + 1) = 10$
$n = 4$; N	$l = 0$; s	$2 \times (2 \times 0 + 1) = 2$
	$l = 1$; p	$2 \times (2 \times 1 + 1) = 6$
	$l = 2$; d	$2 \times (2 \times 2 + 1) = 10$
	$l = 3$; f	$2 \times (2 \times 3 + 1) = 14$

Reviewing our amphitheater analogy, atomic shells are like seating tiers (levels), subshells are like rows of seats within each tier, and subshell electron capacity is like the number of seats in each row. This simple illustration shows an atom with three shells (K, L, and M) with the respective subshells (s, p, and d) represented by differently-shaded rings within each shell, having different numbers of places for electrons within each one:



This illustration is vastly over-simplified, failing to show the diverse shapes of each subshell, serving only to show how each successive shell grows in subshells and electron capacities.

The complete electron configuration for an atom may be expressed using *spectroscopic notation*, showing the shell numbers, subshell letters, and number of electrons residing within each subshell as a superscript. For example, the element helium (with an atomic number of 2) would be expressed as $1s^2$, with just two electrons in the “s” subshell of the first shell. The following table shows the electron structures of the first nineteen elements in the periodic table, from the element hydrogen (atomic number¹⁴ = 1) to potassium (atomic number = 19):

Element	Atomic number	Electron configuration
Hydrogen	1	$1s^1$
Helium	2	$1s^2$
Lithium	3	$1s^2 2s^1$
Beryllium	4	$1s^2 2s^2$
Boron	5	$1s^2 2s^2 2p^1$
Carbon	6	$1s^2 2s^2 2p^2$
Nitrogen	7	$1s^2 2s^2 2p^3$
Oxygen	8	$1s^2 2s^2 2p^4$
Fluorine	9	$1s^2 2s^2 2p^5$
Neon	10	$1s^2 2s^2 2p^6$
Sodium	11	$1s^2 2s^2 2p^6 3s^1$
Magnesium	12	$1s^2 2s^2 2p^6 3s^2$
Aluminum	13	$1s^2 2s^2 2p^6 3s^2 3p^1$
Silicon	14	$1s^2 2s^2 2p^6 3s^2 3p^2$
Phosphorus	15	$1s^2 2s^2 2p^6 3s^2 3p^3$
Sulfur	16	$1s^2 2s^2 2p^6 3s^2 3p^4$
Chlorine	17	$1s^2 2s^2 2p^6 3s^2 3p^5$
Argon	18	$1s^2 2s^2 2p^6 3s^2 3p^6$
Potassium	19	$1s^2 2s^2 2p^6 3s^2 3p^6 4s^1$

In order to avoid having to write unwieldy spectroscopic descriptions of each element’s electron structure, it is customary to write the notation only for subshells that are unfilled. For example, instead of writing the electron structure of the element aluminum¹⁵ as $1s^2 2s^2 2p^6 3s^2 3p^1$, we might just as well write a condensed version showing only the last subshell ($3p^1$), since all the previous subshells are completely full.

Another way to abbreviate the spectroscopic notation for elements is to condense all the shells below the newest (unfilled) shell as the corresponding noble element, in brackets. To use the example of aluminum again, we could write its spectroscopic notation as $[\text{Ne}]3s^2 3p^1$ since its shell 1 and shell 2 configurations are completely described by the electron configuration of Neon.

¹⁴The atomic number is the quantity of protons found in an atom’s nucleus, and may only be a whole number. Since any electrically balanced atom will have the same number of electrons as protons, we may look at the atomic number of an element as being the number of electrons in each atom of that element.

¹⁵Building on the amphitheater analogy for one atom of the element aluminum, we could say that there are two electrons occupying the “s” seating row on the first level, plus two electrons occupying the “s” seating row on the second level, plus six electrons occupying the “p” seating row on the second level, plus two electrons occupying the “s” seating row on the third level, plus one electron occupying the “p” seating row on the third level.

Re-writing our electron shell table for the first nineteen elements using this condensed notation:


Element	Atomic number	Electron configuration
Hydrogen	1	1s ¹
Helium	2	1s ²
Lithium	3	[He]2s ¹
Beryllium	4	[He]2s ²
Boron	5	[He]2s ² 2p ¹
Carbon	6	[He]2s ² 2p ²
Nitrogen	7	[He]2s ² 2p ³
Oxygen	8	[He]2s ² 2p ⁴
Fluorine	9	[He]2s ² 2p ⁵
Neon	10	[He]2s ² 2p ⁶
Sodium	11	[Ne]3s ¹
Magnesium	12	[Ne]3s ²
Aluminum	13	[Ne]3s ² 3p ¹
Silicon	14	[Ne]3s ² 3p ²
Phosphorus	15	[Ne]3s ² 3p ³
Sulfur	16	[Ne]3s ² 3p ⁴
Chlorine	17	[Ne]3s ² 3p ⁵
Argon	18	[Ne]3s ² 3p ⁶
Potassium	19	[Ar]4s ¹

If we progress from element to element in increasing atomic number, we see that no new shell begins to form until after we reach the noble element for that period¹⁶ at the far right-hand column. With the beginning of each new period at the far-left end of the Table, we see the beginning of the next higher-order electron shell. The shell(s) below are represented by whichever noble element shares that same configuration¹⁷, indicating a “noble core” of electrons residing in extremely stable (low-energy) regions around the nucleus.

¹⁶Recall the definition of a “period” in the Periodic Table being a horizontal row, with each vertical column being called a “group”.

¹⁷Building on the amphitheater analogy once again for one atom of the element aluminum, we could say that all seats within levels 1 and 2 are occupied (just like an atom of neon), plus two electrons occupying the “s” seating row on the third level, plus one electron occupying the “p” seating row on the third level.

The beginning of the next higher-order shell is what accounts for the periodic cycle of ionization energies we see in elements of progressing atomic number. The first electron to take residence in a new shell is very easy to remove, unlike the electrons residing in the “noble” configuration shell(s) below:




Noble elements found at the end of each period have the most stable (i.e. difficult-to-ionize) electron configurations.

Not only is the “noble core” notation convenient for tersely describing the electron structure of an element, but it also reveals an important concept in chemistry: the idea of *valence*. Electrons residing in lower-order shells are, by definition, at lower energy states than electrons residing in higher-order shells and are therefore much more difficult to dislodge. Therefore, the electrons in unfilled shells, being easier to dislodge, play a far greater role in chemical bonds than electrons residing in filled shells below. These “outer” electrons are called *valence electrons*, and their number determines how readily an atom will chemically interact with another atom. This is why elements found in the same group (having similar outer-electron configurations) bear similar chemical characteristics: the electrons lying below in the “noble core” configurations have little effect on how the atom will bond with other atoms. A lithium atom, with its outer-most electron configuration being $2s^1$, reacts in much the same way as an atom of sodium having an outer-most configuration of $3s^1$, and much the same as a potassium atom having an outer-most configuration of $4s^1$, all because those outer-shell electrons are the most available for interaction with electrons of other atoms.

If we examine the electron structures of atoms with successively greater atomic numbers (more protons in the nucleus, therefore more electrons in orbit to balance the electrical charge), we notice that the shells and subshells fill up in an interesting pattern. One might think that all the lower-order shells get completely filled before any electrons go into a higher-order shell – just as we might expect people to fill every seat in all the lower tiers of an amphitheater before filling seats in any of the higher tiers – but this is not always the case. Instead, the energy levels of subshells within shells is split, such that certain subshells within a higher shell will have a lower energy value than certain subshells within a lower shell. Referring back to our amphitheater analogy, where seating tiers represented shells and seat rows of various shape represented subshells, it is as though people choose to fill the more comfortable rows in higher-level tiers before sitting in the less-comfortable rows in lowest available tiers, the desire for comfort trumping the desire for proximity to the stage.

A rule commonly taught in introductory chemistry courses called the *Madelung rule* (also referred to as *Aufbau order*, after the German verb *aufbauen* meaning “to build up”) is that subshells fill with increasing atomic number in such an order that the subshell with the lowest $n + l$ value, in the lowest shell, gets filled before any others.

The following graphic illustrates this ordering:



Madelung filling order: $1s \rightarrow 2s \rightarrow 2p \rightarrow 3s \rightarrow 3p \rightarrow 4s \rightarrow 3d \rightarrow 4p \rightarrow 5s \rightarrow 4d \rightarrow 5p \rightarrow 6s \rightarrow 4f \rightarrow 5d \rightarrow 6p \rightarrow 7s \rightarrow 5f \rightarrow 6d \rightarrow 7p \rightarrow 8s \rightarrow (etc.)$

It should be noted that exceptions exist for this rule. We see one of those exceptions with the element chromium (${}_{24}\text{Cr}$). Strictly following the Madelung rule in progressing from vanadium (atomic number = 23, valence electron structure $3d^34s^2$) to chromium (atomic number = 24), we would expect the next electron to take residence in the “3d” subshell making chromium’s valence structure be $3d^44s^2$, but instead we find *two more* electrons residing in chromium’s 3d subshell with one less in the 4s subshell ($3d^54s^1$). The sequence resumes its expected progression with the next element, manganese (atomic number = 25, valence electron structure $3d^54s^2$). The general principle

of energy minimization still holds true . . . it's just that the relative energies of succeeding subshells do not follow a simple rule structure. In other words, the Aufbau order is an over-simplified view of reality. To use the amphitheater analogy again, it's as if someone gave up one of the nice chairs in tier 4 to sit next to a friend who just occupied one of the less comfortable chairs in tier 3.

The practical importance of electron configurations in chemistry is the potential energy possessed by electrons as they reside in different shells and subshells. This is extremely important in the formation and breaking of chemical bonds, which occur due to the interaction of electrons between two or more atoms. A chemical bond occurs between atoms when the outer-most (valence) electrons of those atoms mutually arrange themselves in energy states that are collectively lower than they would be individually. The ability for different atoms to join in chemical bonds completely depends upon the default energy states of electrons in each atom, as well as the next available energy states in the other atoms. Atoms will form stable bonds only if the union allows electrons to assume stable (low-energy) levels. This is why different elements are very selective regarding which elements they will chemically bond with to form compounds: not all combinations of atoms result in favorable potential energy levels.

The amount of energy required to break a chemical bond (i.e. separate the atoms from each other) is the same amount of energy required to restore the atoms' electrons to their previous (default) states before they joined. This is the same amount of energy released by the atoms as they come together to form the bond. Thus, we see the foundation of the general principle in chemistry that forming chemical bonds releases energy, while breaking chemical bonds requires an input of energy from an external source. We also see in this fact an expression of the Conservation of Energy: the amount of energy "invested" in breaking bonds is precisely the same as the amount of energy "returned" when those same bonds re-form.

In summary, the whole of chemistry is a consequence of electrons not being able to assume arbitrary positions around the nucleus of an atom. Instead, they seek the lowest possible energy levels within a framework allowing them to retain unique quantum states. Atoms with mutually agreeable electron structures readily bond together to form molecules, and they release energy in the process of joining. Molecules may be broken up into their constituent atoms, if sufficient energy is applied to overcome the bond. Atoms with incompatible electron structures do not form stable bonds with each other.

3.5 Spectroscopy

Much of our knowledge about atomic structure comes from experimental data relating the interaction between *light* and atoms of the different elements. Light may be modeled as an electromagnetic wave, consisting of an oscillating electric field and an oscillating magnetic field. Like any wave, the relationship between propagation velocity, wavelength, and frequency is described by the following equation:

$$v = \lambda f$$

Where,

v = Velocity of propagation (e.g. meters per second)

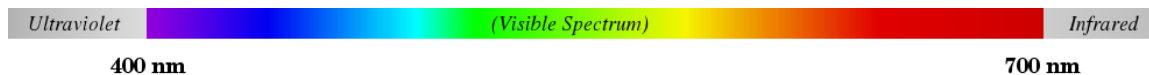
λ = Wavelength (e.g. meters)

f = Frequency of wave (e.g. Hz, or 1/seconds)

When applied to light waves, this equation is typically written as $c = \lambda f$, where c is the speed of light in a vacuum ($\approx 3 \times 10^8$ meters per second): one of the fundamental constants of physics.

Light that is visible to the human eye has wavelengths approximately between 400 nm (400 *nanometers*) at the violet end of the spectrum and 700 nm at the red end of the spectrum. Given the speed of light, this equates to a frequency range for visible light between 7.5×10^{14} Hz and 4.286×10^{14} Hz.

A computer-generated image of the visible light spectrum (plus the ultraviolet and infrared regions outside of the visible range, shown in grey) appears here. A real spectrum may be generated by taking “white” light and passing it through either a prism or a diffraction grating so that the different wavelengths separate from each other:



Just as buoyant objects are moved up and down by waves of water, electrically-charged objects may be moved about by waves of electrical fields such as light. In the case of electrons, their positions around the nucleus of an atom may be altered if struck by light of the necessary wavelength.

One of the major breakthrough discoveries of modern physics was the realization that a ray of light could be modeled as a stream of *particles* – each of these “photon” particles possessing a definite amount of energy – in addition to being modeled as a continuous *wave* possessing a definite frequency. The combined work of physicists Max Planck in 1900 and Albert Einstein in 1905 resulted in the following equation relating a photon’s energy to its frequency:

$$E = hf$$

Where,

E = Energy carried by a single “photon” of light (joules)

h = Planck’s constant (6.626×10^{-34} joule-seconds)


f = Frequency of light wave (Hz, or 1/seconds)

We may re-write this equation to express a photon’s energy in terms of its wavelength (λ) rather than its frequency (f), knowing the equation relating those two variables for waves of light ($c = \lambda f$):

$$E = \frac{hc}{\lambda}$$

Physicists knew that light carried energy, but now they understood that the energy carried by a beam of light was finely divided into fixed (“quantized”) amounts corresponding to the wavelength of each particle-wave (photon). That is to say, a beam of monochromatic (single-color, or single-wavelength) light consists of photons having exactly the same energies, and total beam power is simply a matter of how many of those photons per second pass by. Varying the intensity of a monochromatic light beam without changing its wavelength (color) only changes the number of photons per second, not the amount of energy carried by each photon.


If the amount of energy carried by a photon happens to match the energy required to make an atomic electron “jump” from one energy level to another within the atom, the photon will be consumed in the work of that task when it strikes the atom. Conversely, when that “excited” electron returns to its original (lower) energy level in the atom, it releases a photon of the same frequency as the original photon that excited the electron:



Since the energy levels available for an electron to “jump” within an atom are limited to certain fixed values by virtue of the atom’s shell and subshell structure, this means only certain specific frequencies or wavelengths of light will be able to make an electron of a particular atom move to new shells and/or subshells¹⁸. A startling consequence of this *quantum theory* of light was that the ability of a light beam to dislodge electrons from an atom depended on the *color* (wavelength or frequency) of the photons, and not the intensity (total power) of the light beam. A light beam consisting of photons with insufficient individual energy (i.e. frequency too low; wavelength too long; color too far shifted toward red if visible) is incapable of boosting electrons from a lower energy level to a higher energy level, no matter how intense that beam may be. This is analogous to shooting an armored target with slow-moving bullets: so long as the velocity (kinetic energy) of each bullet is insufficient to penetrate the armor, it does not matter how many of those low-energy bullets are fired at the target, or how frequently they are fired. However, just a single bullet with sufficient kinetic energy will be sufficient to penetrate the armor.

¹⁸This is the reason silicon-based photovoltaic solar cells are so inefficient, converting only a fraction of the incident light into electricity. The energy levels required to create an electron-hole pair at the P-N junction correspond to a narrow portion of the natural light spectrum. This means most of the photons striking a solar cell do *not* transfer their energy into electrical power because their individual energy levels are insufficient to create an electron-hole pair in the cell’s P-N junction. For photovoltaic cells to improve in efficiency, some way must be found to harness a broader spectrum of photon frequencies (light colors) than silicon P-N junctions can do, at least on their own.


The discovery of photons having discrete energy values was a major shift in scientific thought, setting physics down a new path of understanding matter and energy in *quantum* terms. It was this new quantum theory of matter and energy that led to the modern understanding of atomic electron structure, with all its shells, subshells, and orbitals. Later mathematical contributions to quantum theory from physicists such as Louis de Broglie, Werner Heisenberg, and especially Erwin Schrödinger provided tools to calculate the probability distributions of electrons within atoms. The oddly-shaped orbital electron “clouds” discussed earlier in this chapter are in fact solutions to Schrödinger’s wave equation for electrons at different energy levels:



This is why the notation used in the previous section to describe electron configurations (e.g. $1s^2 2s^2 2p^1$) is called *spectroscopic* notation: the discovery of shells, subshells, and orbitals owes itself to the analysis of light wavelengths associated with different types of atoms, studied with a device called a *spectroscope* constructed to analyze the wavelengths of light across the visible spectrum. Just as the telescope was the first tool scientists used to explore outer space, the spectroscope was one of the first tools used by scientists to explore the “inner space” of atomic structure.

3.5.1 Emission spectroscopy

If we take a sample of atoms, all of the same element and at a low density¹⁹ (e.g. a gas or vapor), and “excite” them with a source of energy such as an electric arc, we will notice those atoms emit colors of light that are characteristically unique to that element:



This phenomenon is used to make colored discharge (“neon”) lights. While neon gas glows with a characteristic pink-orange color, other gases glow with their own signature colors. By filling glass tubes with the right gas(es), a wide variety of colors may be produced.

These colors are unique to their respective gases because the unique electron configurations of each element creates a unique set of energy values between which atomic electrons of that element may “jump.” Since no two elements have the exact same electron configurations, no two elements will have the exact same set of available energy levels for their electrons to occupy. When excited electrons fall back into lower shell levels, the photons they emit will have distinct wavelengths. The result is an *emission spectrum* of light wavelengths, much like a “fingerprint” unique to that element. Indeed, just as fingerprints may be used to identify a person, the spectrum of light emitted by an “excited” sample of an element may be used to identify that element.

For example, we see here the emission spectrum for *hydrogen*, shown immediately below the continuous spectrum of visible light for convenient reference²⁰:




Each of the colored “lines” in the emission spectrum for hydrogen represents the photon wavelength emitted when the excited electron loses energy and falls back into a lower-level position. The larger the energy difference between energy levels (i.e. the bigger the jump), the more energy the photon carries away, and consequently the shorter the wavelength (higher the frequency) of the photon. The violet color line, therefore, represents one of the larger “jumps” while the red color line represents one of the smaller. Hydrogen happens to emit four different wavelengths within the visible range (656 nm, 486 nm, 434 nm, and 410 nm), and many others outside the visible range.

¹⁹Solids and liquids tend to emit a broad spectrum of wavelengths when heated, in stark contrast to the distinct “lines” of color emitted by isolated atoms.

²⁰To create these spectra, I used a computer program called *Spectrum Explorer*, or SPEX.

This next illustration shows a simplified view of a hydrogen atom, with the lowest-level shell ($n = 1$, K) representing the ground state and higher-level shells representing “excited” energy states for its single electron:




Wavelengths of light emitted when an excited electron falls from any high-level shell down to the second shell of hydrogen ($n = 2$; L) are called the *Balmer series* of spectral lines. The four wavelengths previously mentioned are Balmer lines visible to the human eye: 410 nm resulting from an electron jumping from the sixth shell ($n = 6$; P) to the second shell, 434 nm resulting from a transition between the fifth and second shells, 486 nm from a transition between the fourth and second shells, and finally the 656 nm wavelength resulting from a transition between the third and second shells. Other Balmer-series wavelengths exist²¹ (electrons transitioning from even higher shells than the sixth, down to the second), but these wavelengths lie within the ultraviolet range and are therefore not visible to the human eye. Note the inverse relationship between jump distance and wavelength: the shortest “jump” (shell 3 to shell 2) yields the photon with the longest wavelength (656 nm). This is because the shortest jump represents the smallest energy change, which then

²¹Including wavelengths of 397 nm, 389 nm, and 384 nm.

results in a photon of comparatively little energy, having a low frequency and therefore a long wavelength.

You will note that the Balmer series of wavelengths do not involve an electron falling all the way back to the hydrogen atom's "ground state" (the normal, or un-excited state of shell $n = 1$, the "K" shell). Electrons falling down to the first shell ($n = 1$; K) from any higher-level shells will also emit photons, but these photons will be of a far shorter wavelength (higher frequency, higher energy) than any in the Balmer series, owing to the larger energy gap between the first shell and all the others. This so-called *Lyman series* of light wavelengths lies within the region of wavelengths referred to as "far-ultraviolet," well outside the range of human vision.

This next graphic shows the emission spectra of several elements contrasted against a continuous spectrum covering both visible light and portions of the ultraviolet and infrared ranges:




Note how complex the emission spectra are for some of the elements. Since we know each spectral line represents a unique change in energy (i.e. a unique "jump distance" from one energy level to another), the multitude of lines we see for each element shows us the range of "jumps" possible within certain atoms. Note also how spectral lines for most elements (including hydrogen) extend past the visible light range. Lines in the ultraviolet range comes from large electron transitions, as electrons fall from high-level shells to low-level shells and lose much energy. Lines in the infrared range originate from small electron transitions, as electrons transition between adjacent shells and lose little energy.

Not only may the wavelengths of photons emitted from "excited" electrons returning to lower-energy conditions be used to positively identify different elements, but we may also use those wavelengths as universal standards, since the fundamental properties of elements are not liable to change. For example, the SI (Système International) definition for the base unit of the *meter* is standardized as 1650763.73 wavelengths of light emitted by a krypton-86 (^{86}Kr) atom as its electrons transition between the $2p^{10}$ and $5d^5$ subshells²².

²²The wavelength of this light happens to lie within the visible range, at approximately 606 nm. Note the shell levels involved with this particular electron transition: between $2p^{10}$ and $5d^5$. Krypton in its ground (un-excited) state has a valence electron configuration of $4p^6$, which tells us the electron's transition occurs between an inner shell of the Krypton atom and an excited shell (higher than the ground-state outer shell of the atom). The wavelength of this photon (606 nm) resulting from a shell 5 to shell 2 transition also suggests different energy levels for those shells of a Krypton atom compared to shells 5 and 2 of a hydrogen atom. Recall that the Balmer line corresponding to a transition from $n = 5$ to $n = 2$ of a hydrogen atom had a wavelength value of 434 nm, a higher energy than 606 nm and therefore a larger jump between those corresponding shells.

3.5.2 Absorption spectroscopy

If we take a sample of atoms, all of the same element and at a low density (e.g. a gas or vapor), and pass a continuous (“white”) spectrum of light wavelengths through that sample, we will notice certain colors of light *missing* from the light exiting the sample:



Not only are these missing wavelengths characteristically unique to that element, but they are *the exact same wavelengths of light found in the emission spectrum for that element!* The same photon wavelengths produced by an atom when “excited” by an external energy source will be readily *absorbed* by that atom if exposed to them. Thus, the spectrum of light missing characteristic wavelengths after passing through a gas sample is called an *absorption spectrum*, and may be used to identify elements just as easily²³ as an emission spectrum.

The absorption spectrum of hydrogen gas is shown at the bottom of this three-spectrum graphic image, contrasted against the continuous spectrum of visible light (top) and the emission spectrum for hydrogen (middle):



Note how the four colored lines in the emission spectrum characteristic of hydrogen appear as *missing* colors (black lines) in the absorption spectrum. It is almost as though one hydrogen spectrum were a photographic “negative” of the other: each of the colors present in the emission spectrum is distinctly absent²⁴ in the absorption spectrum. Although the color patterns may be inverted, the positions of the lines within the spectrum are the same, and are *uniquely* representative of hydrogen.

The effect is analogous to fingerprints made two different ways: one by pressing a pre-inked finger onto a clean sheet of paper; the other by pressing a clean finger onto pre-inked paper. In the first method, the result is a set of dark ink-marks where the fingerprint ridges touched the paper to apply ink and light areas where skin and paper never touched. In the second method, the result is a set

²³In fact, it is often easier to obtain an absorption spectrum of a sample than to create an emission spectrum, due to the relative simplicity of the absorption spectrometer test fixture. We don’t have to energize a sample to incandescence to obtain an absorption spectrum – all we must do is pass white light through enough of it to absorb the characteristic colors.

²⁴One student described this to me as a “shadow” image of the hydrogen gas. The missing colors in the absorption spectrum are the *shadows* of hydrogen gas molecules blocking certain frequencies of the incident light from reaching the viewer.

of *inverse* ink-marks: light where the fingerprint ridges touched the paper to remove ink and dark where skin and paper never touched. The fingerprint patterns in both cases – if made using the same finger – will be identical in form, just inverted in color. Likewise, the *patterns* of emission and absorption spectroscopy will be the same for any given substance, just inverted in color: emission spectroscopy shows select wavelengths against an otherwise dark field, while absorption spectroscopy shows a nearly-full spectrum of color missing (the same) select wavelengths.

Individual atoms are not the only forms of matter possessing uniquely identifying spectra – many *molecules* have spectral “signatures” of their own as well. The absorption spectra for molecular substances are substantially more complex than the absorption spectra of pure elements, owing to the many more different ways in which light energy may be absorbed by a molecule. In addition to electron shell and subshell “jumps” capable of absorbing a photon’s energy, the atoms within a molecule are also able to vibrate, rotate, and twist about each other like mechanical oscillators. Photons of light possessing just the right frequencies are able to “excite” certain molecules in a manner not unlike AC electrical waveforms resonating with tuned LC (inductor-capacitor) circuits. Just as tuned LC circuits absorb and store energy at certain frequencies, molecular oscillators absorb and store energy from photons.

The multiplicity of energy-absorbing modes for certain molecules gives them wide *bands* of absorption in the light spectrum, not just thin “lines” as is the case with individual atoms. These bands are still unique to each molecule type, but they typically cover a far broader swath of wavelengths than is typical for atomic absorption spectra.

The absorption of ultraviolet light by ozone gas (O_3) high in Earth’s atmosphere is an example of absorption spectroscopy on a grand scale. These molecules serve as a protective “blanket” against ultraviolet light rays from the sun which have detrimental effects on life (e.g. sunburn, skin cancer). The ozone does not absorb light in the visible spectrum, and so its protective effects are not visually apparent, but the attenuation of ultraviolet light is definitely measurable. This attenuation also covers far more than just one or two specific wavelengths of ultraviolet light, which is good for life on Earth because otherwise ozone wouldn’t offer much protection.

Many chemical substances of interest in process industries have well-known *absorption signatures* for ultraviolet and infrared light. This makes spectroscopy a powerful tool for the identification (and quantitative measurement) of chemical composition in process fluids, exhaust gases, and sometimes even in solid materials. For more detail on the practical application of spectroscopy to analytical measurement, refer to section 23.4 beginning on page 1805.

An interesting application of optical absorption is the detection of gas leaks using an infrared camera. Many industrial gases are strong absorbers of infrared light, which means if a leaking pipe or vessel is viewed through a camera sensitized to infrared light and there is sufficient ambient infrared light for viewing, the leaking gas will appear on the camera’s image as a dark cloud. The gas plume appears on the camera’s display the way steam or smoke appears to the naked eye. Several paraffinic hydrocarbon compounds such as methane, ethane, propane, butane, pentane, and hexane are detectable with infrared cameras sensitized to light wavelengths of 3.3 to 5 micrometers (μm). Infrared cameras sensitized to longer wavelengths of light (10 μm to 11 μm) are useful for detecting leaks of gases such as sulfur hexafluoride, ammonia, chlorine dioxide, FREON-12, and ethylene to name a few.

3.6 Formulae for common chemical compounds

Most of these formulae appear in *molecular chemical* form rather than structural form. For example, ethanol appears here as C_2H_6O rather than C_2H_5OH . Also, the entries for fructose and glucose are identical ($C_6H_{12}O_6$) despite the two compounds having different structures. This means most of the formulae shown in this section merely represent the ratios of each element in a compound, making little or no attempt to convey the *structure* of the molecule.

It should be noted that this list is definitely *not* exhaustive, but merely attempts to show formulae for some common compounds.

- Acetone: C_3H_6O
- Acetylene: C_2H_2
- Alcohol, methyl (methanol): CH_4O
- Alcohol, ethyl (ethanol): C_2H_6O
- Alcohol, isopropyl (isopropanol): C_3H_8O
- Alcohol, butyl (butanol): $C_4H_{10}O$
- Alcohol, phenol: C_6H_6O
- Aluminum oxide (alumina): Al_2O_3
- Ammonia: NH_3
- Ammonium carbonate: $(NH_4)_2CO_3$
- Ammonium chloride (sal ammoniac): NH_4Cl
- Ammonium nitrate: $N_2H_4O_3$
- Aromatic hydrocarbons:
 - Acetylene: C_2H_2
 - Ethylene: C_2H_4
 - Propylene: C_3H_6
 - Butylene: C_4H_8
 - Benzene: C_6H_6
 - Toluene: C_7H_8
 - Styrene: C_8H_8
 - Napthalene: $C_{10}H_8$
- Calcium carbonate (limestone, marble): $CaCO_3$
- Calcium chloride: $CaCl_2$
- Calcium hydroxide: $Ca(OH)_2$

- Calcium oxide (lime or quicklime): CaO
- Calcium sulfate (gypsum): CaSO_4
- Carbon monoxide: CO
- Carbon dioxide: CO_2
- Carbon tetrachloride: CCl_4
- Carbonic acid: H_2CO_3
- Cellulose: $(\text{C}_6\text{H}_{10}\text{O}_5)_n$
- Clay (or shale): $\text{H}_4\text{Al}_2\text{Si}_2\text{O}_9$
- Copper oxide (cuprite): Cu_2O
- Copper oxide (tenorite): CuO
- Cyanic acid: HOCN
- Dextrose (synonym for biological glucose): $\text{C}_6\text{H}_{12}\text{O}_6$
- Ethyl mercaptan: $\text{C}_2\text{H}_6\text{S}$
- Ethylene glycol: $\text{C}_2\text{H}_6\text{O}_2$
- Ethylene oxide: $\text{C}_2\text{H}_4\text{O}$
- Ferrous chloride: FeCl_2
- Ferric chloride: FeCl_3
- Formaldehyde: CH_2O
- Folic acid: $\text{C}_{19}\text{H}_{19}\text{N}_7\text{O}_6$
- Formaldehyde: CH_2O
- Formic acid: CH_2O_2
- Fructose (same molecular formula as glucose): $\text{C}_6\text{H}_{12}\text{O}_6$
- Glycerol: $\text{C}_3\text{H}_8\text{O}_3$
- Hydrazine: N_2H_4
- Hydrocyanic acid: HCN
- Hydrofluoric acid: HF
- Hydrochloric acid: HCl
- Hydrogen peroxide: H_2O_2

- Hydrogen sulfide: H_2S
- Iron oxide: Fe_2O_3
- Magnesium hydroxide (milk of magnesia): $\text{Mg}(\text{OH})_2$
- Nitric acid: HNO_3
- Nitric oxide: NO
- Nitrogen dioxide: NO_2
- Nitrogen trioxide: NO_2
- Nitroglycerine: $\text{C}_3\text{H}_5\text{N}_3\text{O}_9$
- Nitromethane: CH_3NO_2
- Nitrous oxide: N_2O
- Dinitrogen dioxide: N_2O_2
- Dinitrogen trioxide: N_2O_3
- Ozone: O_3
- Paraffinic hydrocarbons:
 - Methane: CH_4
 - Ethane: C_2H_6
 - Propane: C_3H_8
 - Butane: C_4H_{10}
 - Pentane: C_5H_{12}
 - Hexane: C_6H_{14}
 - Heptane: C_7H_{16}
 - Octane: C_8H_{18}
 - Nonane: C_9H_{20}
 - Decane: $\text{C}_{10}\text{H}_{22}$
- Phosgene: COCl_2
- Phosphoric acid: H_3PO_4
- Potassium chloride: KCl
- Potassium cyanide: KCN
- Potassium hydroxide: KOH
- Potassium sulfate: K_2SO_4

- Silane: SiH_4
- Silica: SiO_2
- Silicon carbide: SiC
- Sodium chloride (table salt): NaCl
- Sodium hydroxide: NaOH
- Sodium fluoride: NaF
- Strychnine: $\text{C}_{21}\text{H}_{22}\text{N}_2\text{O}_2$
- Sucrose: $\text{C}_{12}\text{H}_{22}\text{O}_{11}$
- Sulfuric acid: H_2SO_4
- Sulfur dioxide: SO_2
- Sulfur hexafluoride: SF_6
- Testosterone: $\text{C}_{19}\text{H}_{28}\text{O}_2$
- Turpentine: $\text{C}_{10}\text{H}_{16}$ (approx.)
- Zinc sulfate: ZnSO_4

3.7 Molecular quantities

Sample sizes of chemical substances are often measured in *moles*. One mole of a substance is defined as a sample having 6.022×10^{23} (*Avogadro's number*) molecules²⁵. This number is not arbitrary – it was chosen²⁶ such that 1 mole of carbon-12 (6.022×10^{23} individual ^{12}C atoms together in one sample) would have a mass of exactly 12 grams. In other words, Avogadro's number is a proportionality between an element's atomic mass (measured in *amu* or *Daltons*) and the mass of a sample (measured in *grams*).

With Avogadro's number defined as such, we may interpret any element's atomic mass value as a conversion factor relating moles to grams. For example, if we look up the element *potassium* in a periodic table, we see that it has an average atomic mass of 39.0983 amu (39.0983 Daltons) as found in nature. This means 1 mole of naturally-occurring potassium atoms equals 39.0983 grams of mass. Likewise, 5 moles of potassium atoms will have a mass of 195.4915 grams. Note the use of the equivalence 1 mol potassium = 39.0983 g as a “unity fraction” in the following calculation, used to cancel the given unit of moles to yield an answer in the unit of grams:

$$\left(\frac{5 \text{ mol potassium}}{1}\right) \left(\frac{39.0983 \text{ g}}{1 \text{ mol potassium}}\right) = 195.4915 \text{ g}$$

Molar quantities make it convenient to relate macroscopic samples of elements and compounds with each other. We know, for instance, that one mole of naturally occurring iron (Fe) atoms will have a mass of 55.8 grams, and that one mole of naturally occurring oxygen (O) atoms will have a mass of 16.0 grams, because the average atomic mass of naturally occurring iron is 55.8 amu, and the average atomic mass of naturally occurring oxygen is 16.0 amu. One mole of naturally occurring oxygen *molecules* (O_2) will have a mass of 32.0 grams, since each molecule is a *pair* of oxygen atoms at 16 amu each, and “moles” counts the number of discrete entities which in the case of molecular oxygen is the number of O_2 *molecules* rather than the number of O *atoms*. Applying the same reasoning, one mole of ozone (O_3) molecules will have a mass of 48.0 grams.

The same mathematical proportions apply to compounds as they do to elements, since compounds are nothing more than different elements bound together in whole-number ratios, and the Conservation of Mass tells us a molecule cannot have a mass greater or less than the sum total of the constituent elements' masses. To illustrate this principle, we may calculate the mass of one mole of iron oxide (Fe_2O_3), the principal component of *rust*: $55.8 \times 2 + 16.0 \times 3 = 159.6$ grams. Likewise, we may calculate the mass of five moles of pure glucose ($\text{C}_6\text{H}_{12}\text{O}_6$): $5 \times (12.01 \times 6 + 1.01 \times 12 + 16.0 \times 6) = 900.9$ grams. The sum of the atomic masses of a molecule's constituent atoms is called the *molecular weight* or *formula weight* for that molecule. In the case of iron oxide, the molecular

²⁵Truth be told, a “mole” is 602,200,000,000,000,000,000,000 counts of literally *any* discrete entities. Moles do not represent mass, or volume, or length, or area, but rather a *quantity of individual units*. There is nothing wrong with measuring the amount of eggs in the world using the unit of the mole, or the number of grains of sand in moles, or the number of bits in a collection of digital data. Think of “mole” as nothing more than a *really* big dozen, or more precisely, a really big *half-dozen*!

²⁶Another way to define one mole is that it is the number of individual nucleons (i.e. protons and/or neutrons) necessary to comprise one gram of mass. Since protons and neutrons comprise the vast majority of an atom's mass, we may essentially ignore the mass of an atom's electrons when tabulating its mass and pay attention only to the nucleus. This is why one mole of Hydrogen atoms, each atom having just one lone proton in its nucleus, will have a combined mass of one gram. By extension, one mole of Carbon-12 atoms, each atom with 6 protons and 6 neutrons, will have a combined mass of twelve grams.

weight is 159.6 (typically rounded up to 160 grams per mole). In the case of glucose, the molecular weight is 180.18 (typically rounded down to 180 grams per mole).

When referring to liquid solutions, the concentration of a solute is often expressed as a *molarity*, defined as the number of moles of solute per liter of solution. Molarity is usually symbolized by an italicized capital letter *M*. It is important to bear in mind that the volume used to calculate molarity is that of the total solution (solute plus solvent) and not the solvent alone.

Suppose we had a solution of salt-water, comprised of 33.1 grams of table salt thoroughly mixed with pure water to make a total volume of 1.39 liters. In order to calculate the molarity of this solution, we first need to determine the equivalence between moles of salt and grams of salt. Since table salt is sodium chloride (NaCl), and we know the atomic masses of both sodium (23.0 amu) and chlorine (35.5 amu), we may easily calculate the mass of one mole of table salt:

$$1 \text{ mole of NaCl} = 23.0 \text{ g} + 35.5 \text{ g} = 58.5 \text{ g}$$

Another way to state this is to say that table salt (sodium chloride, or NaCl) has a molecular weight of 58.5 amu (58.5 grams of mass per mole).

We may use this equivalence as a unity fraction to help us convert the number of grams of table salt per unit volume of solution into a molarity (moles of table salt molecules per liter):

$$\left(\frac{33.1 \text{ g}}{1.39 \text{ l}} \right) \left(\frac{1 \text{ mol NaCl}}{58.5 \text{ g}} \right) = 0.407 \frac{\text{mol NaCl}}{\text{l}} = 0.407 \text{ M (NaCl)}$$

Another common expression for the concentration of a solute in either a liquid or a gas solution is related to the concept of *percent*, expressing the presence of the solute as a ratio of how many “parts” of solute exist per “parts” of solution. Earth’s atmosphere, for example, contains approximately 20.9% oxygen gas by volume. This means that for every 100 molecules found in a sample of air, approximately 21 of those are oxygen molecules. When the concentration of a solute is very small, however, percent becomes an awkward unit of measurement. In such cases it is common to see low concentrations of solute expressed as *parts per million* (ppm) or even *parts per billion* (ppb). The volumetric concentration of methane gas in Earth’s atmosphere is a good example where parts-per-million is a more appropriate expression than percent: for every million molecules found in a sample of air, approximately 2 of them are methane molecules (i.e. methane has an atmospheric concentration of 2 ppm). As a percentage, this equates to only 0.0002%.

We may use parts-per-unit concentration values as unity fractions just like molecular weights and just like molarity values, to relate total solution quantity to solute quantity. For example, if we need to calculate the total mass of hydrogen gas in a compressed air cylinder storing 47000 standard cubic feet of air, we could multiply the total volume of that air sample (47000 SCF) by the volumetric concentration of hydrogen naturally found in Earth’s atmosphere which is 0.5 ppm:

$$\left(\frac{47000 \text{ SCF air}}{1} \right) \left(\frac{0.5 \text{ parts hydrogen}}{1000000 \text{ parts air}} \right) = 0.0235 \text{ SCF hydrogen}$$

Note how the units of “parts” and “air” cancel out to leave “SCF hydrogen”.



An important caveat when using percent, ppm, or ppb is that we must clearly define the “parts” proportion as either being volume or mass. The concentration of hydrogen gas in the atmosphere (0.5 ppm) was specified as a *volumetric* concentration, and so it is appropriate to use this 0.5 ppm figure to calculate a proportion of the 47000 SCF total *volume*. If, however, we were given a ppm *mass* concentration for hydrogen, we could only use that figure in conjunction with a total *mass* quantity for the air sample.

The following photograph illustrates this concept, showing the label on a *calibration gas bottle* containing a certified mixture of gases used to check the accuracy of air-safety monitoring instruments. Note the concentrations of each gas type within the mixture – some expressed in percent, others in ppm – and how the label states “MOLE %” in the upper-right corner using large bold print to let you know the concentration values refer to molar quantities (e.g. 18% oxygen means 18% of the *molecules* contained in this bottle are oxygen molecules), which for gases closely corresponds to volumetric quantity rather than mass:



3.8 Stoichiometry

Stoichiometry is the accounting of atoms before and after a chemical reaction. It is an expression of the *Law of Mass Conservation*, in that elements are neither created nor destroyed in a chemical reaction, and that mass is an intrinsic property of every element. Thus, the numbers, types of atoms, and total mass exiting a chemical reaction (i.e. the “products” of that reaction) must be the same as the numbers, types of atoms, and total mass entering that chemical reaction (i.e. the “reactants”). For example, in the combustion of natural gas in an oxygen-rich environment, the fuel (CH_4) and oxidizer (O_2) are the reactants, while water vapor (H_2O) and carbon dioxide gas (CO_2) are the products:




Reactants	Products	Mass (per mole of CH_4)
Carbon = 1×1	Carbon = 1×1	12 grams
Hydrogen = 1×4	Hydrogen = 2×2	4 grams
Oxygen = 2×2	Oxygen = $(1 \times 2) + (2 \times 1)$	64 grams

As you can see in this example, every single reactant atom (and its mass) entering the reaction is accounted for in the product molecules. The only exception to this rule is in *nuclear reactions* where elements transmute into different elements, with gains or losses in nuclear particles. No such transmutation occurs in any mere *chemical* reaction, and so we may safely assume equal numbers and types of atoms before and after any chemical reaction. Chemical reactions strictly involve re-organization of molecular bonds, with electrons as the constituent particles comprising those bonds. Nuclear reactions involve the re-organization of atomic nuclei (protons, neutrons, etc.), with far greater energy levels associated.

Often in chemistry, we know both the reactant and product molecules, but we need to determine their relative numbers before and after a reaction. The task of writing a general chemical equation and then assigning multiplier values for each of the molecules is called *balancing the equation*.

3.8.1 Balancing chemical equations by trial-and-error

Balancing a chemical equation is a task that may be done by trial-and-error. For example, let us consider the case of complete combustion for the hydrocarbon fuel *ethane* (C_2H_6) with oxygen (O_2). If combustion is complete, the only products will be water vapor (H_2O) and carbon dioxide (CO_2). The unbalanced equation representing all reactants and products for this reaction is shown here, along with a table showing the numbers of atoms on each side of the equation:




Reactants	Products
Carbon = 2	Carbon = 1
Hydrogen = 6	Hydrogen = 2
Oxygen = 2	Oxygen = 3

Clearly, this is not a balanced equation, since the numbers of atoms for each element are unequal between the two sides of the equation.


A good place to start in balancing this equation is to look for an element represented by only one molecule on each side of the equation. Carbon is an example (present in the ethane but not in the oxygen molecule on the left-hand side, and in the carbon dioxide but not the water on the right-hand side) and hydrogen is another.

Beginning with carbon, we see that each ethane molecule contains two carbon atoms while each carbon dioxide molecule contains just one carbon atom. Therefore, we may conclude that the ratio of carbon dioxide to ethane must be 2-to-1, no matter what the other ratios might be. So, we double the number of carbon dioxide molecules on the right-hand side and re-check our atomic quantities:



Reactants	Products
Carbon = 2	Carbon = 2
Hydrogen = 6	Hydrogen = 2
Oxygen = 2	Oxygen = 5


Next, we will balance the hydrogen atom numbers, since we know hydrogen is an element found in only one molecule on each side of the equation. Our hydrogen ratio is now 6:2 (left:right), so we know we need three times as many hydrogen-containing molecules on the right-hand side. Tripling the number of water molecules gives us:



Reactants	Products
Carbon = 2	Carbon = 2
Hydrogen = 6	Hydrogen = 6
Oxygen = 2	Oxygen = 7


Unfortunately, the numbers of oxygen atoms on each side of the equation are unequal, and it is not immediately obvious how to make them equal. We need five more atoms of oxygen on the left-hand side, but we cannot add exactly five more because oxygen atoms only come to us in pairs (O_2), limiting us to even-number increments.

However, if we *double* all the other molecular quantities, it will make the disparity of oxygen atoms an even number instead of an odd number:



Reactants	Products
Carbon = 4	Carbon = 4
Hydrogen = 12	Hydrogen = 12
Oxygen = 2	Oxygen = 14

Now it is a simple matter to balance the number of oxygen atoms, by adding six more oxygen molecules to the left-hand side of the equation:




Reactants	Products
Carbon = 4	Carbon = 4
Hydrogen = 12	Hydrogen = 12
Oxygen = 14	Oxygen = 14

Now the equation is balanced: the quantities of each type of atom on both sides of the equation are equal.

3.8.2 Balancing chemical equations using algebra

A more mathematically sophisticated approach to stoichiometry involves the use of *simultaneous systems of linear equations*. The fundamental problem chemists must solve when balancing reaction equations is to determine the ratios of reactant and product molecules. If we assign a variable to each molecular quantity, we may then write a mathematical equation for each element represented by the reaction, and use algebra to solve for the variable values.

To illustrate, let us balance the equation describing the attack of aluminum metal's protective "passivation" layer of oxide by acid rain. When aluminum metal is exposed to oxygen, the outer surface of the metal quickly forms a layer of aluminum oxide (Al_2O_3) which acts to impede further oxidation of the metal. This protective layer, however, may be attacked by the presence of sulfuric acid (H_2SO_4). This acid finds its way into rainwater by way of sulfur compounds emitted during the combustion of sulfur-laden fuels. The products of this reaction between sulfuric acid and aluminum oxide are a sulfate molecule ($\text{Al}(\text{SO}_4)_3$) and water (H_2O), as illustrated in this *unbalanced* chemical equation:



This equation contains four different compounds (acid, aluminum oxide, sulfate, and water), which means we ultimately must solve for four different multiplier quantities. It also contains four different elements (H, S, O, and Al). Since the mathematical requirement for solving a system of linear equations is to have at least one equation per variable, it would first appear as though we could set up a 4×4 matrix (four equations of four variables). However, this will not work. If we tried to solve for four unknown quantities, we would ultimately be foiled by an infinite number of solutions. This makes sense upon further inspection, since any stoichiometric solution to this chemical reaction will have an infinite number of correct *proportions* to satisfy it²⁷. What we need to do is arbitrarily set one of these molecular quantities to a constant value (such as 1), then solve for the quantities of the other three. The result will be ratios or proportions of all the other molecules to the fixed number we assigned to the one molecule type.

²⁷Take the combustion of hydrogen and oxygen to form water, for example. We know we will need two H_2 molecules for every one O_2 molecule to produce two H_2O molecules. However, *four* hydrogen molecules combined with *two* oxygen molecules will make *four* water molecules just as well! Similarly, *six* hydrogen molecules combined with *three* oxygen molecules also perfectly balance, making *six* water molecules. So long as we consider all three molecular quantities to be unknown, we will never be able to solve for just *one* correct answer, because there is no one correct set of absolute quantities, only one correct set of *ratios* or *proportions*.

As an example, I will choose to set the number of acid molecules to 1, and use the variables x , y , and z to solve for the numbers of the other molecules (oxide, sulfate, and water, respectively):

1	x	=	y	z
H ₂ SO ₄	Al ₂ O ₃	→	Al ₂ (SO ₄) ₃	H ₂ O

Now, I will write four algebraic equations, each algebraic equation representing the stoichiometric balance of a single element in the chemical equation. Focusing on the element hydrogen as an example, there will be *two* hydrogen atoms for every *one* molecule of acid, $0x$ hydrogen atoms for every x molecules of aluminum oxide, $0y$ atoms of hydrogen for every y molecules of aluminum sulfate, and $2z$ atoms of hydrogen for every z molecules of water. The following table shows each of the four elements with their respective balance equations:

Element	Balance equation
Hydrogen	$2 + 0x = 0y + 2z$
Sulfur	$1 + 0x = 3y + 0z$
Oxygen	$4 + 3x = 12y + 1z$
Aluminum	$0 + 2x = 2y + 0z$

Simplifying each equation by eliminating all zero values and “1” coefficients:

Element	Balance equation
Hydrogen	$2 = 2z$
Sulfur	$1 = 3y$
Oxygen	$4 + 3x = 12y + z$
Aluminum	$2x = 2y$


We can see by examination of the first, second, and fourth equations that z must be equal to 1, y must be equal to $\frac{1}{3}$, and that x and y are equal to each other (therefore, x must be equal to $\frac{1}{3}$ as well). Plugging these values into the variables of the third equation confirms this ($4 + 1 = 4 + 1$). Thus, our solution to this multi-variable system of equations is:

$$x = \frac{1}{3} \quad y = \frac{1}{3} \quad z = 1$$

It makes little sense to speak of *fractions* of a molecule, which is what the values of x and y seem to suggest, but we must recall these values represent *proportions* only. In other words, we need only one-third as many oxide and sulfate molecules as acid and water molecules to balance this equation. If we multiply all these values by three (as well as the initial constant we chose for the number of acid molecules), the quantities will be whole numbers and the chemical reaction will still be balanced:

$$x = 1 \quad y = 1 \quad z = 3$$

Thus, our final (balanced) equation showing the attack of aluminum metal’s passivation layer by acid rain is as follows:



Another example to illustrate this method of balancing chemical equations is the oxidation of wastewater (sewage) sludge. Here, the reactant is not a single type of molecule, but rather a complex mixture of carbohydrates, proteins, fats, and other organic compounds. A practical way of dealing with this problem is to represent the average quantities of carbon, hydrogen, and oxygen in the form of a *compositional formula*²⁸ based on a gross analysis of the wastewater sludge:



We know that the products will be carbon dioxide and water, but the question is how much oxygen will be required to completely oxidize the mixture. The following (unbalanced) chemical equation shows the reactants and products:



The non-integer subscripts greatly complicate trial-and-error stoichiometry, but they pose absolutely no obstacle at all to simultaneous equations. Assigning variables x , y , and z to the unknown molecular quantities:

1	x	=	y	z
$\text{C}_{4.8}\text{H}_{8.4}\text{O}_{2.2}$	O_2	\rightarrow	CO_2	H_2O


Now, we may write three algebraic equations, each representing the stoichiometric balance of one element in the chemical equation. The following table shows each of the three elements with their respective balance equations:

Element	Balance equation
Carbon	$4.8 + 0x = 1y + 0z$
Hydrogen	$8.4 + 0x = 0y + 2z$
Oxygen	$2.2 + 2x = 2y + 1z$

Simplifying each equation by eliminating all zero values and “1” coefficients:

Element	Balance equation
Carbon	$4.8 = y$
Hydrogen	$8.4 = 2z$
Oxygen	$2.2 + 2x = 2y + z$

We may tell from the first and second equations that $y = 4.8$ and $z = 4.2$, which then leads to a solution of $x = 5.8$ once the values for y and z have been inserted into the third equation. The final result is this balanced compositional equation for the oxidation of wastewater sludge:




My own personal experience with the use of simultaneous linear equations as a tool for stoichiometry is that it is much faster (especially when balancing complex reaction equations) than trial-and-error, and relatively easy to set up once the general principles are understood.

²⁸Note that you cannot have a molecule comprised of 4.8 carbon atoms, 8.4 hydrogen atoms, and 2.2 oxygen atoms, since atoms exist in whole numbers only! This compositional formula merely shows us the *relative proportions* of each element in the complex mixture of molecules that make up sewage sludge.

3.8.3 Stoichiometric ratios

Regardless of the technique used to balance the equation for a chemical reaction, the most practical *purpose* of balancing the equation is to be able to relate the reactant and product quantities to each other. For instance, we may wish to know how much oxygen will be required to completely combust with a given quantity of fuel, so that we will be able to engineer a burner system capable of handling the necessary flow rates of fuel and oxygen. Balancing the chemical reaction is just the first part of the solution. Once we have a balanced equation, we need to consider the ratios of the substances to each other.

For example, let us consider the balanced (stoichiometric) chemical equation for the combustion of ethane fuel with pure oxygen:



From the balanced chemical equation we can see that for every 2 molecules of ethane, we will need 7 molecules of oxygen gas to completely combust, producing 6 molecules of water vapor and 4 molecules of carbon dioxide gas. The numerical multipliers preceding each molecule in the balanced equation tell us the *molar ratios* of those substances to each other. For oxygen to ethane the ratio is 7:2, for water to ethane the ratio is 6:2 (or 3:1), for carbon dioxide to water the ratio is 4:6 (2:3), etc. If for some reason we needed to calculate the number of moles of CO_2 produced after burning 80 moles of ethane, we could easily calculate that by multiplying the 80 moles of ethane by the 2:4 (1:2) ethane-to-carbon dioxide ratio to arrive at a figure of 160 moles of CO_2 . If we wished, we could even solve this using the same method of *unity fractions* we commonly apply in unit-conversion problems, writing the carbon dioxide-to-ethane ratio as a fraction of two equivalent quantities:

$$\left(\frac{80 \text{ mol ethane}}{1}\right) \left(\frac{4 \text{ molecules carbon dioxide}}{2 \text{ molecules ethane}}\right) = 160 \text{ mol carbon dioxide}$$

If any substances involved in the reaction happen to be gases at nearly the same pressure and temperature²⁹, the molar ratios defined by the balanced equation will similarly define the *volumetric ratios* for those substances. For example, knowing our ideal oxygen-to-ethane molar ratio is 7:2 tells us that the volumetric flow rate of oxygen to ethane should also be (approximately) 7:2, assuming both the oxygen and ethane are gases flowing through their respective pipes at the same pressure and at the same temperature. Recall that the Ideal Gas Law ($PV = nRT$) is approximately true for *any* gas far from its critical phase-change values. So long as pressure (P) and temperature (T) are the same for both gases, each gas's volume (V) will be directly proportional to its molar quantity (n), since R is a constant. This means any molar ratio ($\frac{n_1}{n_2}$) for two gases under identical pressure and temperature conditions will be equal to the corresponding volumetric ratio ($\frac{V_1}{V_2}$) for those gases.

²⁹These assumptions are critically important to equating volumetric ratios with molar ratios. First, the compared substances must both be *gases*: the volume of one mole of steam is huge compared to the volume of one mole of liquid water. Next, we cannot assume temperatures and pressures will be the same after a reaction as before. This is especially true for our example here, where ethane and oxygen are *burning* to produce water vapor and carbon dioxide: clearly, the products will be at a greater temperature than the reactants!

It is important to understand that these molar ratios are not the same as the *mass* ratios for the reactants and products, simply because the different substances do not all have the same mass per mole.

If we regard each of the multipliers in the balanced equation as a precise molar quantity (i.e. exactly 2 moles of ethane, 7 moles of oxygen, etc.) and calculate the mass of the reactants, we will find this value precisely equal to the total mass of the products because the Law of Mass Conservation holds true for this (and all other) chemical reactions:

$$2\text{C}_2\text{H}_6 = 2[(12)(2) + (1)(6)] = 60 \text{ grams}$$

$$7\text{O}_2 = 7[(16)(2)] = 224 \text{ grams}$$

$$6\text{H}_2\text{O} = 6[(1)(2) + 16] = 108 \text{ grams}$$

$$4\text{CO}_2 = 4[12 + (16)(2)] = 176 \text{ grams}$$

Calculating mass based on 2 moles of ethane, we have a total reactant mass of 284 grams (60 grams ethane plus 224 grams oxygen), and a total product mass of 284 grams as well (108 grams water plus 176 grams carbon dioxide gas). We may write the mass ratios for this chemical reaction as such:

$$(\text{ethane}) : (\text{oxygen}) : (\text{water}) : (\text{carbon dioxide})$$

$$60 : 224 : 108 : 176$$

If for some reason we needed to calculate the mass of one of these substances in relation to the other for this reaction, we could easily do so using the appropriate mass ratios. For example, assume we were installing a pair of mass flowmeters to measure the mass flow rates of ethane and pure oxygen gas flowing into the combustion chamber of some industrial process. Supposing the ethane flowmeter had a calibrated range of 0 to 20 kg/min, what range should the oxygen's mass flowmeter be calibrated to in order to match in perfect stoichiometric ratio (so that when one flowmeter is at the top of its range, the other flowmeter should be also)?

The answer to this question is easy to calculate, knowing the required mass ratio of oxygen to ethane for this chemical reaction:

$$\left(\frac{20 \text{ kg ethane}}{1}\right) \left(\frac{224 \text{ g oxygen}}{60 \text{ g ethane}}\right) = 74.67 \text{ kg oxygen}$$

Therefore, the oxygen mass flowmeter should have a calibrated range of 0 to 74.67 kg/min. Note how the unit of mass used in the initial quantity (20 *kilograms* ethane) does not have to match the mass units used in our unity fraction (grams). We could have just as easily calculated the number of *pounds* per minute of oxygen given pounds per minute of ethane, since the mass ratio (like all ratios) is a unitless quantity³⁰.

³⁰Looking at the unity-fraction problem, we see that "grams" (g) will cancel from top and bottom of the unity fraction, and "ethane" will cancel from the given quantity and from the bottom of the unity fraction. This leaves "kilograms" (kg) from the given quantity and "oxygen" from the top of the unity fraction as the only units remaining after cancellation, giving us the proper units for our answer: *kilograms of oxygen*.

3.9 Energy in chemical reactions

A chemical reaction resulting in a net release of energy is called *exothermic*. Conversely, a chemical reaction requiring a net input of energy to occur is called *endothermic*. The relationship between chemical reactions and energy exchange corresponds to the breaking or making of chemical bonds. Atoms bonded together represent a lower state of total energy than those same atoms existing separately, all other factors being equal. Thus, when separate atoms join together to form a molecule, they go from a high state of energy to a low state of energy, releasing the difference in energy in some form (heat, light, etc.). Conversely, an input of energy is required to break that chemical bond and force the atoms to separate.



Chemical bonds are considered “strong” if a large input of energy is required to break them. “Weak” chemical bonds, by contrast, only require modest inputs of energy to disrupt. Thus, the strength of a bond is inversely proportional to the energy state of the molecule: atoms falling into very low energy states when joining together to form a molecule enjoy a strong bond because a large investment of energy is required to raise those atoms’ energy states high enough to sever their bond.

An example of a strong bond is that which exists between two atoms of hydrogen (H) and one atom of oxygen (O) when forming water (H₂O). When hydrogen and oxygen atoms bond together to form water, they release energy. This, by definition, is an exothermic reaction, but we know it better as *combustion*: hydrogen is flammable in the presence of oxygen. A reversal of this reaction occurs when water is subjected to an electrical current, breaking water molecules up into hydrogen and oxygen gas molecules. This process of forced separation requires a substantial input of energy to accomplish, which by definition makes it an *endothermic* reaction. Specifically, the use of electricity to cause a chemical reaction is called *electrolysis*.

An even stronger bond is that formed between aluminum (Al) and oxygen (O) to make alumina (Al₂O₃), a ceramic powder at room temperature. The energy state of this molecule is so low that the aluminum-oxygen bonds resist dissolution even at extremely high temperatures, explaining the high melting point and relative non-reactivity of this substance.

3.9.1 Heats of reaction and activation energy

The amount of energy exchanged (either absorbed or released) in a chemical reaction is often expressed as a numerical quantity to the right of the equation, labeled ΔH , usually defined at a reference temperature of 298 Kelvin (25 degrees Celsius). A negative ΔH value signifies an exothermic (heat-releasing) reaction, while a positive ΔH value signifies an endothermic (heat-absorbing) reaction. The combustion of hydrogen and oxygen to form liquid water is an example of the former, and the electrolysis of water to yield hydrogen and oxygen gas is an example of the latter:




This energy value, commonly referred to as the *heat of reaction* or *enthalpy of reaction*, is expressed *per mole* of the reactants and products shown. The “−1” exponent applied to “mole” is simply a fancy way of saying “per mole”³¹, as an alternative to using a fraction bar.

While the mathematical sign of ΔH may seem confusing at first (“Why is it *negative* when energy is *released*?”), it makes sense from the perspective of energy states before and after the reaction. In an exothermic (heat-releasing) reaction, the products are left at a lower state of energy than the reactants began with, and a negative value of ΔH signifies this. The sign of ΔH , then, is an expression of the change in energy state from reactants (input) to products (output), *not* an expression of the energy liberated from the reaction. Even though the term ΔH is called the “heat of reaction” it really refers to the change in potential energy of the *matter* as a consequence of the chemical reaction.


³¹This notation is quite common in scientific and engineering literature, as a way to avoid having to typeset fractions in a text document. Instead of writing $\frac{\text{kJ}}{\text{mol}}$ which requires a fraction bar, we may write kJ mol^{-1} which is mathematically equivalent. Another common example of this notation is to express frequency in the unit of s^{-1} (per second) rather than the unit of Hertz (Hz). Perhaps the most compelling reason to use negative exponents in unit expressions, though, is sociological: scientific studies have shown the regular use of this unit notation makes you appear 37.5% smarter than you actually are. Questioning statistical results of scientific studies, on the other hand, reduces your apparent intelligence by over 63%! Now, aren't you glad you took the time to read this footnote?

The fact that hydrogen and oxygen as separate gases possess potential energy does not mean they are guaranteed to spontaneously combust when brought together. By analogy, just because rocks sitting on a hillside possess potential energy (by virtue of being elevated above the hill's base) does not mean all rocks in the world spontaneously roll downhill. Some rocks need a push to get started because they are caught on a ledge or resting in a depression on the hillside. Likewise, many exothermic reactions require an initial investment of energy before they can proceed. In the case of hydrogen and oxygen, what is generally needed is a spark to initiate the reaction. This initial requirement of input energy is called the *activation energy* of the reaction.

Activation energy may be shown in graphical form. For an exothermic reaction, it appears as a "hill" that must be climbed before the total energy can fall to a lower (than original) level:



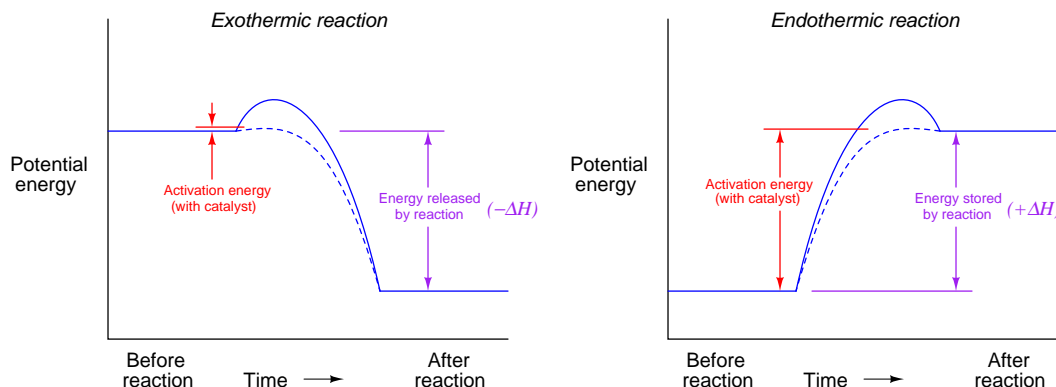
For an endothermic reaction, activation energy is much greater, a part of which never returns but is stored in the reaction products as potential energy:



A *catalyst* is a substance that works³² to minimize activation energy in a chemical reaction without being altered by the reaction itself. Catalysts are popularly used in industry to accelerate both exothermic and endothermic reactions, reducing the gross amount of energy that must be initially input to a process to make a reaction occur. A common example of a catalyst is the *catalytic converter* installed in the exhaust pipe of an automobile engine, helping to oxidize unburnt fuel molecules and certain combustion products such as carbon monoxide (CO) to compounds which are not as polluting. Without a catalytic converter, the exhaust gas temperature is not hot enough to overcome the activation energy of these reactions, and so they will not occur (at least not at the rate necessary to make a significant difference). The presence of the catalyst allows the reactions to progress quickly at typical engine exhaust temperatures.

³²Just how catalysts perform this trick is a subject of continuing research. Catalysts used in industrial process industries are usually selected based on the results of empirical tests rather than by theory, since a general theoretical understanding of catalysis is lacking at this present time. Indeed, the specific selection of catalysts for high-value chemical processes is often a patented feature of those processes, reflecting the investment of time, finances, and effort finding a suitable catalyst for optimizing each chemical reaction.

The effect of a catalyst on activation energy may be shown by the following graphs, the dashed-line curve showing the energy progression with a catalyst and the solid-line curve showing the reaction progressing without the benefit of a catalyst:



It should be noted that the presence of a catalyst has absolutely no effect on the *net* energy loss or gain resulting from a chemical reaction. That is to say, the heat of reaction (ΔH) stands independent of catalytic assistance: with or without a catalyst, the difference in potential energy before and after a reaction will be the same³³. The only difference a catalyst makes to a chemical reaction is how much energy must be *initially invested* to spark the reaction. To use the example of hydrogen and oxygen gas once again, the presence of a catalyst does not cause the combustion of hydrogen and oxygen to release more energy. All the catalyst does is make it easier for the combustion to begin.

³³If this were not true, one could construct an over-unity ("perpetual motion") machine by initiating an endothermic reaction and then reversing that reaction (exothermic) using a catalyst in either or both portions of the cycle to reap a net energy release from the system. So trustworthy is the Law of Energy Conservation that we may safely invoke the impossibility of over-unity energy production as a *disproof* of any given hypothesis permitting it. In other words, if any hypothesis allows for an over-unity process (i.e. violates the Law of Energy Conservation), we may reject that hypothesis with confidence! This form of disproof goes by the name *reductio ad absurdum* (Latin: "reducing to an absurdity").

3.9.2 Heats of formation and Hess's Law

As we have seen, the formation of new chemical bonds between atoms is an energy-releasing process (i.e. exothermic), while the dissolution of chemical bonds is an energy-absorbing process (i.e. endothermic). Given the fact that the Law of Energy Conservation is universal, it stands to reason we ought to be able to mathematically balance the potential energy held by reactants, the potential energy held by the products, and the amount of energy either released or absorbed by the reaction.

Let us begin our exploration of this concept with the formation of carbon dioxide (CO_2) from the combustion of elemental carbon (C) and oxygen molecules (O_2):



As we can see, this reaction is exothermic³⁴: the products have a lower energy than the reactants, losing 393.5 kilojoules of energy for every mole of carbon dioxide formed by this reaction.

If we are to account for all the energy entering and exiting a chemical reaction, we must have some means of quantifying the amount of energy stored within both the reactants and the products as well as the amount of energy released or absorbed by the reaction itself. Quantifying the amount of chemical potential possessed by atoms and molecules is difficult if not impossible to do in any absolute sense, and so the common practice is to arbitrarily assign an energy value of *zero* to chemical elements in their normal states at standard temperature and pressure (293.15 Kelvin and 1 atmosphere, abbreviated “STP”). This point of reference will be the norm for any subsequent determinations of chemical potential energy. The *standard heat of formation* or *standard enthalpy of formation* (ΔH_f° , or sometimes $\Delta_f H^\circ$) for any substance is thus defined as the amount of energy gained or lost when one mole of that substance is formed from its constituent elements at STP. A superscripted “o” symbol represents conditions of standard temperature and pressure.

We know that the *phase* of a substance (i.e. solid, liquid, gas) affects how much energy it contains, and therefore in order to accurately account for all energy we must represent the phase of each substance when we specify heats of formation. Since the natural state of carbon is solid (s) at STP while the natural state of oxygen is gas (g) at STP, we will represent those states as letters within parentheses when defining their heats of formation:

$$\Delta H_f^\circ (\text{C}, \text{s}) = 0 \text{ kJ mol}^{-1} \quad \text{Heat of formation for solid carbon at STP}$$

$$\Delta H_f^\circ (\text{O}_2, \text{g}) = 0 \text{ kJ mol}^{-1} \quad \text{Heat of formation for gaseous oxygen at STP}$$

It takes no gain or loss of energy at all to form solid carbon (C) or gaseous oxygen molecules (O_2) at STP because those elements are already in those forms at STP. The only way we will ever have a non-zero ΔH_f° value is if the substance in question is a compound (i.e. comprised of multiple elements joined by chemical bonds) or if the substance in question is an element in some unusual energy state (e.g. ionization).

³⁴The heat-releasing nature of this reaction should be evident by inspection of the reaction equation alone, without reference to the ΔH figure. Note how the carbon (C) and oxygen (O_2) begin as separate entities, but end up as one combined molecule (CO_2). Since we know the formation of new chemical bonds is exothermic, this reaction between carbon and oxygen to produce carbon dioxide must therefore be exothermic.

Since we already know the combustion of one atom of carbon with one molecule of oxygen liberates 393.5 kilojoules of heat energy, we may conclude the heat of formation for carbon dioxide gas (cooled down to the standard temperature of 293.15 Kelvin) must be $-393.5 \text{ kJ mol}^{-1}$, since this is precisely how much energy is liberated when carbon dioxide is formed from its constituent elements. Representing all these figures in a table helps us make sense of it all:

Reactant	Reactant	Reaction	Product
C(s)	O ₂ (g)	→	CO ₂ (g)
$\Delta H_f^\circ = 0 \text{ kJ mol}^{-1}$	$\Delta H_f^\circ = 0 \text{ kJ mol}^{-1}$	$\Delta H^\circ = -393.5 \text{ kJ mol}^{-1}$	$\Delta H_f^\circ = -393.5 \text{ kJ mol}^{-1}$

Stated verbally, the combined heats of formation for all reactants plus the heat of reaction yields the combined heats of formation for all products. Put into simpler terms, the energy contained by the reactants plus the change in energy wrought by the reaction gives us the energy left³⁵ inside the products. The mathematical formulation of this principle is as follows:

$$\Sigma [\Delta H_f^\circ (\text{Reactants})] + \Delta H^\circ (\text{Reaction}) = \Sigma [\Delta H_f^\circ (\text{Products})]$$


The practical application of this is that we may calculate³⁶ the exact amount of heat liberated or absorbed by *any* chemical reaction, if only we know in advance the heats of formation for all the reactants and products. Fortunately for our reference, chemists have tabulated standard heats of formation for a great many substances.

Hess's Law states that this accounting of energy is true regardless of the reaction path. For example, if the combustion of carbon with oxygen proceeds in one step ($\text{C} + \text{O}_2 \rightarrow \text{CO}_2$), the overall heat of reaction will be precisely the same as for any other series of steps resulting in the same product(s) from the same reactant(s), for example the partial combustion of carbon to form carbon monoxide ($\text{C} + \text{O} \rightarrow \text{CO}$) followed by the subsequent combustion of carbon monoxide ($\text{CO} + \text{O} \rightarrow \text{CO}_2$). Just as we saw with catalytically-aided chemical reactions, the total heat of reaction is strictly a function of the reactants and the products, not of any process or path by which the reaction may proceed. The Law of Energy Conservation is indeed iron-clad!

³⁵At first it may seem non-sensical for the carbon dioxide product of this reaction to have a *negative* energy, until you realize the zero values given to both the carbon and oxygen reactants are entirely arbitrary. Viewed in this light, the negative heat of formation for CO₂ is nothing more than a *relative* expression of chemical potential energy in reference to the elements from which CO₂ originated. Therefore, a negative ΔH_f° value for any molecule simply tells us that molecule has less energy (i.e. is more stable) than its constituent elements.

³⁶We may also readily tell whether any given reaction will be exothermic or endothermic, based on the mathematical sign of this ΔH value.

Let us investigate a practical application where we employ heats of formation to calculate the heat of a chemical reaction. In this case, we will consider the combustion of propane fuel gas (C_3H_8) in the presence of pure oxygen gas (O_2), producing liquid water (H_2O) and gaseous carbon dioxide (CO_2) as products:



To begin, we must identify the standard heats of formation for each of these substances at STP from a suitable reference³⁷:

- Propane gas $\Delta H_f^\circ = -103.8 \text{ kJ mol}^{-1}$
- Oxygen gas $\Delta H_f^\circ = 0 \text{ kJ mol}^{-1}$
- Water $\Delta H_f^\circ = -285.8 \text{ kJ mol}^{-1}$
- Carbon dioxide gas $\Delta H_f^\circ = -393.5 \text{ kJ mol}^{-1}$

Setting these quantities into a table for ease of organization (all heats of formation given in units of kilojoules per mole):

Reactant	Reactant	Reaction	Product	Product
$\text{C}_3\text{H}_8(\text{g})$	$5\text{O}_2(\text{g})$	\rightarrow	$4\text{H}_2\text{O}(\text{l})$	$3\text{CO}_2(\text{g})$
(1)(-103.8)	(5)(0)	ΔH°	(4)(-285.8)	(3)(-393.5)

Solving for the unknown heat of reaction (ΔH°):

$$\Sigma [\Delta H_f^\circ (\text{Reactants})] + \Delta H^\circ (\text{Reaction}) = \Sigma [\Delta H_f^\circ (\text{Products})]$$

$$[(1)(-103.8) + (5)(0)] + \Delta H^\circ = [(4)(-285.8) + (3)(-393.5)]$$

$$[-103.8] + \Delta H^\circ = [-1143.2 + -1180.5]$$

$$[-103.8] + \Delta H^\circ = [-2323.7]$$

$$\Delta H^\circ = -2323.7 + 103.8$$

$$\Delta H^\circ = -2219.9 \text{ kJ per mole of propane fuel}$$

The large, negative value of ΔH° tells us the reaction of propane with oxygen will be highly exothermic.

³⁷Of course, it is not necessary to look up ΔH_f° for oxygen gas, as that is an element in its natural state at STP and therefore its standard heat of formation is defined to be zero. The heat of formation for carbon dioxide gas may be found from the preceding example, while the heat of formation for water may be found in the “Heats of Reaction and Activation Energy” subsection of this book. The only substance in this list of which the heat of formation is not defined as zero or given in this book is propane. Note that many thermochemical reference books will give heats of formation in units of *kilocalories* per mole rather than kilojoules per mole. The conversion factor between these is 1 calorie = 4.184 joules.

3.10 Periodic table of the ions

Periodic Table of the Ions

Ionization state

Symbol: **K**

Name: **Potassium**

Electron configuration: **4s¹**

Atomic number: **19**

Atomic mass: **39.0983**
(averaged according to occurrence on earth)

Metals																		Metalloids						Nonmetals																																																																		
H 1 Hydrogen 1.00794 1s ¹	He 2 Helium 4.00260 1s ²	Li 3 Lithium 6.941 2s ²	Be 4 Beryllium 9.012182 2s ²	B 5 Boron 10.81 2p ¹	C 6 Carbon 12.011 2p ²	N 7 Nitrogen 14.0067 2p ³	O 8 Oxygen 15.9994 2p ⁴	F 9 Fluorine 18.9984 2p ⁵	Ne 10 Neon 20.179 2p ⁶	Na 11 Sodium 22.989768 3s ¹	Mg 12 Magnesium 24.3050 3s ²	Al 13 Aluminum 26.9815 3p ¹	Si 14 Silicon 28.0855 3p ²	P 15 Phosphorus 30.9738 3p ³	S 16 Sulfur 32.06 3p ⁴	Cl 17 Chlorine 35.453 3p ⁵	Ar 18 Argon 39.948 3p ⁶	K 19 Potassium 39.0983 4s ¹	Ca 20 Calcium 40.078 4s ²	Sc 21 Scandium 44.955910 3d ¹ 4s ²	Ti 22 Titanium 47.88 3d ² 4s ²	V 23 Vanadium 50.9415 3d ³ 4s ²	Cr 24 Chromium 51.9961 3d ⁵ 4s ¹	Mn 25 Manganese 54.93805 3d ⁵ 4s ²	Fe 26 Iron 55.847 3d ⁶ 4s ²	Co 27 Cobalt 58.93320 3d ⁷ 4s ²	Ni 28 Nickel 58.69 3d ⁸ 4s ²	Cu 29 Copper 63.546 3d ¹⁰ 4s ¹	Zn 30 Zinc 65.39 3d ¹⁰ 4s ²	Ga 31 Gallium 69.723 4p ¹	Ge 32 Germanium 72.61 4p ²	As 33 Arsenic 74.92159 4p ³	Se 34 Selenium 78.96 4p ⁴	Br 35 Bromine 79.904 4p ⁵	Kr 36 Krypton 83.80 4p ⁶	Rb 37 Rubidium 85.4678 5s ¹	Sr 38 Strontium 87.62 5s ²	Y 39 Yttrium 88.90585 4d ¹ 5s ²	Zr 40 Zirconium 91.224 4d ² 5s ²	Nb 41 Niobium 92.90638 4d ⁴ 5s ¹	Mo 42 Molybdenum 95.94 4d ⁵ 5s ¹	Tc 43 Technetium (98) 4d ⁵ 5s ²	Ru 44 Ruthenium 101.07 4d ⁷ 5s ¹	Rh 45 Rhodium 102.90550 4d ⁸ 5s ¹	Pd 46 Palladium 106.42 4d ¹⁰ 5s ⁰	Ag 47 Silver 107.8682 4d ¹⁰ 5s ¹	Cd 48 Cadmium 112.411 4d ¹⁰ 5s ²	In 49 Indium 114.82 5p ¹	Sn 50 Tin 118.710 5p ²	Sb 51 Antimony 121.75 5p ³	Te 52 Tellurium 127.60 5p ⁴	I 53 Iodine 126.905 5p ⁵	Xe 54 Xenon 131.30 5p ⁶	Cs 55 Cesium 132.90543 6s ¹	Ba 56 Barium 137.327 6s ²	Lanthanide series						Hf 72 Hafnium 178.49 5d ² 6s ²	Ta 73 Tantalum 180.9479 5d ³ 6s ²	W 74 Tungsten 183.85 5d ⁴ 6s ²	Re 75 Rhenium 186.207 5d ⁵ 6s ²	Os 76 Osmium 190.2 5d ⁶ 6s ²	Ir 77 Iridium 192.22 5d ⁷ 6s ²	Pt 78 Platinum 195.08 5d ⁹ 6s ¹	Au 79 Gold 196.96654 5d ¹⁰ 6s ¹	Hg 80 Mercury 200.59 5d ¹⁰ 6s ²	Tl 81 Thallium 204.3833 6p ¹	Pb 82 Lead 207.2 6p ²	Bi 83 Bismuth 208.98037 6p ³	Po 84 Polonium (209) 6p ⁴	At 85 Astatine (210) 6p ⁵	Rn 86 Radon (222) 6p ⁶	Fr 87 Francium (223) 7s ¹	Ra 88 Radium (226) 7s ²	Actinide series						Unq 104 Unnilquadium (261) 6d ⁴ 7s ²	Unp 105 Unnilpentium (262) 6d ⁴ 7s ²	Unh 106 Unnilhexium (263) 6d ⁴ 7s ²	Uns 107 Unnilseptium (262) 6d ⁴ 7s ²	108	109

Lanthanide series											
La 57 Lanthanum 138.9055 5d ¹ 6s ²	Ce 58 Cerium 140.115 4f ¹ 5d ¹ 6s ²	Pr 59 Praseodymium 140.90765 4f ³ 6s ²	Nd 60 Neodymium 144.24 4f ⁴ 6s ²	Pm 61 Promethium (145) 4f ⁵ 6s ²	Sm 62 Samarium 150.36 4f ⁶ 6s ²	Eu 63 Europium 151.965 4f ⁷ 6s ²	Gd 64 Gadolinium 157.25 4f ⁷ 5d ¹ 6s ²	Tb 65 Terbium 158.92534 4f ⁹ 6s ²	Dy 66 Dysprosium 162.50 4f ¹⁰ 6s ²	Ho 67 Holmium 164.93032 4f ¹¹ 6s ²	Er 68 Erbium 167.26 4f ¹² 6s ²

Actinide series											
Ac 89 Actinium (227) 6d ¹ 7s ²	Th 90 Thorium 232.0381 6d ² 7s ²	Pa 91 Protactinium 231.03688 5f ² 6d ¹ 7s ²	U 92 Uranium 238.02891 5f ³ 6d ¹ 7s ²	Np 93 Neptunium (237) 5f ⁴ 6d ¹ 7s ²	Pu 94 Plutonium (244) 5f ⁶ 6d ¹ 7s ²	Am 95 Americium (243) 5f ⁷ 6d ¹ 7s ²	Cm 96 Curium (247) 5f ⁷ 6d ¹ 7s ²	Bk 97 Berkelium (247) 5f ⁹ 6d ¹ 7s ²	Cf 98 Californium (251) 5f ¹⁰ 6d ¹ 7s ²	Es 99 Einsteinium (252) 5f ¹¹ 6d ¹ 7s ²	Fm 100 Fermium (257) 5f ¹² 6d ¹ 7s ²

3.11 Ions in liquid solutions

Many liquid substances undergo a process whereby their constituent molecules split into positively and negatively charged ion pairs, the positively-charge ion called a *cation* and the negatively-charged ion called an *anion*³⁸. Liquid *ionic* compounds³⁹ split into ions completely or nearly completely, while only a small percentage of the molecules in a liquid *covalent* compound⁴⁰ split into ions. The process of neutral molecules separating into ion pairs is called *dissociation* when it happens to ionic compounds, and *ionization* when it happens to covalent compounds.

Molten salt (NaCl) is an example of the former, while pure water (H₂O) is an example of the latter. In liquid salt, practically every NaCl molecule splits up into an Na⁺ and Cl⁻ ion pair, whereas with liquid water only a very small percentage of molecules split up into positively and negatively charged ions – most remain as whole H₂O molecules. All the ions present in molten salt serve as electrical charge carriers, making molten salt a very good conductor of electricity. The scarcity of ions in a sample of pure water explains why it is often considered an insulator. In fact, the electrical conductivity of a liquid substance is the definitive test of whether it is an ionic or a covalent (“molecular”) substance.

The few water molecules that do ionize split into positive hydrogen ions⁴¹ (H⁺) and negative hydroxyl ions (OH⁻). At room temperature, the concentration of hydrogen and hydroxyl ions in a sample of pure water is quite small: a molarity of 10⁻⁷ M (moles of hydrogen ions per liter of solution) each.

Given the fact that pure water has a mass of 1 kilogram (1000 grams) per liter, and one mole of pure water has a mass of 18 grams, we must conclude that there are approximately 55.56 moles of water molecules in one liter (55.56 M). If only 10⁻⁷ moles of those molecules ionize at room temperature, that represents an extremely small percentage of the total:

$$\frac{10^{-7} \text{ mol hydrogen ions}}{55.56 \text{ mol solution}} = 0.000000018 = 0.0000018\% = 0.0018 \text{ ppm (parts per million)}$$

It is not difficult to see why pure water is such a poor conductor of electricity. With so few ions available to act as charge carriers, pure water is practically an insulator. The vast majority of water molecules remain un-ionized and therefore cannot transport electric charges from one point to another.

The molarity of both hydrogen and hydroxyl ions in a pure water sample increases with increasing temperature. For example, at 60 °C, the molarity of hydrogen and hydroxyl ions increases to 3.1 ×

³⁸These names have their origin in the terms used to classify positive and negative electrodes immersed in a liquid solution. The positive electrode is called the “anode” while the negative electrode is called the “cathode.” An *anion* is an ion attracted to the anode. A *cation* is an ion attracted to the cathode. Since opposite electrical charges tend to attract, this means “anions” are negatively charged and “cations” are positively charged.

³⁹Ionic compounds are formed when oppositely charged atomic ions bind together by mutual attraction. The distinguishing characteristic of an ionic compound is that it is a conductor of electricity in its pure, liquid state. That is, it readily separates into anions and cations all by itself. Even in its solid form, an ionic compound is already ionized, with its constituent atoms held together by an imbalance of electric charge. Being in a liquid state simply gives those atoms the physical mobility needed to dissociate.

⁴⁰Covalent compounds are formed when electrically neutral atoms bind together by the mutual sharing of valence electrons. Such compounds are not good conductors of electricity in their pure, liquid states.

⁴¹Actually, the more common form of positive ion in water is *hydronium*: H₃O⁺, but we often simply refer to the positive half of an ionized water molecule as hydrogen (H⁺).

$10^{-7} M$, which is still only 0.0056 parts per million, but definitely larger than the concentration at room temperature (25 °C).

The electrical conductivity of water may be greatly enhanced by dissolving an ionic compound in it, such as table salt. When dissolved, the table salt molecules (NaCl) immediately dissociate into sodium cations (Na^+) and chlorine anions (Cl^-), becoming available as charge carriers for an electric current. In industry, we may exploit this relationship between electrical conductivity and ionic dissociation to detect the presence of ionic compounds in otherwise pure water.

3.12 pH

Hydrogen ion activity in aqueous (water-solvent) solutions is a very important parameter for a wide variety of industrial processes. A number of reactions important to chemical processing are inhibited or significantly slowed if the hydrogen ion activity of a solution lies outside a narrow range. Some additives used in water treatment processes (e.g. flocculants) will fail to function efficiently if the hydrogen ion activity in the water is not kept within a certain range. Alcohol and other fermentation processes strongly depend on tight control of hydrogen ion activity, as an incorrect level of ion activity will not only slow production but may also spoil the product. The concentration of active hydrogen ions⁴² in a solution is always measured on a logarithmic scale, and referred to as *pH*.

pH is mathematically defined as the negative common logarithm of active hydrogen ion concentration in a solution⁴³. Hydrogen ion concentration is expressed as a molarity (number of moles of ions per liter of total liquid solution volume), with “pH” being the unit of measurement for the logarithmic result:

$$\text{pH} = -\log[\text{H}^+]$$

For example, an aqueous solution with an active hydrogen concentration of 0.00044 *M* has a pH value of 3.36 pH.

Water is a covalent compound, and so there is little ionization of water molecules in liquid form. Most of the molecules in a sample of pure water remain as whole molecules (H_2O) while a very small percentage ionize into positive hydrogen ions (H^+) and negative hydroxyl ions (OH^-). The mathematical product of hydrogen and hydroxyl ion molarity in water is known as the *ionization constant* (K_w), and its value varies with temperature:

$$K_w = [\text{H}^+] \times [\text{OH}^-]$$

At 25 degrees Celsius (room temperature), the value of K_w is very nearly equal to 1.0×10^{-14} . Since each one of the water molecules that does ionize in this absolutely pure water sample separates into exactly one hydrogen ion (H^+) and one hydroxyl ion (OH^-), the molarities of hydrogen and hydroxyl ions must be equal to each other. The equality between hydrogen and hydroxyl ions in a pure water sample means that pure water is *neutral*, and that the molarity of hydrogen ions is equal to the square root of K_w :

$$[\text{H}^+] = \sqrt{K_w} = \sqrt{1.0 \times 10^{-14}} = 1.0 \times 10^{-7} \text{ M}$$

⁴²Free hydrogen ions (H^+) are rare in a liquid solution, and are more often found attached to whole water molecules to form a positive ion called *hydronium* (H_3O^+). However, process control professionals usually refer to these positive ions simply as “hydrogen” even though the truth is a bit more complicated.

⁴³The letter “p” refers to “potential,” in reference to the logarithmic nature of the measurement. Other logarithmic measurements of concentration exist for molecular species, including pO_2 and pCO_2 (concentration of oxygen and carbon dioxide molecules in a liquid solution, respectively).

Since we know pH is defined as the negative logarithm of hydrogen ion activity, and we can be assured all hydrogen ions present in a pure water sample will be “active” since there are no other positive ions to interfere with them, the pH value for water at 25 degrees Celsius is:

$$\text{pH of pure water at } 25\text{ }^\circ\text{C} = -\log(1.0 \times 10^{-7} M) = 7.0 \text{ pH}$$

As the temperature of a pure water sample changes, the ionization constant changes as well. Increasing temperature causes more of the water molecules to ionize into H^+ and OH^- ions, resulting in a larger K_w value and a lower pH value. The following table shows K_w and pH values for pure water at different temperatures:

Temperature	K_w	pH
0 °C	1.139×10^{-15}	7.47 pH
5 °C	1.846×10^{-15}	7.37 pH
10 °C	2.920×10^{-15}	7.27 pH
15 °C	4.505×10^{-15}	7.17 pH
20 °C	6.809×10^{-15}	7.08 pH
25 °C	1.008×10^{-14}	6.998 pH
30 °C	1.469×10^{-14}	6.92 pH
35 °C	2.089×10^{-14}	6.84 pH
40 °C	2.919×10^{-14}	6.77 pH
45 °C	4.018×10^{-14}	6.70 pH
50 °C	5.474×10^{-14}	6.63 pH
55 °C	7.296×10^{-14}	6.57 pH
60 °C	9.614×10^{-14}	6.51 pH

This means that while any pure water sample is *neutral* (an equal number of positive hydrogen ions and negative hydroxyl ions) at any temperature, the pH value of pure water actually changes with temperature, and is only equal to 7.0 pH⁴⁴ at one particular (“standard”) temperature: 25 °C. Based on the K_w values shown in the table, pure water will be 6.51 pH at 60 °C and 7.47 pH at freezing.

⁴⁴Often, students assume that the 7 pH value of water is an arbitrary assignment, using water as a universal standard just like we use water as the standard for the Celsius temperature scale, viscosity units, specific gravity, etc. However, this is not the case here. Pure water at room temperature just happens to have an hydrogen ion molarity equivalent to a (nearly) round-number value of 7 pH.

If we add an electrolyte to a sample of pure water, molecules of that electrolyte will separate into positive and negative ions⁴⁵. If the positive ion of the electrolyte happens to be a hydrogen ion (H^+), we call that electrolyte an *acid*. If the negative ion of the electrolyte happens to be a hydroxyl ion (OH^-), we call that electrolyte a *caustic*, or *alkaline*, or *base*. Some common acidic and alkaline substances are listed here, showing their respective positive and negative ions in solution:

Sulfuric acid is an *acid* (produces H^+ in solution)



Nitric acid is an *acid* (produces H^+ in solution)



Hydrocyanic acid is an *acid* (produces H^+ in solution)



Hydrofluoric acid is an *acid* (produces H^+ in solution)



Lithium hydroxide is a *caustic* (produces OH^- in solution)



Potassium hydroxide is a *caustic* (produces OH^- in solution)



Sodium hydroxide is a *caustic* (produces OH^- in solution)



Calcium hydroxide is a *caustic* (produces OH^- in solution)



When an acid substance is added to water, some⁴⁶ of the acid molecules dissociate into positive hydrogen ions (H^+) and negative ions (the type of negative ions depending on what type of acid it is). This increases the molarity of hydrogen ions (the number of moles of H^+ ions per liter of solution), therefore driving the pH value of the solution down to a smaller number. For example, a sample of acid added to a sample of neutral water at room temperature (7 pH) will drive the pH value down below 7 due to the increasing molarity of hydrogen ions in the solution. The addition of hydrogen ions to the solution also decreases the molarity of hydroxyl ions (the number of moles of OH^- ions per liter of solution) because some of the water's OH^- ions combine with the acid's H^+ ions to form deionized water molecules (H_2O).

If an alkaline substance (otherwise known as a *caustic*, or a *base*) is added to water, some⁴⁷ of the alkaline molecules dissociate into negative hydroxyl ions (OH^-) and positive ions (the type of positive ions depending on what type of alkaline it is). This increases the molarity of OH^- ions in

⁴⁵If the electrolyte is considered *strong*, all or nearly all of its molecules will dissociate into ions. A *weak* electrolyte is one where only a mere portion of its molecules dissociate into ions.


⁴⁶For “strong” acids, all or nearly all molecules dissociate into ions. For “weak” acids, just a portion of the molecules dissociate.

⁴⁷For “strong” bases, all or nearly all molecules dissociate into ions. For “weak” bases, just a portion of the molecules dissociate.

the solution, as well as decreases the molarity of hydrogen ions (again, because some of the caustic's OH^- ions combine with the water's H^+ ions to form deionized water molecules, H_2O). This decrease in hydrogen ion molarity will raise the pH value of the solution. For example, if we were to add a sample of caustic to a sample of neutral water at room temperature (7 pH), the pH of the solution would increase with the decreasing hydrogen ion molarity.

The result of this complementary effect (increasing one type of water ion, decreasing the other) keeps the overall ionization constant relatively constant, at least for dilute solutions. In other words, the addition of an acid or a caustic to water may change $[\text{H}^+]$, but it has little effect on K_w .

A simple way to envision this effect is to think of a laboratory balance scale, balancing the number of hydrogen ions in a solution against the number of hydroxyl ions in the same solution:



When the solution is pure water, this imaginary scale is balanced (neutral), with $[\text{H}^+] = [\text{OH}^-]$. Adding an acid to the solution tips the scale to the left (lower pH value), while adding a caustic to the solution tips the scale to the right (higher pH value)⁴⁸.

⁴⁸It should be noted that the solution never becomes *electrically* imbalanced with the addition of an acid or caustic. It is merely the balance of hydrogen to hydroxyl ions we are referring to here. The net electrical charge for the solution should still be zero after the addition of an acid or caustic, because while the balance of hydrogen to hydroxyl ions does change, that electrical charge imbalance is made up by the other ions resulting from the addition of the electrolyte (anions for acids, cations for caustics). The end result is still one negative ion for every positive ion (equal and opposite charge numbers) in the solution no matter what substance(s) we dissolve into it.

If an electrolyte has no effect on either the hydrogen and hydroxyl ion activity of an aqueous solution, we call it a *salt*. The following is a list of some common salts, showing their respective ions in solution:


Potassium chloride is a *salt* (produces neither H^+ nor OH^- nor O^{2-} in solution)



Sodium chloride is a *salt* (produces neither H^+ nor OH^- nor O^{2-} in solution)



Zinc sulfate is a *salt* (produces neither H^+ nor OH^- nor O^{2-} in solution)



The addition of a salt to an aqueous solution should have no effect on pH, because the ions created neither add to nor take away from the hydrogen ion activity⁴⁹.

Acids and caustics tend to neutralize one another, the hydrogen ions liberated by the acid combining (and canceling) with the hydroxyl ions liberated by the caustic. This process is called *pH neutralization*, and it is used extensively to adjust the pH value of solutions. If a solution is too acidic, just add caustic to raise its pH value. If a solution is too alkaline, just add acid to lower its pH value.

The result of a perfectly balanced mix of acid and caustic is deionized water (H_2O) and a salt formed by the combining of the acid's and caustic's *other* ions. For instance, when hydrochloric acid (HCl) and potassium hydroxide (KOH) neutralize one another, the result is water (H_2O) and potassium chloride (KCl), a salt. This production of salt is a necessary side-effect of pH neutralization, which may require addressing in later stages of solution processing. Such neutralizations are exothermic, owing to the decreased energy states of the hydrogen and hydroxyl ions after combination. Mixing of pure acids and caustics together without the presence of substantial quantities of water (as a solvent) is often violently exothermic, presenting a significant safety hazard to anyone near the reaction.

Both acidic and caustic solutions pose safety hazards to human and animal life. Concentrated acids will cause burns to living tissue, while concentrated caustics chemically reduce fat within tissue to *soap*. These hazards are not just related to external skin contact, but also to internal contact in the form of ingestion or inhalation.

For more information on pH and its measurement, refer to section 23.2 beginning on page 1761 discussing different technologies for measuring the concentration of hydrogen ions in liquid solutions.

⁴⁹Exceptions do exist for strong concentrations, where hydrogen ions may be present in solution yet unable to react because of being "crowded out" by other ions in the solution.

References

Chase, Malcolm W. Jr., *NIST-JANAF Thermochemical Tables*, Fourth Edition, Part I, Al-Co, Journal of Physical and Chemical Reference Data, Monograph No. 9, American Institute of Physics, American Chemical Society, 1998.

Dolmalski, Eugene S., *Selected Values of Heats of Combustion and Heats of Formation of Organic Compounds Containing the Elements C, H, N, O, P, and S*, Chemical Thermodynamics Data Center, National Bureau of Standards, Washington, D.C., 1972.

“Fundamental Physical Constants – Extensive Listing”, from <http://physics.nist.gov/constants>, National Institute of Standards and Technology (NIST), 2006.

“Gas Detection – the professional guide”, FLIR Systems AB, 2009.

Geddes, L.A. and Baker, L.E., *Principles of Applied Biomedical Instrumentation*, John Wiley & Sons, Inc., New York, NY, 1968.

Giancoli, Douglas C., *Physics for Scientists & Engineers*, Third Edition, Prentice Hall, Upper Saddle River, NJ, 2000.

Haug, Roger Tim, *The Practical Handbook of Compost Engineering*, CRC Press, LLC, Boca Raton, FL, 1993.

Mills, Ian; Cvitaš, Tomislav; Homann, Klaus; Kallay, Nikola; Kuchitsu, Kozo, *Quantities, Units and Symbols in Physical Chemistry* (the “Green Book”), Second Edition, International Union of Pure and Applied Chemistry (IUPAC), Blackwell Science Ltd., Oxford, England, 1993.

“NIOSH Pocket Guide to Chemical Hazards”, DHHS (NIOSH) publication # 2005-149, Department of Health and Human Services (DHHS), Centers for Disease Control and Prevention (CDC), National Institute for Occupational Safety and Health (NIOSH), Cincinnati, OH, September 2005.

Pauling, Linus, *General Chemistry*, Dover Publications, Inc., Mineola, NY, 1988.

Rosman, K.J.R. and Taylor, P.D.P., *Isotopic Compositions of the Elements 1997*, International Union of Pure and Applied Chemistry (IUPAC), 1997.

Scerri, Eric R., “How Good Is the Quantum Mechanical Explanation of the Periodic System?”, *Journal of Chemical Education*, Volume 75, Number 11, pages 1384-1385, 1998.

Theory and Practice of pH Measurement, PN 44-6033, Rosemount Analytical, 1999.

Weast, Robert C.; Astel, Melvin J.; and Beyer, William H., *CRC Handbook of Chemistry and Physics*, 64th Edition, CRC Press, Inc., Boca Raton, FL, 1984.


Whitten, Kenneth W.; Gailey, Kenneth D.; and Davis, Raymond E., *General Chemistry*, Third Edition, Saunders College Publishing, Philadelphia, PA, 1988.

Chapter 4

DC electricity

4.1 Electrical voltage

Voltage is the amount of *specific potential energy* available between two points in an electric circuit. Potential energy is energy that is potentially available to do work. Looking at this from a classical physics perspective, potential energy is what we accumulate when we lift a weight above ground level, or when we compress a spring:



In either case, potential energy is calculated by the work done in exerting a force over a parallel distance. In the case of the weight, potential energy (E_p) is the simple product of weight (gravity g acting on the mass m) and height (h):

$$E_p = mgh$$

Energy calculations for springs are more complex than for weights. The force exerted by the spring against the compressing motion increases with compression ($F = kx$, where k is the elastic constant of the spring). It does not remain steady as the force of weight does for the lifted mass. Therefore, the potential energy equation is nonlinear:

$$E_p = \frac{1}{2}kx^2$$

Releasing the potential energy stored in these mechanical systems is as simple as dropping the mass, or letting go of the spring. The potential energy will return to the original condition (zero) when the objects are at rest in their original positions. If either the mass or the spring were attached to a machine to harness the return-motion, that stored potential energy could be used to do useful tasks.


Potential energy may be similarly defined and quantified for *any* situation where we exert a force over a parallel distance, regardless of where that force or the motivating distance comes from. For instance, the static cling you experience when you pull a wool sock out of a clothes dryer is an example of a force. By pulling that sock away from another article of clothing against the force of “static cling,” you are doing *work*, and storing *potential energy* in the tension between that sock and the rest of the clothing. In a similar manner, that stored energy could be released to do useful tasks if we placed the sock in some kind of machine harnessing the sock’s motion as it returns to its original position on the pile of laundry, pulled by the force of static electrical attraction.

If we make use of non-mechanical means to move electric charge from one location to another, the result is no different. Moving attracting charges apart from one another means doing *work* (a force exerted over a parallel distance) and storing potential energy in that physical tension. When we use chemical reactions to move electrons from one metal plate to another in a solution, or when we spin a generator and electro-magnetically motivate electrons to seek other locations, we impart potential energy to those electrons. We could express this potential energy in the same unit as we do for mechanical systems (the *Joule*). However, it is actually more useful to express the potential energy in an electric system in terms of how many joules are available per a specific quantity of electric charge (a certain number of electrons). This measure of *specific* potential energy is simply called *electric potential* or *voltage*, and we measure it in units of *Volts*, in honor of the Italian physicist Alessandro Volta, inventor of the first electrochemical battery.

$$1 \text{ Volt} = \frac{1 \text{ Joule of potential energy}}{1 \text{ Coulomb of electric charge}}$$

In other words, if we forced 1 Coulomb's worth of electrons (6.24×10^{18} of them, to be exact) away from a positively-charged place, and did one Joule's worth of work in the process, we would have generated one Volt of electric potential.

Electric potential (voltage) and potential energy share a common, yet confusing property: both quantities are fundamentally *relative* between two physical locations. There is really no such thing as specifying a quantity of potential energy at a single location. The amount of potential energy in any system is always relative between two different points. If I lift a mass off the ground, I can specify its potential energy, *but only in relation to its former position on the ground*. The amount of energy that mass is potentially capable of releasing by free-fall depends on how far it could possibly fall. To illustrate, we will perform some “thought experiments” where we lift a 1 kilogram mass 1 meter off the ground and then drop that mass to assess its energy. That 1-kilo mass weighs 9.81 Newtons on Earth, and the distance lifted was 1 meter, so the potential energy stored in the mass is 9.81 joules, right? Consider the following scenario:




If we drop the mass over the spot we first lifted it from, it will release all the potential energy we invested in it: 9.81 joules. From this example it seems very clear that the mass possessed 9.81 joules of potential energy when lifted. However, the situation becomes more complicated if we carry the mass over to the table and release it there. Since the mass can now only fall *half* a meter, it will only release 4.9 joules of energy in the process. How much potential energy did the mass have while suspended above that table? If you argue that the mass possessed a full 9.81 joules of potential energy at that original height, then why can't we realize that full potential when dropped over the table? Can we really say that an object possess potential energy if we cannot realize that potential?

Next, imagine lifting the 1 kg mass 1 meter off the ground, and then carrying it over the edge of a cliff with a 300 meter drop. If we were to release the mass here, it would fall a full 301 meters before hitting the ground. This represents *2.95 kilojoules (kJ)* of energy! How can we possibly argue that the mass had a potential energy of 9.81 joules when lifted one meter off the ground, and then without imparting any more energy to that mass we suddenly are able to realize over 300 times as much energy?

The answer to this paradox is that potential energy is always *relative* between the object's initial height and its height after falling. It is technically incorrect to say that the mass possesses 9.81 joules of potential energy when lifted 1 meter off the ground, as though potential energy were an absolute

quantity. Instead, it is more accurate to say that the mass *gained* 9.81 joules of potential energy after being lifted 1 meter off the ground, implying that the mass could have possessed potential energy (with reference to the cliff’s bottom) even while sitting on the ground (above the cliff).

Likewise, electrical voltage (as an expression of potential energy) is a quantity relative between two points. Electrical voltage is symbolized by the variable V and measured in the unit of the *Volt*. Just as we must know the mass’s position relative to its falling point before we can quantify its potential energy, we must know an electric charge’s position relative to its return point before we can quantify its voltage. Consider a series of batteries¹ connected as shown:



The voltage as measured between any two points directly across a single battery will be 1.5 volts:

$$V_{AB} = 1.5 \text{ volts}$$

$$V_{BC} = 1.5 \text{ volts}$$

$$V_{CD} = 1.5 \text{ volts}$$

If, however, we span more than one battery with our voltmeter connections, our voltmeter will register more than 1.5 volts:

$$V_{AC} = 3.0 \text{ volts}$$


$$V_{BD} = 3.0 \text{ volts}$$

$$V_{AD} = 4.5 \text{ volts}$$

Thus we see there is no such thing as “voltage” at a single point in a circuit. The concept of voltage has meaning only *between* pairs of points in a circuit, just as the concept of potential energy for a mass has meaning only *between* two physical locations: where the mass is, and where it may potentially fall.

¹A *battery* is an electrochemical device producing an electrical voltage as the result of a chemical reaction.

Things get interesting when we connect voltage sources in different configurations. Consider the following example, identical to the previous illustration except the middle battery has been reversed:




Note the “+” and “-” signs next to the ends of the batteries. These signs show the *polarity* of each battery’s voltage. Also note how the two voltmeter readings are different from before. Here we see an example of *negative potential* with the middle battery connected in opposition to the other two batteries. While the top and bottom batteries are both “lifting” electric charges to greater potential (going from point **D** to point **A**), the middle battery is decreasing potential from point **C** to point **B**. It’s like taking a step forward, then a step back, then another step forward. Or, perhaps more appropriately, like lifting a mass 1.5 meters up, then setting it down 1.5 meters, then lifting it 1.5 meters up again. The first and last steps accumulate potential energy, while the middle step releases potential energy.

This explains why it is important to install multiple batteries the same way into battery-powered devices such as radios and flashlights. The batteries’ voltages are supposed to add to make a larger total required by the device. If one or more batteries are placed backwards, potential will be lost instead of gained, and the device will not receive enough voltage.

Here we must pay special attention to how we use our voltmeter, since polarity matters. All voltmeters are standardized with two colors for the test leads: red and black. To make sense of the voltmeter's indication, especially the positive or negative *sign* of the indication, we must understand what the red and black test lead colors mean:


A positive reading indicates a **gain**
in potential from black to red.



Connecting these test leads to different points in a circuit will tell you whether there is potential gain or potential loss from one point (black) to the other point (red).

4.2 Electrical current


Current is the name we give to the motion of electric charges from a point of high potential to a point of low potential. All we need to form an electric current is a source of potential (voltage) and some electric charges that are free to move between the poles of that potential. For instance, if we connected a battery to two metal plates, we would create an electric field between those plates, analogous to a gravitational field except it only acts on electrically charged objects, while gravity acts on anything with mass. A free charge placed between those plates would “fall” toward one of the plates just as a mass would fall toward a larger mass:



An electric charge will "fall" in an electric field just as a mass will fall in a gravitational field.

Some solid substances, most notably metals, have very mobile electrons. That is, the outer (valence) electrons are very easily dislodged from the parent atoms to drift to and fro throughout the material. In fact, the electrons within metals are so free that physicists sometimes refer to the structure of a metal as atomic nuclei “floating in a sea of electrons”. The electrons are almost fluid in their mobility throughout a solid metal object, and this property of metals may be exploited to form definite pathways for electric currents. This is why electrical wires are made of metal such as copper or aluminum: substances with a high degree of electron mobility. Any substance whose electrons are mobile is called a *electrical conductor*, while any substance lacking mobile electrons is called an *electrical insulator*.

If the poles of a voltage source are joined by a conductor, the free electrons within that conductor will drift toward the positive pole (electrons having a negative charge, opposite charges attracting one another). For each electron reaching the positive pole, an electron exits the negative pole of the source to replenish the total number of electrons in the flow:




If the source of this voltage is continually replenished by chemical energy, mechanical energy, or some other form of energy, the free electrons will continually loop around this circular path. We call this unbroken path an *electric circuit*. The drifting motion of electrons in a circuit has the same average rate of flow (current) at all points in that circuit, because there is only one pathway for the current. You may think of this like liquid flowing through a circular loop of pipe: since there is only one pathway for the liquid to flow, the rate of flow at all points in that pathway must be the same.

We typically measure the amount of current in a circuit by the unit of *Ampères*, or *Amps* for short (named in honor of the French physicist André Ampère, and symbolize it in formulae using the variable I). One ampere of current is equal to one coulomb of electric charge (6.24×10^{18} electrons) moving past a point in a circuit for every second of time.

Like masses falling toward a source of gravity, these electrons continually “fall” toward the positive pole of a voltage source. After arriving at that source, the energy imparted by that source “lifts” the electrons to a higher potential state where they once again “fall down” to the positive pole through the circuit.

Like rising and falling masses in a gravitational field, these electrons act as carriers of energy within the electric field of the circuit. This is very useful, as we can use them to convey energy from one place to another, using metal wires as conduits for this energy. This is the basic idea behind electric power systems: a source of power (a *generator*) is turned by some mechanical engine (windmill, water turbine, steam engine, etc.), creating an electric potential. This potential is then used to motivate free electrons inside the metal wires to drift in a common direction. The electron drift is conveyed in a circuit through long wires, where they can do useful work at a *load* device such as an electric motor, light bulb, or heater.




Given the proper metal alloys, the friction that electrons experience within the metal wires may be made very small, allowing nearly all the energy to be expended at the load (motor), with very little wasted along the path (wires). This makes electricity the most efficient means of energy transport known.

The electric currents common in electric power lines may range from hundreds to thousands of amperes. The currents conveyed through power receptacles in your home typically are no greater than 20 amperes. The currents in the small battery-powered circuits you will build are even less: fractions of an ampere. For this reason, we commonly use the metric prefix *milli* (one one-thousandth) to express these small currents. For instance, 10 milliamperes is 0.010 amperes, and 500 milliamperes is one-half of an ampere.


4.2.1 Electron versus conventional flow

When Benjamin Franklin proposed his single-fluid theory of electricity, he defined “positive” and “negative” as the surplus and deficiency of electric charge, respectively. These labels were largely arbitrary, as Mr. Franklin had no means of identifying the actual nature of electric charge carriers with the primitive test equipment and laboratory techniques of his day. As (bad) luck would have it, his hypothesis was precisely opposite of the truth for metallic conductors, where electrons are the dominant charge carrier.


This means that in an electric circuit consisting of a battery and a light bulb, electrons slowly move from the negative side of the battery, through the metal wires, through the light bulb, and on to the positive side of the battery as such:



Unfortunately, scientists and engineers had grown accustomed to Franklin’s false hypothesis long before the true nature of electric current in metallic conductors was discovered. Their preferred notation was to show electric current flowing from the positive pole of a source, through the load, returning to the negative pole of the source:




This relationship between voltage polarity marks and conventional flow current tends to make more intuitive sense than electron flow notation, because it is reminiscent of fluid pressure and flow direction:



If we take the “+” sign to represent *more* pressure and the “-” sign to represent *less* pressure, it makes perfect sense that fluid should move from the high-pressure (discharge) port of the pump through the hydraulic “circuit” and back to the low-pressure (suction) port of the pump. It also makes perfect sense that the upstream side of the valve (a fluid restriction) will have a greater pressure than the downstream side of the valve. In other words, conventional flow notation best honors Mr. Franklin’s original intent of modeling current as though it were a fluid, even though he was later proven to be mistaken in the case of metallic conductors where electrons are the dominant charge carrier.

This convention was so well-established in the electrical engineering realm that it held sway despite the discovery of electrons. Engineers, who create the symbols used to represent the electronic devices they invent, consistently chose to draw arrows in the direction of conventional flow rather than electron flow. In each of the following symbols, the arrow heads point in the direction that *positive* charge carriers would move (opposite the direction that electrons actually move):



This stands in contrast to electronics technicians, who historically have been taught using electron flow notation. I remember sitting in a technical school classroom being told by my teacher to always imagine the electrons moving *against the arrows* of the devices, and wondering why it mattered.

It is truly a sad situation when the members of two branches within the same field do not agree on something as fundamental as the convention used to denote flow in diagrams. It is even worse when people within the field argue over which convention is best. So long as one is consistent with their convention and with their thinking, *it does not matter!* Many fine technologists may be found on either side of this “fence,” and some are adept enough to switch between both without getting confused.


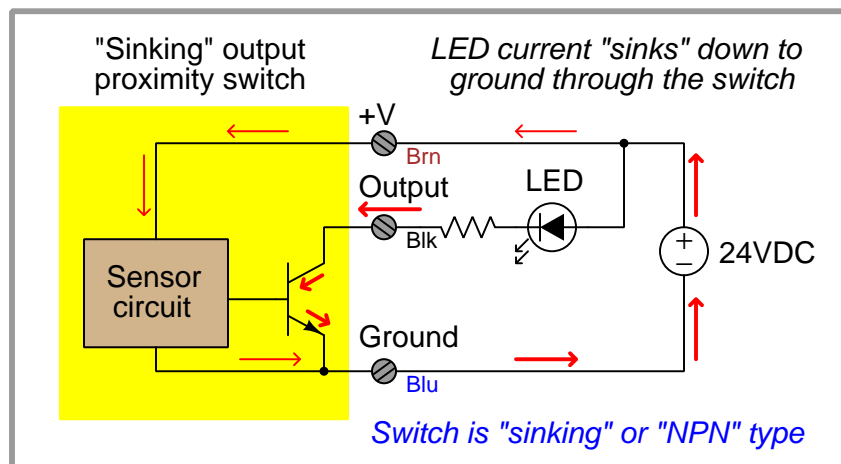
My personal preference is conventional flow notation. The only objective arguments I have in favor of this preference are as follows:

- Conventional flow notation is a closer analogue to fluid flow in pneumatic, hydraulic, and process flow systems. Since instrument technicians need to understand the relationships between pressure and flow in fluid systems as well as electrical circuits, using conventional flow notation for electrical circuits helps to reinforce these analogous concepts.
- Conventional flow notation is the *standard* for modern manufacturers' documentation (reference manuals, troubleshooting guides, datasheets, etc.)².
- Conventional flow notation is consistent with the “right-hand rule” for vector cross products (which are essential for understanding electromagnetics at advanced academic levels). The so-called “left-hand rule” taught to students learning electron flow notation is mathematically wrong, and must be un-learned if the student ever progresses to the engineering level in his or her studies.
- Conventional flow notation matches all device arrows; no need to “go against the arrow” when tracing current in a schematic diagram.
- Conventional flow notation makes sense of the descriptive terms *sourcing* and *sinking*.

This last point merits further investigation. The terms “sourcing” and “sinking” are often used in the study of digital electronics to describe the direction of current in a switching circuit. A circuit that “sources” current to a load is one where the direction of conventional flow points outward from the sourcing circuit to the load device. These terms are also used to characterize the DC input and output modules of certain industrial control equipment such as programmable logic controllers (PLCs), and so any technologist working with such equipment will need to properly identify and connect these module types based on the directions of electric current for which they are designed.

²I have yet to read a document of any kind written by an equipment manufacturer using electron flow notation, and this is after scrutinizing literally hundreds of documents looking for this exact detail! For the record, though, most technical documents do not bother to draw a direction for current at all, leaving it to the imagination of the reader instead. It is only when a direction must be drawn that one sees a strong preference in industry for conventional flow notation.

For example, here are two schematic diagrams showing two different kinds of electronic proximity switch. The first switch *sinks* current in from the LED through its output terminal, through its transistor, and down to ground. The second switch *sources* current from the positive supply terminal through its transistor and out to the LED through its output terminal (note the direction of the thick arrow near the output screw terminal in each circuit):



These terms simply make no sense when viewed from the perspective of electron flow notation. If you were to actually trace the directions of the electrons, you would find that a device "sourcing" current has electrons flowing *into* its connection terminal, while a device "sinking" current sends electrons *out* to another device where they travel (up) to a point of more positive potential.

In fact, the association between conventional flow notation and sourcing/sinking descriptions is so firm that I have yet to see a professionally published textbook on digital circuits that uses electron flow³. This is true even for textbooks written for technicians and not engineers!

³If by chance I have missed anyone's digital electronics textbook that does use electron flow, please accept my

Once again, though, it should be understood that either convention of current notation is adequate for circuit analysis. I dearly wish this horrible state of affairs would come to an end, but the plain fact is it will not. Electron flow notation may have the advantage of greater correspondence to the actual state of affairs (in the vast majority of circuits), but conventional flow has the weight of over a hundred years of precedent, cultural inertia, and convenience. No matter which way you choose to think, at some point you will be faced with the opposing view.

Choose the notation you like best, and may you live long and prosper.

4.3 Electrical sources and loads


By definition, a *source* is a device delivering energy into a system, while a *load* is a device extracting energy from a system. Examples of typical electrical sources include generators, photovoltaic cells, thermopiles, and primary-cell batteries. These devices create electrical voltage, which in turn motivates electrical current to flow in a circuit. Examples of typical electrical loads include resistors, lamps, and electric motors. These devices resist the flow of electrical current through them, creating a voltage *drop* as a result.

In a working circuit, electrical sources and loads may be easily distinguished by comparison of their current directions and voltage polarities. An electrical source always manifests a voltage polarity in a direction *aiding* the direction of charge flow; i.e. a source is “pushing” the current along. An electrical load always manifests a voltage polarity in a direction *opposing* the direction of charge flow; i.e. a load “resists” the current.

The way in which we designate the direction of current (charge flow) becomes very important here. Since there are two commonly accepted notations – electron flow and “conventional” flow, exactly opposite of each other – it is easy to become confused.

apologies. I can only speak of what I have seen myself.


First we see a diagram showing a source and a load, using electron flow notation. Electrons, being negatively charged particles, are repelled by the negative (−) poles of both source and load, and attracted to the positive (+) poles of both source and load. The difference between source and load is that the source device *motivates* the flow of electrons while the load device *resists* the flow of electrons:



Electrons are repelled by the (−) poles
and attracted to the (+) poles

In the case of the source (battery), the polarity of the voltage works *for* the direction of charge motion. In the case of the load (resistor), the polarity of the voltage drop works *against* the direction of charge motion.


Next we see a diagram showing the same source and load, this time using “conventional” flow notation to designate the direction of current. Here we must imagine positively-charged carriers moving through the wires instead of electrons. These positive charge carriers are repelled by any positive (+) pole and attracted to any negative (−) pole. Viewed in this light, we see the exact same principle at work: the source device is seen to *motivate* the flow of these positive charge carriers while the load device *resists* the flow:



Positive charge carriers are repelled by the
(+) poles and attracted to the (−) poles

Despite using a different notation for charge motion, the concept of source and load remains the same. In the case of the source (battery), the polarity of the voltage works *for* the direction of charge motion. In the case of the load (resistor), the polarity of the voltage drop works *against* the direction of charge motion.

An alternative notation for voltage (other than using “+” and “-” symbols) that many students find particularly illustrative is the use of curved arrows, where the tip of the curved arrow is the positive pole and the tail of the curved arrow is the negative pole. This notation is intended to be used when the direction of current (using “straight” or “angular” arrows) is shown using conventional flow notation:






Curved arrows show direction of voltage
(electro-motive force, or EMF)

Straight arrows show direction of current
(conventional flow assumed; positive charges moving)

Using arrows to represent voltage polarity in addition to using arrows to represent current direction is highly intuitive. It shows which way each component in the DC circuit is “pushing” in relation to the flow of charge carriers. Note how the source’s voltage arrow points in the same direction as the current: this means the source is motivating the current, causing charge carriers to flow in this circuit. Note how the resistor’s voltage arrow points opposite to the direction of current: this means the resistor is opposing the current, in a sense “fighting against” the flow of charge carriers. This comparison of voltage-arrow versus current-arrow direction makes the distinction between sources and loads rather obvious: sources push in the direction of current while loads push against the direction of current.

I personally lament the obscurity of this “curved-arrow” notation for voltage, as it greatly aids comprehension of this critically important distinction between sources and loads. When the voltage and current arrows point in the same direction, it means the component in question is motivating charge carriers along and therefore imparts energy to the circuit. When voltage and current arrows point in opposite directions, it means the component in question opposes charge carrier motion and therefore acts to extract energy from the circuit.

This directly relates to the fundamental physics concept of *work*, specifically in relation to the mathematical sign of work being a function of the relative angle between force and displacement. When a force acts in the same direction as motion, the work done is positive; when a force acts in the opposite direction as motion, the work done is negative:




$W_{battery}$ = positive value, because its force is in the same direction as the current


$W_{resistor}$ = negative value, because its force opposes the direction of the current

Positive work represents an infusion of energy into a system (source), while negative work represents an extraction of energy from that system (load).

If we examine a hydraulic system, where a pump pushes fluid around a pipe loop and an orifice (called a “restrictor”) restricts the flow of this fluid, we see this same concept in action: the pump’s pressures at its discharge and suction ports work *for* the direction of fluid flow, while the pressures at the upstream and downstream ports of the restrictor work *against* the direction of fluid flow. The pump acts as a power *source* in this hydraulic “circuit” (infusing energy into the system) while the restrictor acts as a power *load* (extracting energy from the system):




We may even see this concept revealed in a simple mechanical system where work is being done. Examine the case of a crane lifting a heavy weight into the air, shown below. As the crane lifts the weight upward, the crane’s upward force on the weight is clearly working *for* the direction of motion, while the weight’s downward force against the crane is clearly working *against* the direction of motion:



Thus, the crane is doing positive work (acting as a *source*, infusing potential energy into the weight) while the weight is doing negative work (acting as a *load*, absorbing potential energy from the crane). A mathematically rigorous way to demonstrate this is to calculate the work done by each using the formula $W = \vec{F} \cdot \vec{x}$ or $W = Fx \cos \theta$. Since the crane’s force and motion vectors both point in the same direction, $\theta = 0$ and work is a positive quantity Fx . The weight’s force vector, however, points 180° away from the motion vector, and so its work calculation is $Fx \cos(180^\circ)$ or $-Fx$.

Some electric components have the ability to act as sources and loads at different times. Both capacitors (see section 4.12 starting on page 346) and inductors (see section 4.13 starting on page 350) have the ability to temporarily contribute to and extract energy from electrical circuits, both having the ability to act as energy *storage* devices. One of the key concepts necessary to grasp the energy-storing abilities of capacitors and inductors is being able to recognize sources and loads at a glance based on the relationship between voltage polarity and charge motion. A set of three schematic diagrams shows how a capacitor is able to play the role of either source or load depending on what other component it is connected to:



Rechargeable batteries (called “secondary-cell” batteries as opposed to “primary-cell” batteries which cannot be recharged) may also behave as either sources or loads depending on external conditions. If a secondary-cell battery is connected to a resistor, the battery will discharge its energy (i.e. act as a source) while the resistor will dissipate that energy (i.e. act as a load). If a depleted secondary-cell battery is connected to an electrical generator of greater voltage, the generator will source energy to the battery while the re-charging battery will load down the generator (i.e. conventional flow entering the battery’s positive terminal and exiting the negative terminal).

Another practical benefit of clearly comprehending the distinction between electrical sources and electrical loads is being able to understand and troubleshoot 4-20 mA signal “loop” circuits used extensively in industrial instrumentation, especially circuits containing 2-wire (“loop-powered”) process transmitters. A “2-wire transmitter” is a device designed to regulate the amount of electrical current through it to a value determined by some physical variable such as a sensed pressure, temperature, or flow rate. The purpose of such a device is to represent that physical measurement in the form of an electric current that may be carried far distances through wires. What makes this device so troublesome for people to understand is that despite its function to *set* the value of current in the circuit, it is actually an electrical *load* and not an electrical *source* as one might assume. That is, a 2-wire transmitter relies wholly on some other source of electrical power in the circuit to function, although the transmitter solely defines how much current will flow in the circuit by virtue of its function as a regulator. For more information on this subject, refer to section 13.5 starting on page 891.

4.4 Electrical power

Since we often use electricity to convey energy from one location to another, it is important to be able to quantify and calculate the rate at which energy is delivered by an electric circuit. The *rate* at which energy transfers from one location to another is technically referred to as *power* (P), and it is typically⁴ measured in the unit of the *watt* (1 watt = 1 joule of energy per second of time).

The total amount of power (energy per unit time) carried by an electric current is proportional to the strength of that current as well as the amount of potential (voltage) between the poles of the electrical source. Thus, power is equal to the product of current and voltage in a formula called *Joules' Law*:

$$P = IV$$

Where,

P = Power in watts (joules of energy per second of time)

I = Current in amperes

V = Voltage in volts

We may prove the correctness of this formula by verifying all the units of measurement agree:

$$P = IV$$

$$[\text{Watts}] = [\text{Amperes}] \times [\text{Volts}] \quad \text{or} \quad [\text{W}] = [\text{A}][\text{V}]$$

$$\left[\frac{\text{Joules}}{\text{Seconds}} \right] = \left[\frac{\text{Coulombs}}{\text{Seconds}} \right] \times \left[\frac{\text{Joules}}{\text{Coulombs}} \right] \quad \text{or} \quad \left[\frac{\text{J}}{\text{s}} \right] = \left[\frac{\text{C}}{\text{s}} \right] \left[\frac{\text{J}}{\text{C}} \right]$$

Note how the basic units for power (joules per second) are indeed equal to the product of voltage (joules per coulomb) and current (coulombs per second). This process of checking for agreement between units of measurement in a physics formula is called *dimensional analysis*.

⁴Although the unit of the “watt” is commonly used for electrical power, other units are valid as well. The British unit of *horsepower* is every bit as valid for expressing electrical power as “watts,” although this usage is less common. Likewise, the “watt” may be used to express measurements of non-electrical power as well, such as the mechanical power output of an engine. European automobile manufacturers, for example, rate the power output of their cars’ engines in *kilowatts*, as opposed to American automobile manufacturers who rate their engines in *horsepower*. This choice of units is strictly a cultural convention, since any valid unit for power may be applied to any form of energy rate.

4.5 Electrical resistance and Ohm's Law

To review, *voltage* is the measure of potential energy available to electric charges. *Current* is the uniform drifting of electric charges in response to a voltage. We can have a voltage without having a current, but we cannot have a current without first having a voltage to motivate it⁵. Current without voltage would be equivalent to motion without a motivating force.

When electric charges move through a material such as metal, they will naturally encounter some friction, just as fluid moving through a pipe will inevitably encounter friction⁶. We have a name for this friction to electrical charge motion: *resistance*. Like voltage and current, resistance has its own special unit of measurement: the *ohm*, named in honor of the German physicist Georg Simon Ohm.

At this point it would be good to summarize and compare the symbols and units we use for voltage, current, and resistance:

Quantity	Algebraic symbol	Unit	Unit abbreviation
Voltage	V (or E)	Volt	V
Current	I	Ampere (or Amp)	A
Resistance	R	Ohm	Ω

Ohm defined resistance as the mathematical ratio between applied voltage and resulting current. This formula became known as *Ohm's Law*, perhaps the most basic formula in all of electrical science (shown here in three different forms, each one solving for a different variable):

$$R = \frac{V}{I} \quad V = IR \quad I = \frac{V}{R}$$

Verbally expressed, resistance is how much voltage it takes to force a certain rate of current through a conductive material. Many materials have relatively stable resistances, while others do not. Devices called *resistors* are sold which are manufactured to possess a very precise amount of resistance, for the purpose of limiting current in circuits (among other things).

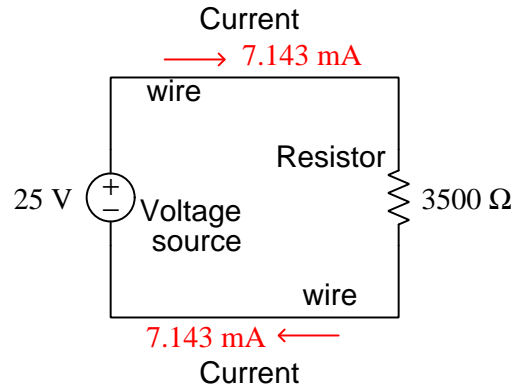
Here is an example of Ohm's Law in action: calculate the amount of current in a circuit with a voltage source of 25 V and a total resistance of 3500 Ω . Taking 25 volts and dividing by 3500 ohms, you should arrive at a result of 0.007143 amperes, or 7.143 milliamperes (7.143 mA).

One of the most challenging aspects of Ohm's Law is remembering to *keep all variables in context*. This is a common problem for many students when studying physics as well: none of the equations learned in a physics class will yield the correct results unless all the variables relate to the same object or situation. For instance, it would make no sense to try to calculate the kinetic energy of a moving object ($E = \frac{1}{2}mv^2$) by taking the mass of one object (m) and multiplying it by the square of the velocity of some *other* object (v^2). Likewise, with Ohm's Law, we must make sure the voltage, current, and resistance values we are using all relate to the same portion of the same circuit.

⁵Except in the noteworthy case of *superconductivity*, a phenomenon occurring at extremely low temperatures.


⁶Except in the noteworthy case of *superfluidity*, another phenomenon occurring at extremely low temperatures.

If the circuit in question has only one source of voltage, one resistance, and one path for current, we cannot mis-apply Ohm's Law. Expressing the previous example in a schematic diagram:




Note: arrows point in the direction of conventional flow

However, if we look at a more complex circuit, we encounter the potential for mis-application of Ohm's Law because we have multiple resistances in the circuit rather than just one resistance:



Which resistance do we use to calculate current in this circuit? Do we divide our 25 volts by 3500 ohms like we did last time, or do we divide it by 1500 ohms, or something entirely different? The answer to this question lies in the identification of voltages and currents. We know that the 25 volt potential will be impressed across the *total* of the two resistances R_1 and R_2 , and since there is only one path for current they must share the same current. Thus, we actually have *three* voltages (V_1 , V_2 , and V_{total}), *three* resistances (R_1 , R_2 , and R_{total}), and only *one* current (I):



Note: arrows point in the direction of conventional flow

Using the $V = IR$ form of Ohm's Law to relate these three voltages (V_1 , V_2 , and V_{total}) to the one current (I), we end up with three equations for this circuit:

$$V_1 = IR_1$$

$$V_2 = IR_2$$

$$V_{total} = IR_{total} = I(R_1 + R_2)$$

We can only solve for one unknown variable at a time in any equation. This means we cannot solve for V_1 yet, because although we do know the value of resistor R_1 (3500 ohms), we do not yet know the circuit's current (I). Likewise for V_2 , because we do not yet know the value of I . However, the third equation is solvable, since we happen to know the total voltage as well as both resistor values, leaving circuit current I as the only unknown variable. Manipulating this equation and solving for I :


$$I = \frac{V_{total}}{R_{total}} = \frac{V_{total}}{R_1 + R_2} = \frac{25 \text{ V}}{3500 \Omega + 1500 \Omega} = 0.005 \text{ A} = 5 \text{ mA}$$

Now that we know the amount of current in this circuit, we may solve for V_1 and V_2 in the other equations to find that the voltage across resistor R_1 is 17.5 volts, and that the voltage across resistor R_2 is 7.5 volts.


4.6 Series versus parallel circuits

In addition to Ohm's Law, we have a set of rules describing how voltages, currents, and resistances relate in circuits comprised of multiple resistors. These rules fall neatly into two categories: *series* circuits and *parallel* circuits. The two circuit types are shown here, with squares representing any type of two-terminal electrical component:

Series circuit



Parallel circuit




The defining characteristic of a series electrical circuit is it provides just one path for current. This means there can be only one value for current anywhere in the circuit, the exact same current for all components at any given time⁷. The principle of current being the same everywhere in a series circuit is actually an expression of a more fundamental law of physics: the *Conservation of Charge*, which states that electric charge cannot be created or destroyed. In order for current to have different values at different points in a series circuit indefinitely, electric charge would have to somehow appear and disappear to account for greater rates of charge flow in some areas than in others. It would be the equivalent of having different rates of water flow at different locations along one length of pipe⁸.

⁷Interesting exceptions do exist to this rule, but only on very short time scales, such as in cases where we examine the a transient (pulse) signal nanosecond by nanosecond, and/or when very high-frequency AC signals exist over comparatively long conductor lengths.

⁸Those exceptional cases mentioned earlier in the footnote are possible only because electric charge may be temporarily stored and released by a property called *capacitance*. Even then, the law of charge conservation is not violated because the stored charges re-emerge as current at later times. This is analogous to pouring water into a bucket: just because water is poured into a bucket but no water leaves the bucket does not mean that water is magically disappearing. It is merely being stored, and can re-emerge at a later time.

Series circuits are defined by having only one path for current, and this means the steady-state current in a series circuit must be the same at all points of that circuit. It also means that the sum of all voltages dropped by load devices must equal the sum total of all source voltages, and that the total resistance of the circuit will be the sum of all individual resistances:




The defining characteristic of a parallel circuit, by contrast, is that all components share the same two equipotential points. “Equipotential” simply means “at the same potential” which points along an uninterrupted conductor must be⁹. This means there can be only one value of voltage anywhere in the circuit, the exact same voltage for all components at any given time¹⁰. The principle of voltage being the same across all parallel-connected components is (also) an expression of a more fundamental law of physics: the *Conservation of Energy*, in this case the conservation of specific potential energy which is the definition of voltage. In order for voltage to differ between parallel-connected components, the potential energy of charge carriers would have to somehow appear and disappear to account for lesser and greater voltages. It would be the equivalent of having a “high spots” and “low spots” of water mysteriously appear on the quiet surface of a lake, which we know cannot happen because water has the freedom to move, meaning any high spots would rush to fill any low spots¹¹.

⁹An ideal conductor has no resistance, and so there is no reason for a difference of potential to exist along a pathway where nothing stands in the way of charge motion. If ever a potential difference developed, charge carriers within the conductor would simply move to new locations and neutralize the potential.

¹⁰Again, interesting exceptions do exist to this rule on very short time scales, such as in cases where we examine the a transient (pulse) signal nanosecond by nanosecond, and/or when very high-frequency AC signals exist over comparatively long conductor lengths.

¹¹The exceptional cases mentioned in the previous footnote exist only because the electrical property of *inductance* allows potential energy to be stored in a magnetic field, manifesting as a voltage different along the length of a conductor. Even then, the Law of Energy Conservation is not violated because the stored energy re-emerges at a later time.

The sum of all component currents must equal the total current in a parallel circuit, and total resistance will be *less* than the smallest individual resistance value:



The rule for calculating total resistance in a parallel circuit perplexes many students with its weird compound reciprocal notation. There is a more intuitive way to understand this rule, and it involves a different quantity called *conductance*, symbolized by the letter G .

Conductance is defined as the reciprocal of resistance; that is, a measure of how *easily* electrical charge carriers may move through a substance. If the electrical resistance of an object doubles, then it now has *half* the conductance it did before:

$$G = \frac{1}{R}$$

It should be intuitively apparent that conductances add in parallel circuits. That is, the total amount of conductance for a parallel circuit must be the sum total of all individual conductances, because the addition of more conductive pathways must make it easier overall for charge carriers to move through the circuit. Thus,

$$G_{\text{total}} = G_1 + G_2 + \dots + G_n$$

The formula shown here should be familiar to you. It has the same form as the total resistance formula for series circuits. Just as resistances add in series (more series resistance makes the overall resistance to current increase), conductances add in parallel (more conductive branches makes the overall conductance increase).

Knowing that resistance is the reciprocal of conductance, we may substitute $\frac{1}{R}$ for G wherever we see it in the conductance equation:

$$\frac{1}{R_{\text{total}}} = \frac{1}{R_1} + \frac{1}{R_2} + \dots + \frac{1}{R_n}$$

Now, to solve for R_{total} , we need to reciprocate both sides:

$$R_{total} = \frac{1}{\frac{1}{R_1} + \frac{1}{R_2} + \cdots + \frac{1}{R_n}}$$


For both series and parallel circuits, total power dissipated by all load devices is equal to the total power delivered by all source devices. The configuration of a circuit is irrelevant to the balance between power supplied and power lost, because this balance is an expression of the Law of Energy Conservation.

4.7 Kirchhoff's Laws

Two extremely important principles in electric circuits were codified by Gustav Robert Kirchhoff in the year 1847, known as *Kirchhoff's Laws*. His two laws refer to voltages and currents in electric circuits, respectively.

Kirchhoff's Voltage Law states that the algebraic sum of all voltages in a closed loop is equal to zero. Another way to state this law is to say that for every rise in potential there must be an equal fall, if we begin at any point in a circuit and travel in a loop back to that same starting point.

An analogy for visualizing Kirchhoff's Voltage Law is hiking up a mountain. Suppose we start at the base of a mountain and hike to an altitude of 5000 feet to set up camp for an overnight stay. Then, the next day we set off from camp and hike farther up another 3500 feet. Deciding we've climbed high enough for two days, we set up camp again and stay the night. The next day we hike down 6200 feet to a third location and camp once gain. On the fourth day we hike back to our original starting point at the base of the mountain. We can summarize our hiking adventure as a series of rises and falls like this:




Day	Path	Altitude gain/loss
Day 1	A to B	+5000 feet
Day 2	B to C	+3500 feet
Day 3	C to D	-6200 feet
Day 4	D to A	-2300 feet
(Total)	ABCD A	0 feet

Of course, no one would tell their friends they spent *four days hiking a total altitude of 0 feet*, so people generally speak in terms of the *highest* point reached: in this case 8500 feet. However, if we track each day's gain or loss in algebraic terms (maintaining the mathematical sign, either positive or negative), we see that the end sum is zero (and indeed *must always be zero*) if we finish at our starting point.

If we view this scenario from the perspective of potential energy as we lift a constant mass from point to point, we would conclude that we were doing work on that mass (i.e. investing energy in it by lifting it higher) on days 1 and 2, but letting the mass do work on us (i.e. releasing energy by lowering it) on days 3 and 4. After the four-day hike, the net potential energy imparted to the mass is zero, because it ends up at the exact same altitude it started at.

Let's apply this principle to a real circuit, where total current and all voltage drops have already been calculated for us:



Arrow shows current in the direction of conventional flow notation

If we trace a path ABCDEA, we see that the algebraic voltage sum in this loop is zero:

Path	Voltage gain/loss
A to B	- 4 volts
B to C	- 6 volts
C to D	+ 5 volts
D to E	- 2 volts
E to A	+ 7 volts
ABCDEA	0 volts

We can even trace a path that does not follow the circuit conductors or include all components, such as EDCBE, and we will see that the algebraic sum of all voltages is still zero:

Path	Voltage gain/loss
E to D	+ 2 volts
D to C	- 5 volts
C to B	+ 6 volts
B to E	- 3 volts
EDCBE	0 volts

Kirchhoff's Voltage Law is often a difficult subject for students, precisely because voltage itself is a difficult concept to grasp. Remember that there is no such thing as voltage at a single point; rather, voltage exists only as a *differential* quantity. To intelligently speak of voltage, we must refer to either a *loss* or *gain* of potential between **two points**.


Our analogy of altitude on a mountain is particularly apt. We cannot intelligently speak of some point on the mountain as having a specific altitude unless we assume a point of reference to measure from. If we say the mountain summit is 9200 feet high, we usually mean 9200 feet *higher than sea level*, with the level of the sea being our common reference point. However, our hiking adventure where we climbed 8500 feet in two days did not imply that we climbed to an absolute altitude of 8500 feet above sea level. Since I never specified the sea-level altitude at the base of the mountain,

it is impossible to calculate our absolute altitude at the end of day 2. All you can tell from the data given is that we climbed 8500 feet *above* the mountain base, wherever that happens to be with reference to sea level.

So it is with electrical voltage as well: most circuits have a point labeled as *ground* where all other voltages are referenced. In DC-powered circuits, this ground point is often the negative pole of the DC power source¹². Voltage is fundamentally a quantity relative between two points: a measure of how much potential has *increased* or *decreased* moving from one point to another.

Kirchhoff's Current Law is a much easier concept to grasp. This law states that the algebraic sum of all currents at a junction point (called a *node*) is equal to zero. Another way to state this law is to say that for every electron entering a node, one must exit somewhere.


An analogy for visualizing Kirchhoff's Current Law is water flowing into and out of a "tee" fitting:



So long as there are no leaks in this piping system, every drop of water entering the tee must be balanced by a drop exiting the tee. For there to be a continuous mis-match between flow rates would imply a violation of the Law of Mass Conservation.

¹²But not always! There do exist positive-ground systems, particularly in telephone circuits and in some early automobile electrical systems.

Let's apply this principle to a real circuit, where all currents have been calculated for us:



Arrows show currents in the direction of conventional flow notation

At nodes where just two wires connect (such as points A, B, and C), the amount of current going in to the node exactly equals the amount of current going out (4 mA, in each case). At nodes where three wires join (such as points D and E), we see one large current and two smaller currents (one 4 mA current versus two 2 mA currents), with the directions such that the sum of the two smaller currents form the larger current.


Just as the balance of water flow rates into and out of a piping “tee” is a consequence of the Law of Mass Conservation, the balance of electric currents flowing into and out of a circuit junction is a consequence of the Law of *Charge* Conservation, another fundamental conservation law in physics.

4.8 Circuit fault analysis

Perhaps the most valuable skill an instrument technician can possess is the ability to efficiently diagnose malfunctioning systems: determining in as short a time as possible the cause of a system's malfunction. Since most instrumentation and control systems are based on electricity, a solid understanding of electrical faults is the foundation of this skill set. In this section we will explore the two basic types of electrical faults (*opens* and *shorts*) and analyze their respective effects in DC circuits.


An *open* is an electrical fault whereby the pathway for electrical current is broken, thus preventing the passage of current. A *short* is an electrical fault whereby two points in a circuit that are supposed to be separated are joined together by a conductive pathway. It should be noted that this definition for the word “short” both is technical and specific. This is important to understand, as many people tend to use the word “short” to refer to any electrical problem in general. In technical parlance, a “short” fault is the exact opposite of an “open” fault, and should never be confused one for the other.

Let us examine the effects of both faults in a simple two-resistor DC circuit. We will begin with an analysis of the circuit in a healthy condition, showing all values of voltage and current:



Being a series circuit, the current is the same through all components while the two loads' voltage drops add to equal the source voltage.

Now suppose resistor R_1 fails open. We know that a continuous electric current is able to exist only where there is a continuous path for that current. If a component fails open, the continuity of the circuit will be broken, and current must halt. Since this is a series circuit, if current halts in one location it must likewise halt in all locations:




We would still expect to measure 15 volts between test points A and B because it is the nature of a voltage source to maintain a constant voltage across its terminals (i.e. to maintain the same amount of energy difference per unit charge from one side to the other). We may apply Ohm's Law to calculate the voltage across the healthy resistor R_2 : since $V = IR$ and $I = 0$, we may safely conclude that $V = 0$ for resistor R_2 .


Unfortunately, we cannot apply Ohm's Law to the failed resistor R_1 since the resistance of any failed-open component is infinite. Infinities do not lend themselves to arithmetic calculations (how much is zero times infinity?), and so we must find some other way to determine the voltage dropped across the open R_1 . Here we find that Kirchhoff's Voltage Law applies very nicely: if we know that the algebraic sum of all voltages in a loop must equal zero, and we know all but one of the voltages in the loop BACDB, we may calculate the last voltage by simple addition and subtraction:

Path	Voltage gain/loss	Reason
B to A	+ 15 volts	Source voltage value
A to C	-15 volts	Necessary to satisfy KVL
C to D	0 volts	These points are equipotential
D to B	0 volts	Ohm's Law: $V = 0 \text{ amps} \times 500 \text{ ohms}$
BACDB	0 volts	KVL sum = 0

Clearly, the voltage between points A and C (across the failed-open R_1) must be -15 volts in order to satisfy Kirchhoff's Voltage Law. The negative sign of V_{AC} tells us that point C must be at a lesser potential than point A (i.e. A is positive and C is negative).




Let us suppose now resistor R_1 fails shorted. This means its two terminals are now directly connected to each other rather than being separated by 1000 ohms of electrical resistance. In other words, the two terminals of resistor R_1 have been made *equipotential* by the shorted fault. By extension, this forces points A and C to be equipotential as well. Since points C and D are already equipotential to each other by virtue of the wire connecting them, points A through D must now be equipotential.



If points A through D are all equipotential, it means the right-hand terminal of resistor R_2 is directly connected to the positive terminal of the 15 volt source. The left-hand terminal of R_2 is already equipotential with the negative terminal of the source. This means the shorted fault at R_1 has placed the full source potential across resistor R_2 . An analysis using Kirchhoff's Law confirms this:

Path	Voltage gain/loss	Reason
B to A	+ 15 volts	Source voltage value
A to C	0 volts	Shorting points makes them equipotential
C to D	0 volts	These points are equipotential
D to B	-15 volts	Necessary to satisfy KVL
BACDB	0 volts	KVL sum = 0


The negative sign of V_{DB} tells us that point B must be at a lesser potential than point D (i.e. D is positive and B is negative). Calculating the amount of current in this circuit is now a simple matter of applying Ohm's Law to resistor R_2 . Given a 15 volt potential across it and 500 ohms of resistance, the current is equal to 30 milliamps:



All DC circuit fault analysis reduces to these simple principles: *open* faults directly affect current by interrupting the continuity of the circuit, while *shorted* faults directly affect voltage by making points equipotential to each other that were not equipotential before. The rest is merely applying Kirchhoff's and Ohm's Laws to determine consequences of the fault throughout the circuit.

4.9 Bridge circuits

A *bridge* circuit is basically a pair of voltage dividers where the circuit output is taken as the difference in potential between the two dividers. Bridge circuits may be drawn in schematic form in an H-shape or in a diamond shape, although the diamond configuration is more common:




The voltage source powering the bridge circuit is called the *excitation* source. This source may be DC or AC depending on the application of the bridge circuit. The components comprising the bridge need not be resistors, either: capacitors, inductors, lengths of wire, sensing elements, and other component forms are possible, depending on the application.

Two major applications exist for bridge circuits, which will be explained in the following subsections.

4.9.1 Component measurement

Bridge circuits may be used to test components. In this capacity, one of the “arms” of the bridge circuit is comprised of the component under test, while at least one of the other “arms” is made adjustable. The common *Wheatstone bridge* circuit for resistance measurement is shown here:




Fixed resistors R_1 and R_2 are of precisely known value and high precision. Variable resistor R_{adjust} has a labeled knob allowing for a person to adjust and read its value to a high degree of precision. When the ratio of the variable resistance to the specimen resistance equals the ratio of the two fixed resistors, the sensitive galvanometer will register exactly zero volts regardless of the excitation source’s value. This is called a *balanced* condition for the bridge circuit:

$$\frac{R_1}{R_2} = \frac{R_{adjust}}{R_{specimen}}$$

When the two resistance ratios are equal, the voltage drops across the respective resistances will also be equal. Kirchhoff’s Voltage Law declares that the voltage differential between two equal and opposite voltage drops must be zero, accounting for the meter’s indication of balance.

It would not be inappropriate to relate this to the operation of a laboratory balance-beam scale, comparing a specimen of unknown mass against a set of known masses. In either case, the instrument is merely comparing an unknown quantity against an (adjustable) known quantity, indicating a condition of equality between the two:




Many legacy instruments were designed around the concept of a *self-balancing* bridge circuit, where an electric servo motor drove a potentiometer to achieve a balanced condition against the voltage produced by some process sensor. Analog electronic paper chart recorders often used this principle. Almost all pneumatic process instruments use this principle to translate the force of a sensing element into a variable air pressure.

Modern bridge circuits are mostly used in laboratories for extremely precise component measurements. Very rarely will you encounter a Wheatstone bridge circuit used in the process industries.


4.9.2 Sensor signal conditioning

A different application for bridge circuits is to convert the output of an electrical sensor into a voltage signal representing some physical measurement. This is by far the most popular use of bridge measurement circuits in industry, and here we see the same circuit used in an entirely different manner from that of the balanced Wheatstone bridge circuit.




Here, the bridge will be balanced only when R_{sensor} is at one particular resistance value. Unlike the Wheatstone bridge, which serves to measure a component's value when the circuit is balanced, this bridge circuit will probably spend most of its life in an unbalanced condition. The output voltage changes as a function of sensor resistance, which makes that voltage a reflection of the sensor's physical condition. In the above circuit, we see that the output voltage increases (positive on the top wire, negative on the bottom wire) as the resistance of R_{sensor} increases.

One of the most common applications for this kind of bridge circuit is in strain measurement, where the mechanical strain of an object is converted into an electrical signal. The sensor used here is a device known as a *strain gauge*: a folded wire designed to stretch and compress with the object under test, altering its electrical resistance accordingly. Strain gauges are typically quite small, as shown by this photograph:



Strain gauges are useful when bonded to metal specimens, providing a means of electrically sensing the strain (“stretching” or “compressing” of that specimen). The following bridge circuit is a typical application for a strain gauge:



When the specimen is stretched along its long axis, the metal wires in the strain gauge stretch with it, increasing their length and decreasing their cross-sectional area, both of which work to

increase the wire's electrical resistance. This stretching is microscopic in scale, but the resistance change is measurable and repeatable within the specimen's elastic limit. In the above circuit example, stretching the specimen will cause the voltmeter to read upscale (as defined by the polarity marks). Compressing the specimen along its long axis has the opposite effect, decreasing the strain gauge resistance and driving the meter downscale.

Strain gauges are used to precisely measure the strain (stretching or compressing motion) of mechanical elements. One application for strain gauges is the measurement of strain on machinery components, such as the frame components of an automobile or airplane undergoing design development testing. Another application is in the measurement of force in a device called a *load cell*. A "load cell" is comprised of one or more strain gauges bonded to the surface of a metal structure having precisely known elastic properties. This metal structure will stretch and compress very precisely with applied force, as though it were an extremely stiff spring. The strain gauges bonded to this structure measure the strain, translating applied force into electrical resistance changes.


You can see what a load cell looks like in the following photograph:



Strain gauges are not the only dynamic element applicable to bridge circuits. In fact, any resistance-based sensor may be used in a bridge circuit to translate a physical measurement into an electrical (voltage) signal. Thermistors (changing resistance with temperature) and photocells (changing resistance with light exposure) are just two alternatives to strain gauges.

It should be noted that the amount of voltage output by this bridge circuit depends both on the amount of resistance change of the sensor *and* the value of the excitation source. This dependency on source voltage value is a major difference between a sensing bridge circuit and a Wheatstone (balanced) bridge circuit. In a perfectly balanced bridge, the excitation voltage is irrelevant: the output voltage is zero no matter what source voltage value you use. In an unbalanced bridge circuit, however, source voltage value matters! For this reason, these bridge circuits are often rated in terms of how many millivolts of output they produce *per volt of excitation* per unit of physical measurement (microns of strain, newtons of stress, etc.).

An interesting feature of a sensing bridge circuit is its ability to cancel out unwanted variables. In the case of a strain gauge, for example, mechanical strain is not the only variable affecting gauge resistance. Temperature also affects gauge resistance. Since we do not wish our strain gauge to also act as a thermometer (which would make measurements very uncertain – how would we differentiate the effects of changing temperature from the effects of changing strain?), we must find some way to nullify resistance changes due solely to temperature, such that our bridge circuit will respond only to changes in strain. The solution is to creatively use a “dummy” strain gauge as another arm of the bridge:




The “dummy” gauge is attached to the specimen in such a way that it maintains the same temperature as the active strain gauge, yet experiences no strain. Thus, any *difference* in gauge resistances must be due solely to specimen strain. The differential nature of the bridge circuit naturally translates the differential resistance of the two gauges into one voltage signal representing strain.

If thermistors are used instead of strain gauges, this circuit becomes a differential temperature sensor. Differential temperature sensing circuits are used in solar heating control systems, to detect when the solar collector is hotter than the room or heat storage mass being heated.

Sensing bridge circuits may have more than one active “arm” as well. The examples you have seen so far in this section have all been *quarter-active* bridge circuits. It is possible, however, to incorporate more than one sensor into the same bridge circuit. So long as the sensors’ resistance changes are coordinated, their combined effect will be to increase the sensitivity (and often the linearity as well) of the measurement.

For example, *full-active bridge* circuits are sometimes built out of four strain gauges, where each strain gauge comprises one arm of the bridge. Two of the strain gauges must compress and the other two must stretch under the application of the same mechanical force, in order that the bridge will become unbalanced with strain:




Not only does a full-active bridge circuit provide greater sensitivity and linearity than a quarter-active bridge, but it also *naturally* provides temperature compensation without the need for “dummy” strain gauges, since the resistances of all four strain gauges will change by the same proportion if the specimen temperature changes.

4.10 Null-balance voltage measurement

A number of different sensors used in instrumentation generate DC voltage signals proportional to the process variable of interest. We call such sensors *potentiometric*, which literally means “voltage-measuring”. Thermocouples are one type of potentiometric sensor, used to measure temperature. Photodiodes are another, used to measure light intensity. Glass pH electrodes are yet another, used to measure the hydrogen ion activity in a liquid solution. It should be obvious that accurate voltage measurement is critical for any instrument based on a potentiometric sensor, for if our measurement of that sensor’s output voltage is not accurate, we will surely suffer inaccurate measurement of any process variable proportional to that voltage (e.g. temperature, light, pH).


One common obstacle to accurate sensor voltage measurement is the internal resistance of the sensor itself. We will explore this concept by way of a practical example: trying to measure the voltage output by a pH electrode pair using a standard digital voltmeter. A pictorial diagram shows the basic concept, where a voltmeter is connected to a pH electrode pair immersed in a liquid solution:



Hydrogen ions within the liquid solution penetrate the round glass bulb of the measurement electrode, generating a potential difference approximately equal to 59 millivolts per pH unit of deviation from 7 pH (neutral). The reference electrode serves the simple purpose of completing the electrical circuit from the voltmeter’s terminals to both sides of the glass bulb (inside and outside).

What should be an elementary task is complicated by the fact that the glass bulb of the measurement electrode has an incredibly high electrical resistance, typically on the order of *hundreds* of mega-ohms. When connected to a common digital multimeter having an input resistance in the order of *tens* of mega-ohms, the voltmeter acts as a rather “heavy” electrical load which causes the measured voltage to be far less than what the glass electrode is actually producing.

If we sketch an equivalent electrical schematic of these components, the problem becomes more evident. Red arrows in this schematic depict the flow of electrical current (conventional notation):



Only a small fraction of the glass electrode's voltage (V_{pH}) will actually be seen at the voltmeter's terminals due to this loading effect. We may treat the voltmeter's internal resistance of $10\text{ M}\Omega$ as one resistance in a voltage divider network, the two probe resistances being the other two divider resistances:

$$V_{meter} = V_{pH} \left(\frac{10\text{ M}\Omega}{300\text{ M}\Omega + 10\text{ M}\Omega + 3\text{ k}\Omega} \right)$$

Supposing the pH-sensing glass bulb outputs 100 millivolts, the voltmeter in this circuit would only register a reading of 3.226 millivolts: just a *few percent* of the actual sensor's potentiometric output. While this is a rather extreme example, it should be clear to see that any potentiometric circuit of the same form will suffer some degree of measurement inaccuracy due to this effect – the only question being how much error.


Lying at the heart of this problem is the fact that voltmeters necessarily draw some electric current in the act of measuring a voltage. It is this current draw, no matter how slight, that causes a voltmeter to register something other than a perfect facsimile of the sensor's voltage signal. The solution to this problem, then, is to minimize or eliminate this current draw. In other words, we need our voltmeter to have as much internal resistance as possible (ideally, an *infinite* amount of internal resistance).

Modern field-effect transistor amplifier circuits go a long way toward addressing this problem by allowing us to manufacture voltmeters having internal resistances in the *trillions* of ohms. So long as the voltmeter's internal resistance far overshadows (i.e. “swamps”) the signal source's resistance, loading errors will be held to a minimum.

In the days before high-resistance semiconductor amplifier circuits, special voltmeters called *vacuum-tube voltmeters* (VTVMs) were used whenever voltages needed to be measured from high-resistance potentiometric sensors.

Prior to the advent of electronic vacuum tubes, though, it was impossible to mitigate the problem of voltmeter loading simply by selecting a better-quality voltmeter. Un-amplified voltmeters rely on the passage of a small current from the source under test to drive their indicating mechanisms. Without this small current drawn from the circuit, the voltmeter simply would not function at all. How then did early experimenters and electrical metrologists overcome this problem of loading?

An ingenious solution to the problem of voltmeter loading is the so-called *null-balance* method of voltage measurement. This technique added two components to the measurement circuit: a highly sensitive ammeter called a *galvanometer* and a variable DC voltage source. Using our pH measurement circuit as an example, a null-balance arrangement would look something like this:



Operation of this circuit follows these two steps:

- Adjust the variable DC voltage source until the galvanometer registers exactly zero (i.e. no current)
- Read the voltmeter to see what the pH sensor's voltage is

So long as the galvanometer registers zero (a “null” condition), there will be no electric current passing through the large resistances of the pH sensor's electrodes because the pH sensor's voltage is perfectly balanced against the variable supply's voltage. With no current passing through those high resistances, they will drop no voltage whatsoever. Thus, V_{pH} must be equal to the voltage of the variable DC source, which the voltmeter registers accurately because its current requirements are met by the variable source and not the pH sensor. The only way this measurement technique can fail in its objective is if the galvanometer is not able to precisely detect a condition of zero current. So long as the galvanometer faithfully tells us when we have reached a condition of zero current, we may measure the voltage of the pH sensor using *any* DC voltmeter regardless of its internal resistance.


Special null-balance voltmeter instruments were manufactured with precision variable voltage sources built into them, called *differential voltmeters*. One such instrument was the Fluke model 801, a gold-plated version of which is shown here in Fluke's museum of measurement:



Note the center-zero analog meter on the face of this instrument, performing the function of the sensitive galvanometer in our schematic diagram. A set of five knobs oriented vertically on the face of this instrument, each one showing one digit of a 5-digit number, adjusted the DC voltage output by the internal voltage source. When the “null” meter registered zero, it meant the voltage of the source or circuit under test was precisely equal to the voltage dialed up by these five knobs. Differential voltmeters such as the Fluke 801 used amplifier circuits to make this “null” detector ultra-sensitive, in order to achieve the most accurate condition of balance between the variable DC voltage source and the source under test possible.

4.11 Electromagnetism


The fundamental principle of *electromagnetism* is that an electric current will create a magnetic field at right angles to the direction of the current. If the electric current travels in a straight path, the lines of magnetic flux will form concentric circles around that path. If the electric current travels in a circular path (i.e. through a loop or coil of wire), the magnetic lines of flux will form straight lines down the center of the coil, wrapping around at the ends to form a complete loop of its own:



Magnetic field strength is directly proportional to the amount of current in the conductor (and also directly proportional to the number of “turns” in a coiled wire), such that the unit of measurement¹³ for magnetic field strength is the *amp-turn*.

¹³Both in the British system of measurement *and* the SI metric system of measurement! The older metric system (called “CGS” for Centimeter-Gram-Second) had a special unit of measurement called the *Gilbert* for expressing magnetic field strength, with 1 Gilbert (Gb) equal to 0.7958 Amp-turns (At).

The directly proportional relationship between current intensity and magnetic field strength is exploited by *clamp-on ammeters*, which are able to measure electric current without the need for contact with the conductor:



The following photograph shows how particles of steel align with the circular magnetic field surrounding one of the electric cables of an electric arc welder on a job site:




Very strong magnetic fields may be generated with wire *coils*, since the magnetic fields surrounding each “turn” of wire in a coil tend to overlap constructively, supporting one another to form a stronger total field. The magnetic field from a wire coil may be so strong, in fact, that it is useful for creating an attractive force upon a ferrous¹⁴ object (called an *armature*) strong enough to move mechanisms. This arrangement of a wire coil and an iron armature is called a *solenoid*¹⁵.

¹⁴The term “ferrous” simply refers to any substance containing the element *iron*. This includes steel, which is a combination of iron and carbon.


¹⁵The word “solenoid” may also be used to describe a wire coil with no armature, but the more common industrial use of the word refers to the complete arrangement of coil and movable armature.

A practical example of a solenoid is a *solenoid valve*: a mechanical valve opened and/or closed by the application of electric current through the coil:



When the coil is energized by an external source of electric current, the magnetic field attracts the movable armature, thereby actuating the valve. A spring typically returns the valve mechanism back to its original position upon de-energization of the coil.

Another very practical application of electromagnetism are *electric motors*, which use magnetic fields generated by wire coils to twist a mechanical shaft and perform useful work.




In the DC motor design, electric current creates a magnetic field in the armature coil which reacts against the magnetic fields of the permanent magnets to twist the armature about its axis. Every half-turn, the brushes break and re-make contact with the commutator bars, reversing the direction of current through the armature coil to keep it spinning in the same direction.

In the AC motor design, alternating current (continually reversing direction many times per second) creates an alternating magnetic field in a set of stationary wire coils called *stator* windings.

This alternating magnetic field causes a rotor to spin at speed proportional to the frequency of the magnetic field's reversals.

Both electric current and magnetic field lines are *vectors*, having both magnitude and direction. As we have seen already, there is a perpendicular relationship between these two vectors. This relationship may be visualized by a simple rule called the *right-hand rule*, whereby the fingers and thumb of a human right hand represent the vector orientations of current and magnetism (or vice-versa). Using the right-hand rule, digits representing current direction assume the use of *conventional flow* rather than electron flow¹⁶, while digits representing magnetism point in the direction of “North.”

To use this rule, curl the four fingers of your right hand such that they point toward the palm of that hand, and extend your right thumb so that it points perpendicularly to your curled fingers. Your right hand should look like this:





Your curled fingers will represent one set of vectors, while your thumb points in the direction of the other. Whether the fingers represent current and the thumb magnetism, or whether the fingers represent magnetism and the thumb current, is irrelevant: the rule works both ways.

¹⁶There is also a left-hand rule for fans of electron flow, but in this book I will default to conventional flow. For a more complete discussion on this matter, see section 4.2.1 beginning on page 301.

This flexibility makes the right-hand rule easy to apply to different situations such as these:


Thumb represents current vector
Fingers represent magnetic field vectors




Fingers represent current vectors
Thumb represents magnetic field vector

Note: red arrows point in the direction of conventional current flow

Physicists have devised a convention for clearly illustrating the directions of perpendicular vectors (arrows) without resorting to illustrations drawn in 3-dimensional perspective. The following two-dimensional illustration shows the magnetic field surrounding a straight, current-carrying wire. The magnetic field, of course, encircles the wire. This is shown by the alternating “dot” and “cross” symbols above and below the wire, representing arrow heads (circles with dots) coming “out” of the page directly at the reader, and representing arrow tails (circles with crosses) headed straight into the page away from the reader:



The same notation may be used to show the perpendicular relationship between current and magnetic flux lines for a coiled conductor. Here, the arrow “tips” and “tails” represent current (conventional flow) entering and exiting the page, while the horizontal arrow represents magnetic field direction:



Note how the individual magnetic fields surrounding each wire in the coil all have their arrows pointing to the right in the coil’s interior, and to the left at the coil’s exterior. This shows how the individual magnetic loops constructively add to generate a large magnetic field through the center of the coil, looping around back to the other end of the coil.


4.12 Capacitors

Any two electrical conductors separated by an insulating medium possess the characteristic called *capacitance*: the ability to store energy in the form of an electric field created by a voltage between those two conductors. Capacitance is symbolized by the capital letter C and is measured in the unit of the *Farad* (F). The relationship between capacitance, stored electric charge (Q), and voltage (V) is as follows:

$$Q = CV$$

For example, a capacitance having a value of 33 microfarads charged to a voltage of 5 volts would store an electric charge of 165 microcoulombs.

Capacitors are devices expressly designed and manufactured to possess capacitance. They are constructed of a “sandwich” of conductive plates separated by an insulating *dielectric*. Capacitors have maximum voltage ratings as well as capacitance ratings. Here are some schematic symbols for capacitors:



Although any two conductors separated from each other by an insulating medium possess capacitance, we may optimize the amount of capacitance by situating the conductors in a manner yielding more electric flux for any given amount of applied voltage. Placing the two conductors closer together (thus creating a more intense electric field between them, measured in *volts per inch*) is one way to increase capacitance. Increasing the overlapping area of the conductors (thus providing more area for parallel electric flux lines to form) is another way. Inserting a strongly dielectric (i.e. an exceptionally good insulating) material between the conductors is a third way to increase capacitance.

These three factors may be expressed in a single mathematical formula for capacitance:

$$C = \frac{\epsilon A}{d}$$

Where,


C = Capacitance in Farads

A = Overlapping area of conductors in square meters

ϵ = Permittivity of the dielectric substance

d = Distance separating the conductors in meters

Capacitance is a non-dissipative quantity. Unlike resistance, a pure capacitance does not dissipate energy in the form of heat; rather, it stores and releases energy from and to the rest of the circuit. We may illustrate the energy-storing behavior of a capacitor by this simple voltage source, switch, and resistor circuit:



Note how the capacitor alternately functions as a source and as a load, depending on what it's connected to. When connected to a source of voltage, the capacitor absorbs (stores) energy in the form of an electric field between its plates. Current flows through the voltage source in the same direction as though it were powering a load (e.g. a resistor). When the capacitor's voltage equals the source voltage, current stops in the circuit. Flipping the switch to the "discharge" position connects the capacitor to a resistor, where it discharges its store of energy, acting as a source to the resistor's load.


Note also how the capacitor maintains the same polarity of voltage during all three steps. Since energy is stored within an *electric field* inside the capacitor, and electric fields are a function of *voltage*, the stored energy manifests itself as a voltage with an unchanging polarity. The Law of Energy Conservation – which states energy cannot spontaneously vanish or appear but must be accounted for – tells us a capacitor's voltage cannot spontaneously vanish, appear, or change polarity. In other words, *voltage is the conserved quantity in a capacitor as it transitions between acting as a source and acting as a load*. Current through a capacitor, however, switches direction depending on whether the capacitor is charging (acting as a load) or discharging (acting as a source).

Capacitance adds when capacitors are connected in parallel. It diminishes when capacitors are connected in series:

$$C_{parallel} = C_1 + C_2 + \dots + C_n \qquad C_{series} = \frac{1}{\frac{1}{C_1} + \frac{1}{C_2} + \dots + \frac{1}{C_n}}$$

Recall that capacitance stores energy in the form of an electric field, as a function of the voltage applied to it. If we wish to increase capacitance by connecting multiple capacitors to each other, we need to do that in such a way that each capacitor receives the same (total) applied voltage so that each additional amount of capacitance included in the network will contribute a proportional amount of energy storage to the network. We know voltage is guaranteed to be equal only among *parallel*-connected components. If we were to connect multiple capacitors in series with one another, their individual voltages would be some fraction of the total voltage (series voltages always adding to equal the total voltage), thus diminishing the energy stored in each capacitor and similarly diminishing the total capacitance.

Another way to comprehend why capacitance increases in parallel and diminishes in series is to consider a network of capacitors as one equivalent capacitor in terms of aggregate plate area and separation distance. Examining the directions of change to the corresponding variables in the permittivity/area/distance formula for capacitance, we see how parallel-connected capacitances must add while series-connected capacitances must diminish:



The relationship between voltage and current for a capacitor is as follows:

$$I = C \frac{dV}{dt}$$

As such, capacitors oppose changes in voltage over time by passing a current. This behavior makes capacitors useful for stabilizing voltage in DC circuits. One way to think of a capacitor in a DC circuit is as a *temporary voltage source*, always “wanting” to maintain voltage across its terminals as a function of the energy stored within its electric field.

The amount of potential energy (E_p , in units of joules) stored by a capacitor may be determined by altering the voltage/current/capacitance equation to express power ($P = IV$) and then applying some calculus (recall that power is defined as the time-derivative of work or energy, $P = \frac{dW}{dt} = \frac{dE}{dt}$):

$$I = C \frac{dV}{dt}$$

$$P = IV = CV \frac{dV}{dt}$$

$$\frac{dE_p}{dt} = CV \frac{dV}{dt}$$

$$\frac{dE_p}{dt} dt = CV dV$$

$$\int \frac{dE_p}{dt} dt = \int CV dV$$

$$\int dE_p = C \int V dV$$

$$E_p = \frac{1}{2} CV^2$$

In an AC circuit, the amount of capacitive reactance (X_C) offered by a capacitor is inversely proportional to both capacitance and frequency:

$$X_C = \frac{1}{2\pi fC}$$

This means an AC signal finds it “easier” to pass through a capacitor (i.e. less ohms of reactance) at higher frequencies than at lower frequencies.

4.13 Inductors

Any conductor possesses a characteristic called *inductance*: the ability to store energy in the form of a magnetic field created by an electric current flowing through that conductor. Inductance is symbolized by the capital letter L and is measured in the unit of the *Henry* (H).

Inductors are devices expressly designed and manufactured to possess inductance. They are typically constructed of a wire coil wound around a ferromagnetic¹⁷ core material. Inductors have current ratings as well as inductance ratings. Due to the effect of *magnetic saturation*, inductance tends to decrease as current approaches the rated maximum value in an iron-core inductor. Here are some schematic symbols for inductors:




An inductor's inductance is related to the magnetic permeability of the core material (μ), the number of turns in the wire coil (N), the cross-sectional area of the coil (A), and the length of the coil (l):

$$L = \frac{\mu N^2 A}{l}$$

¹⁷The term "ferrous" refers to any substance containing the element *iron*. Steel is one such substance, being a combination of iron and carbon.

Inductance is a non-dissipative quantity. Unlike resistance, a pure inductance does not dissipate energy in the form of heat; rather, it stores and releases energy from and to the rest of the circuit. We may illustrate the energy-storing behavior of an ideal¹⁸ inductor by this simple current source, switch, and resistor circuit:



Note how the inductor alternately functions as a source and as a load, depending on what it's connected to. When connected to a source of current, the inductor absorbs (stores) energy in the form of a magnetic field within its core. Voltage is dropped across the current source in the same polarity as though it were powering a load (e.g. a resistor). When the inductor's current equals the rated current of the source, there will no longer be a voltage drop (assuming a perfect inductor with zero internal wire resistance). Flipping the switch to the "discharge" position connects the inductor to a resistor, where it discharges its store of energy, acting as a source to the resistor's load.

Note also how the inductor maintains the same direction of current. Since energy is stored within a *magnetic field* inside the inductor, and magnetic fields are a function of *current*, the stored energy manifests itself as a current with an unchanging direction. The Law of Energy Conservation – which states energy cannot spontaneously vanish or appear but must be accounted for – tells us an inductor's current cannot spontaneously vanish, appear, or change direction. In other words, *current is the conserved quantity in an inductor as it transitions between acting as a source and acting as a load*. Voltage across an inductor, however, switches polarity depending on whether the inductor is charging (acting as a load) or discharging (acting as a source).

¹⁸Unlike the charge/hold/discharge capacitor circuit, this inductor demonstration circuit would not function quite as well in real life. Real inductors contain substantial amounts of electrical resistance (R) in addition to inductance (L), which means real inductors have an inherent capacity to dissipate their own store of energy. If a real inductor were placed in a circuit such as this, it would not maintain its store of energy indefinitely in the switch's "neutral" position as a capacitor would. Realistically, the inductor's energy would likely dissipate in a matter of milliseconds following the switch to the "neutral" position.

Inductance adds when inductors are connected in series. It diminishes when inductors are connected in parallel:

$$L_{series} = L_1 + L_2 + \cdots + L_n \qquad L_{parallel} = \frac{1}{\frac{1}{L_1} + \frac{1}{L_2} + \cdots + \frac{1}{L_n}}$$

Recall that inductance stores energy in the form of a magnetic electric field, as a function of the electric current flowing through it. If we wish to increase inductance by connecting multiple inductors to each other, we need to do that in such a way that each inductor receives the same (total) current so that each additional amount of inductance included in the network will contribute a proportional amount of energy storage to the network. We know current is guaranteed to be equal only among *series*-connected components. If we were to connect multiple inductors in parallel with one another, their individual currents would be some fraction of the total current (parallel currents always adding to equal the total current), thus diminishing the energy stored in each inductor and similarly diminishing the total inductance.

The relationship between voltage and current for an inductor is as follows:

$$V = L \frac{dI}{dt}$$

As such, inductors oppose changes in current over time by dropping a voltage. This behavior makes inductors useful for stabilizing current in DC circuits. One way to think of an inductor in a DC circuit is as a *temporary current source*, always “wanting” to maintain current through its coil as a function of the energy stored within its magnetic field.

The amount of potential energy (E_p , in units of joules) stored by an inductor may be determined by altering the voltage/current/inductance equation to express power ($P = IV$) and then applying some calculus (recall that power is defined as the time-derivative of work or energy, $P = \frac{dW}{dt} = \frac{dE}{dt}$):

$$V = L \frac{dI}{dt}$$

$$P = IV = LI \frac{dI}{dt}$$

$$\frac{dE_p}{dt} = LI \frac{dI}{dt}$$

$$\frac{dE_p}{dt} dt = LI dI$$

$$\int \frac{dE_p}{dt} dt = \int LI dI$$

$$\int dE_p = L \int I dI$$

$$E_p = \frac{1}{2} LI^2$$

In an AC circuit, the amount of inductive reactance (X_L) offered by an inductor is directly proportional to both inductance and frequency:

$$X_L = 2\pi fL$$

This means an AC signal finds it “harder” to pass through an inductor (i.e. more ohms of reactance) at higher frequencies than at lower frequencies.

References

Boylestad, Robert L., *Introductory Circuit Analysis*, 9th Edition, Prentice Hall, Upper Saddle River, NJ, 2000.

Chapter 5

AC electricity

While *direct current* (DC) refers to the flow of electrical charge carriers in a continuous direction, *alternating current* (or *AC*) refers to a periodic reversal of charge flow direction¹. As a mode of transferring electrical power, AC is tremendously useful because it allows us to use *transformers* to easily and efficiently step voltage up or down at will. If an electro-physical sensor senses a physical quantity that oscillates, the electric signal it produces will oscillate (AC) as well. For both these reasons, an instrument technician needs to be aware of how AC circuits work, and how to understand them mathematically.


¹It is also acceptable to refer to electrical voltages and/or currents that vary periodically over time even if their directions never alternate, as AC *superimposed* on DC.

5.1 RMS quantities

It is often useful to express the amplitude of an alternating quantity such as voltage or current in terms that are equivalent to direct current (DC). Doing so provides an “apples-to-apples” comparison between AC and DC quantities that makes comparative circuit analysis much easier.

The most popular standard of equivalence is based on *work* and *power*, and we call this the *root-mean-square* value of an AC waveform, or RMS for short. For example, an AC voltage of 120 volts “RMS” means that this AC voltage is capable of delivering the exact same amount of power (in Watts) at an electrical load as a 120 volt DC source powering the exact same load.

The problem is exactly how to calculate this “RMS” value if all we know about the AC waveform is its peak value. If we compare a sine wave and a DC “wave” side by side, it is clear that the sine wave must peak at a greater value than the constant DC level in order to be equivalent in terms of doing the same *work* in the same amount of time:



At first, it might seem like the correct approach would be to use calculus to integrate (determine the area enclosed by) the sine wave over one-half of a cycle: from 0 to π radians. This is close, but not fully correct. The ability of an electrical voltage to dissipate power at a resistor is not directly proportional to the magnitude of that voltage, but rather proportional to the *square* of the magnitude of that voltage! In mathematical terms, resistive power dissipation is predicted by the following equation:

$$P = \frac{V^2}{R}$$

If we double the amount of voltage applied to a resistor, the power increases four-fold. If we triple the voltage, the power goes up by a factor of nine. If we are to calculate the “RMS” equivalent value of a sine wave, we must take this nonlinearity into consideration.

First let us begin with a mathematical equivalence between the DC and AC cases. For DC, the amount of *work* done is equal to the constant power of that circuit multiplied by time. The unit of measurement for power is the *Watt*, which is defined as 1 Joule of work per second. So, multiplying the steady power rate in a DC circuit by the time we keep it powered will result in an answer of joules (total energy dissipated by the resistor):

$$\text{Work} = \left(\frac{V^2}{R} \right) t$$

Showing this equivalence by dimensional analysis:

$$[\text{Joules}] = \left(\frac{[\text{Joules}]}{[\text{s}]} \right) [\text{s}]$$

We cannot calculate the work done by the AC voltage source quite so easily, because the power dissipation varies over time as the instantaneous voltage rises and falls. Work is still the product of power and time, but we cannot simply multiply one by the other because the voltage in this case is a *function* of time ($V(t)$). Instead, we must use integration to calculate the product of power and time, and sum those work quantities into a total work value.

Since we know the voltage provided by the AC source takes the form of a sine wave ($V(t) = \sin t$ if we assume a sine wave with a peak value of 1 volt), we may write the formula for instantaneous AC power as follows:

$$\text{Power} = \frac{(V(t))^2}{R} = \frac{\sin^2 t}{R}$$

To calculate the work done by this sinusoidal voltage on the resistor, we will integrate this instantaneous power with respect to time, between the intervals of $t = 0$ and $t = \pi$ (one half-cycle of the sine wave):

$$\text{Work} = \int_0^\pi \frac{\sin^2 t}{R} dt$$

In order to solve for the amount of DC voltage equivalent (from the perspective of resistive power dissipation) to a one-volt AC sine wave, we will set the DC work and AC work equations equal to each other, making sure the DC side of the equation has π for the amount of time (being the same time interval as the AC side):

$$\left(\frac{V^2}{R} \right) \pi = \int_0^\pi \frac{\sin^2 t}{R} dt$$

Our goal is to solve for V on the left-hand side of this equation.

First, we know that R is a constant value, and so we may move it out of the integrand:

$$\left(\frac{V^2}{R}\right)\pi = \frac{1}{R} \int_0^\pi \sin^2 t \, dt$$

Multiplying both sides of the equation by R eliminates it completely. This should make intuitive sense, as our RMS equivalent value for an AC voltage is defined strictly by the ability to produce the same amount of power as the same value of DC voltage for *any* resistance value. Therefore the actual value of resistance (R) should not matter, and it should come as no surprise that it cancels:

$$V^2\pi = \int_0^\pi \sin^2 t \, dt$$

Now, we may simplify the integrand by substituting the half-angle equivalence for the $\sin^2 t$ function

$$V^2\pi = \int_0^\pi \frac{1 - \cos 2t}{2} \, dt$$

Factoring one-half out of the integrand and moving it outside (because it's a constant):

$$V^2\pi = \frac{1}{2} \int_0^\pi 1 - \cos 2t \, dt$$

We may write this as the difference between two integrals, treating each term in the integrand as its own integration problem:

$$V^2\pi = \frac{1}{2} \left(\int_0^\pi 1 \, dt - \int_0^\pi \cos 2t \, dt \right)$$

The second integral may be solved simply by using substitution, with $u = 2t$, $du = 2 \, dt$, and $dt = \frac{du}{2}$:

$$V^2\pi = \frac{1}{2} \left(\int_0^\pi 1 \, dt - \int_0^\pi \frac{\cos u}{2} \, du \right)$$

Moving the one-half outside the second integrand:

$$V^2\pi = \frac{1}{2} \left(\int_0^\pi 1 \, dt - \frac{1}{2} \int_0^\pi \cos u \, du \right)$$

Finally we are at a point where we may perform the integrations:

$$V^2\pi = \frac{1}{2} \left(\int_0^\pi 1 dt - \frac{1}{2} \int_0^\pi \cos u du \right)$$

$$V^2\pi = \frac{1}{2} \left([t]_0^\pi - \frac{1}{2} [\sin 2t]_0^\pi \right)$$

$$V^2\pi = \frac{1}{2} \left([\pi - 0] - \frac{1}{2} [\sin 2\pi - \sin 0] \right)$$

$$V^2\pi = \frac{1}{2} \left([\pi - 0] - \frac{1}{2} [0 - 0] \right)$$

$$V^2\pi = \frac{1}{2}(\pi - 0)$$

$$V^2\pi = \frac{1}{2}\pi$$

We can see that π cancels out of both sides:

$$V^2 = \frac{1}{2}$$

Taking the square root of both sides, we arrive at our final answer for the equivalent DC voltage value:

$$V = \frac{1}{\sqrt{2}}$$

So, for a sinusoidal voltage with a peak value of 1 volt, the DC equivalent or “RMS” voltage value would be $\frac{1}{\sqrt{2}}$ volts, or approximately 0.707 volts. In other words, a sinusoidal voltage of 1 volt peak will produce just as much power dissipation at a resistor as a steady DC voltage of 0.7071 volts applied to that same resistor. Therefore, this 1 volt peak sine wave may be properly called a 0.7071 volt RMS sine wave, or a 0.7071 volt “DC equivalent” sine wave.

This factor for sinusoidal voltages is quite useful in electrical power system calculations, where the wave-shape of the voltage is nearly always sinusoidal (or very close). In your home, for example, the voltage available at any wall receptacle is 120 volts RMS, which translates to 169.7 volts peak.

Electricians and electronics technicians often memorize the $\frac{1}{\sqrt{2}}$ conversion factor without realizing it only applies to *sinusoidal* voltage and current waveforms. If we are dealing with a non-sinusoidal wave-shape, the conversion factor between peak and RMS *will* be different! The mathematical procedure for obtaining the conversion factor will be identical, though: integrate the wave-shape’s function (squared) over an interval sufficiently long to capture the essence of the shape, and set that equal to V^2 times that same interval span.

5.2 Resistance, Reactance, and Impedance

Resistance (R) is the dissipative opposition to an electric current, analogous to friction encountered by a moving object. In any example of electrical resistance, the electrical energy is converted into some other form of energy that cannot (or does not) return back to the circuit. Resistance may take the form of an actual resistor, in which case the electrical energy is converted into heat. Resistance may also take the form of an electric motor, an electric light, or an electrochemical cell where the electrical energy is converted into mechanical work, photons, or enables an endothermic chemical reaction, respectively.

Reactance (X) is the opposition to an electric current resulting from energy storage and release between certain components and the rest of the circuit, analogous to inertia of a moving object. Capacitors and inductors are classic examples of “reactive” electrical components, behaving either as electrical loads or as electrical sources depending on whether the applied electrical signal is increasing or decreasing in intensity at that instant in time. When a purely reactive component is subjected to a sinusoidal signal, it will spend exactly half the time behaving as a load (absorbing energy from the circuit) and half the time behaving as a source (returning energy to the circuit). Thus, a purely reactive component neither contributes nor dissipates any net energy in the circuit, but merely exchanges² energy back and forth. Even though the fundamental mechanism of reactance (energy storage and release) is different from the fundamental mechanism of resistance (energy conversion and dissipation), reactance and resistance are both expressed in the same unit of measurement: the ohm (Ω).

Impedance (Z) is the combined total opposition to an electric current, usually some combination of electrical resistance (energy dissipation) and electrical reactance (energy storage and release). It is also expressed in the unit of the ohm. In order to represent how much of a particular impedance is due to resistance and how much is due to reactance, the value of an impedance may be expressed as a *complex number* with a “real” part (representing resistance) and an “imaginary” part (representing reactance). This concept will be explored in much more detail later in this chapter.

The amount of electrical reactance offered by a capacitor or an inductor depends on the *frequency* of the applied signal. The faster the rate at which an AC signal oscillates back and forth, the more a reactive component tends to react to that signal. The formulae for capacitive reactance (X_C) and inductive reactance (X_L) are as follows:

$$X_C = \frac{1}{2\pi fC} \qquad X_L = 2\pi fL$$

Just as conductance (G) is the reciprocal of resistance ($1/R$), a quantity called *susceptance* (B) is the reciprocal of reactance ($1/X$). Susceptance is useful when analyzing parallel-connected reactive components while reactance is useful for analyzing series-connected reactive components, in much the same way that conductance and resistance are useful when analyzing parallel-connected and series-connected resistors, respectively.

Impedance (Z) also has a reciprocal counterpart known as *admittance* (Y).


²Charles Proteus Steinmetz, in his book *Theoretical Elements of Electrical Engineering*, refers to the voltage and current values of a reactive component being “wattless” in honor of the fact that they transfer zero net power to or from the circuit (page 41). The voltage and current values of resistive components, by contrast, constitute real power dissipated in the circuit.

5.3 Series and parallel circuits

Impedance in a series circuit is the orthogonal sum of resistance and reactance:

$$Z = \sqrt{R^2 + (X_L^2 - X_C^2)}$$

Equivalent series and parallel circuits are circuits that have the exact same total impedance as one another, one with series-connected resistance and reactance, and the other with parallel-connected resistance and reactance. The resistance and reactance values of equivalent series and parallel circuits may be expressed in terms of those circuits' total impedance:



If the total impedance of one circuit (either series or parallel) is known, the component values of the equivalent circuit may be found by algebraically manipulating these equations and solving for the desired R and X values:

$$Z^2 = R_{series}R_{parallel}$$


$$Z^2 = X_{series}X_{parallel}$$

5.4 Transformers

The *transformer* is one of the most important components in all of AC circuitry. Principally used to “step” between different values of AC voltage and current in power systems, transformers find uses in many other types of circuits including electronic amplifiers (for impedance matching) and even sensor circuits (sensing physical position).

5.4.1 Basic principles

Before exploring the operation of a transformer, it is useful to review the operation of a simple inductor, which is nothing more than a coil of wire usually wrapped around a ferromagnetic core material:



If we apply an alternating (AC) voltage to this coil, it will generate an alternating magnetic field in the core. Just how much magnetic flux (ϕ) will develop in the core depends on how much voltage we apply to the coil. The fundamental relationship between voltage and magnetic flux for any conductive coil is given by Faraday's Law of Electromagnetic Induction³:

$$V = N \frac{d\phi}{dt}$$

Where,

V = Voltage applied to the coil or induced by the coil (volts)

N = Number of turns of wire

$\frac{d\phi}{dt}$ = Rate of change of magnetic flux (Webers per second)

³At first it may seem strange to apply Faraday's Law here, because this formula is typically used to describe the amount of voltage *produced* by a coil of wire exposed to a changing magnetic field, not the amount of magnetic field produced by an applied voltage. However, the two are closely related because the inductor must produce a voltage drop in equilibrium with the applied voltage just like any other component, in accordance with Kirchhoff's Voltage Law. In a simple circuit such as this where the voltage source directly connects to the inductor (barring any resistive losses in the connecting wires), the coil's induced voltage drop must exactly equal the source's applied voltage at all points in time, and so Faraday's Law works just as well to describe the source's applied voltage as it does to describe the coil's induced voltage. This is the principle of *self-induction*.

If the applied voltage is sinusoidal (i.e. shaped like a sine wave), then the magnetic flux magnitude will trace a cosine wave over time. We may demonstrate this mathematically by substituting $\sin \omega t$ (the sine of some frequency ω at any particular point in time t) for V in Faraday's equation and integrating:

$$V = N \frac{d\phi}{dt}$$

$$\sin \omega t = N \frac{d\phi}{dt}$$

$$\sin \omega t dt = N d\phi$$

$$\int \sin \omega t dt = \int N d\phi$$

$$\int \sin \omega t dt = N \int d\phi$$


$$-\frac{1}{\omega} \cos \omega t + \phi_0 = N\phi$$

$$\phi = -\frac{1}{N\omega} \cos \omega t + \phi_0$$

Thus, the amount of magnetic flux (ϕ) in the core at any point in time t is proportional to the cosine of the frequency-time function ωt plus any residual magnetism (ϕ_0) the core happened to start out with before any voltage was applied to the coil.

The amount of current drawn by this inductor depends on the reluctance of the core's magnetic "circuit" and the number of turns in the coil (N). The less reluctance offered by the magnetic path, the less current will be necessary to generate the requisite magnetic field to balance the applied voltage. If we were to take two perfect inductors (i.e. lacking wire resistance) – one with a heavy ferrous core and one with a light ferrous core (or even an air core) – and apply the same AC voltage to them, they would both generate the exact same strength of alternating magnetic field, but the inductor with the lesser core would draw more current from the source in doing so. In other words, the latter inductor would exhibit less reactance (i.e. fewer ohms) to oppose current.

Things get interesting if we wrap a second coil of wire around the same core as the first. For the sake of analysis we will label voltage polarities at one of the peaks of the AC source:




At that moment in time when the top terminal of the source is positive and the bottom terminal is negative, we see that the first coil drops the same voltage (due to self-induction), and that the second coil drops the same voltage as well (due to *mutual* induction). The polarity of both coils' voltages are identical because they are wrapped in the same direction around the core and they both experience the same magnetic flux (ϕ). When we examine the directions of current through each coil, however, we see they are opposite one another: the left-hand coil acts as a *load* (drawing energy from the AC voltage source) while the right-hand coil acts as a *source* (providing energy to the resistive load).

What we have created here is a true *transformer*: an electromagnetic component transferring energy from electric form to magnetic form and back again to electric form. The AC voltage source is able to energize the resistive load without direct conductive connection between the two, since the magnetic flux serves as the energy "link" between the two circuits.


Transformers are typically drawn as a set of coils sharing a common core. The coil connected to an electrical source is called the *primary*, while the coil connected to an electrical load is called the *secondary*. If the core is ferromagnetic, it is shown as a set of parallel lines between the coils:

A simple transformer circuit



5.4.2 Loading effects

We may explore transformer behavior by observing the effects of powering one with a constant⁴-voltage AC source and varying the load resistance:




Observe how voltage at both coils is unaffected by load, and similarly how the magnetic flux remains unchanged under different load conditions. The secondary coil acts like a voltage source to the resistive load, reflecting the nature of the primary coil's source behavior. The magnetic flux amplitude is unaffected by secondary loading in order to satisfy Kirchhoff's Voltage Law and Faraday's Law at the primary coil: the coil's voltage drop must be equal and opposite to the source's

⁴In this context, "constant" means an alternating voltage with a consistent peak value, not "constant" in the sense that a DC source is constant at all points in time.

applied voltage, and so the magnetic flux must alternate at the same rates and reach the same peaks so long as the primary source voltage does the same.

Continuing our exploration of transformer behavior, we will now power one with a constant⁵-current AC source and vary the load resistance:



Observe how current now is the unaffected quantity, while voltage and magnetic flux are load-dependent. The secondary coil now acts like a current source to the resistive load, reflecting the nature of the primary coil's source behavior. As load resistance varies, the secondary coil's voltage varies proportionately, which in turn demands a commensurate change in magnetic flux.

⁵In this context, "constant" means an alternating voltage with a consistent peak value, not "constant" in the sense that a DC source is constant at all points in time.

5.4.3 Step ratios

Transformers are principally used to step between different levels of voltage and current. This is achieved by building the transformer with primary and secondary coils having different numbers of turns. Since both coils share the same magnetic flux, the number of turns will be proportionate to how much voltage is developed at each coil. We may prove this mathematically with Faraday's Law, using $\frac{d\phi}{dt}$ as the quantity shared between primary and secondary coils:


$$V_P = N_P \frac{d\phi}{dt} \qquad V_S = N_S \frac{d\phi}{dt}$$

$$\frac{V_P}{N_P} = \frac{d\phi}{dt} \qquad \frac{V_S}{N_S} = \frac{d\phi}{dt}$$

$$\frac{V_P}{N_P} = \frac{V_S}{N_S}$$



$$\frac{V_P}{V_S} = \frac{N_P}{N_S}$$

That is to say, the ratio of primary to secondary voltage is the same as the ratio of primary to secondary turns. We may exploit this principle to build transformers delivering the same amount of power to two different load resistances from the same power source, the only difference being the number of turns in the secondary coil:



Whichever way a transformer steps voltage from primary to secondary, it must step current the other way.

Here are some quantitative examples, assuming lossless transformers:



Note how primary and secondary powers are always equal to each other for any given transformer arrangement. Real transformers suffer some internal⁶ power loss, and as such will exhibit secondary power levels slightly less than primary, but assuming equality provides an easy way to check our voltage and current ratio calculations.


⁶These power losses take the form of *core losses* due to magnetic hysteresis in the ferrous core material, and *winding losses* due to electrical resistance in the wire coils. Core losses may be minimized by reducing magnetic flux density (H), which requires a core with a larger cross-section to disperse the flux (ϕ) over a wider area. Winding losses may be minimized by increasing wire gauge (i.e. thicker wire coils). In either case, these modifications make for a bulkier and more expensive transformer.

5.4.4 Transformer impedance

An ideal transformer is completely lossless, conveying electrical power from a connected source (on the primary side) to a connected load (on the secondary side) with 100 percent efficiency. Ideal transformers also pose no limit on the amount of power they may couple from primary to secondary winding – in other words, an ideal transformer imposes no inherent limit to power throughput.

Real transformers, however, are not lossless and in fact do act as current-limiting devices. The mechanisms for this include magnetic hysteresis losses, wire resistance, leakage inductance⁷, etc.

Consider a thought experiment where we short-circuit the secondary winding of an ideal transformer, which is being powered by an AC voltage source of infinite power capacity (i.e. the source has zero impedance). How much current would pass through the shorted secondary circuit?




This question has no realistic answer. If the 480 VAC source has no current limitation (i.e. is capable of supplying infinite current to a shorted load) and the transformer likewise presents no limit at all to current, the shorted secondary circuit would also experience infinite current, at least in principle.

It should be rather obvious that this scenario cannot exist in the real world. Even with a source of infinite current capability, any realistic transformer would act to impede current delivered to a short-circuit on the secondary side. The question of “how much current would pass through the short-circuit” is really a question of how much *impedance* the transformer offers.

⁷Transformers, of course, utilize the principle of electromagnetic induction to generate a voltage at the secondary winding which may power a load. Ideally, 100 percent of the magnetic flux generated by the energized primary winding “links” or “couples” to the secondary winding. However, imperfections in the windings, core material, etc. conspire to prevent every bit of magnetic flux from coupling with the secondary winding, and so any magnetic flux from the primary winding that *doesn't* transfer power to the secondary winding simply absorbs and releases energy like a plain inductor. This is called “leakage” inductance because the flux in question has found a path to “leak” around the secondary winding. Leakage inductance may be modeled in a transformer as a separate series-connected inductance connected to the primary winding. Like any inductance, it presents a reactance equal to $X_L = 2\pi fL$, and in a transformer serves to impede primary current.

Let us consider a different thought experiment, this time using a real transformer with a short-circuited secondary winding, powered by a variable AC voltage source:



Imagine gradually increasing the source voltage until the secondary circuit ammeter registers a current equal to the transformer's full-load rating. For an ideal transformer (perfect power coupling), this would happen at some very small amount of voltage applied to the primary winding. Due to the imperfections and losses of real transformers, though, full secondary current will be obtained at a primary voltage equal to some small percentage of the normal (rated) primary voltage. Suppose, for example, our hypothetical transformer with a primary winding rating of 480 VAC outputs full secondary current through a short-circuit at an applied source voltage of only 22 volts. 22 volts is 4.58% of 480 volts, and so we would say this transformer has a measured impedance of 4.58 percent⁸.

Although a short-circuited secondary winding scenario may seem contrived, it actually is quite relevant to real-world conditions. In electrical power systems we are often concerned with the maximum amount of current which will flow during *fault* conditions. If two power conductors directly touch each other, or if a low-resistance arc develops between them through the air, the effect is very nearly a perfect short-circuit. This means transformer impedance will be the dominant factor in limiting fault current: the more impedance a transformer has, the less fault current will manifest during shorted conditions.

One way to apply the impedance percentage value for a power transformer to a fault scenario is to use it as a multiplying factor for secondary current. For example, if a power transformer has a maximum rated secondary current of 180 amps and an impedance rating of 3.3%, the available secondary current into a bolted⁹ fault will be:

$$\frac{180 \text{ A}}{3.3\%} = 5454.5 \text{ A}$$

Bolted-fault current calculations are very useful when predicting the amount of energy released in an *arc blast* incident, which is what happens when an electric arc develops between two closely-spaced conductors in a high-power electric power system. The arc behaves as an extremely low-resistance

⁸Although it is possible to express transformer impedance in the more familiar unit of *Ohms* (Ω), percentage is greatly preferred for the simple reason that it applies identically to the primary and secondary sides of the transformer. Expressing transformer impedance in ohms would require a different value depending on whether the primary side or secondary side were being considered.

⁹The rather colorful term “bolted” refers to a short-circuit fault consisting of a large copper bus-bar physically attached to the transformer's secondary terminal using bolts. In other words, a “bolted” fault is as close to a perfect short-circuit as you can get.

connection between the conductors, resulting in very large current values with correspondingly high arc temperatures.

Transformer impedance is also useful for calculating the degree to which the output voltage of a power transformer will “sag” below its ideal value when powering a load. Suppose we had a power transformer with a 5:1 turns ratio, designed to receive 120 VAC at its primary winding and output 24 VAC. Under no-load conditions the transformer’s internal impedance will be of no effect, and the transformer will output 24 VAC exactly. However, when a load is connected to the secondary terminals and current begins to flow to power this load, the transformer’s internal impedance will result in the secondary voltage decreasing by a small amount. For example, if this transformer happens to have an impedance of 5.5%, it means the secondary (output) voltage will sag 5.5% below 24 VAC at full load, assuming the primary voltage is maintained at the standard 120 VAC level. 5.5% of 24 volts is 1.32 volts, and so this transformer’s secondary voltage will “sag” from 24 volts down to 22.68 volts (i.e. 1.32 volts less than 24 volts) as load current increases from zero to its full rated value.

5.5 Phasors


Phasors are to AC circuit quantities as *polarity* is to DC circuit quantities: a way to express the “directions” of voltage and current waveforms. As such, it is difficult to analyze AC circuits in depth without using this form of mathematical expression. Phasors are based on the concept of *complex numbers*: combinations of “real” and “imaginary” quantities. The purpose of this section is to explore how complex numbers relate to sinusoidal waveforms, and show some of the mathematical symmetry and beauty of this approach.

Since waveforms are not limited to alternating current electrical circuits, phasors have applications reaching far beyond the scope of this chapter.

5.5.1 Circles, sine waves, and cosine waves

Something every beginning trigonometry student learns (or *should* learn) is how sine and cosine waves may be derived from a circle. First, sketch a circle, then sketch a set of radius vectors from the circle's center to the circle's perimeter at regular angle intervals. Mark each point of intersection between a vector and the circle's perimeter with a dot and label each with the vector angle. Sketch rectangular graphs to the right and below the circle, with regularly-spaced divisions. Label those divisions with the same angles as the vectors and then sketch dashed "projection" lines from each vector tip to the respective divisions on the rectangular graphs, marking each intersection with a dot. Connect the dots with curves to reveal sinusoidal waveshapes.

The following illustration shows the dots and projection lines for the first five vectors (angles 0 through $\frac{\pi}{2}$ radians¹⁰) only. As you can see, the circle's vertical projection forms a *sine* wave, while the circle's horizontal projection forms a *cosine* wave:



¹⁰A full circle contains 360 degrees, which is equal to 2π radians. One "radian" is defined as the angle encompassing an arc-segment of a circle's circumference equal in length to its radius, hence the name "radian". Since the circumference of a circle is 2π times as long as its radius, there are 2π radians' worth of rotation in a circle. Thus, while the "degree" is an arbitrary unit of angle measurement, the "radian" is a more natural unit of measurement because it is defined by the circle's own radius.

The Swiss mathematician Leonhard Euler (1707-1783) developed a symbolic equivalence between polar (circular) plots, sine waves, and cosine waves by plotting the circle on a *complex plane* where the vertical axis is an “imaginary”¹¹ number line and the horizontal axis is a “real” number line. *Euler’s Relation* expresses the vertical (imaginary) and horizontal (real) projections of an imaginary exponential function as a complex (real + imaginary) trigonometric function:

$$e^{j\theta} = \cos \theta + j \sin \theta$$

Where,

e = Euler’s number (approximately equal to 2.718281828)


θ = Angle of vector, in radians

$\cos \theta$ = Horizontal projection of a unit vector (along a real number line) at angle θ

j = Imaginary “operator” equal to $\sqrt{-1}$, represented by i or j

$j \sin \theta$ = Vertical projection of a unit vector (along an imaginary number line) at angle θ


To illustrate, we will apply Euler’s relation to a unit¹² vector having an angular displacement of $\frac{3\pi}{8}$ radians:



¹¹The definition of an imaginary number is the square root of a negative quantity. $\sqrt{-1}$ is the simplest case, and is symbolized by mathematicians as i and by electrical engineers as j .

¹²The term “unit vector” simply refers to a vector with a length of 1 (“unity”).

This mathematical translation from circles to sinusoidal waves finds practical application in AC electrical systems, because the rotating vector directly relates to the rotating magnetic field of an AC generator while the sinusoidal function directly relates to the voltage generated by the field's motion. The principle of an AC generator is that a magnet rotates on a shaft past stationary coils of wire. When these wire coils experience the changing magnetic field produced by the rotating magnet, a sinusoidal voltage is induced:



If we mark the shaft of this simple two-pole generator such that an angle of zero is defined as the position where the stator coils develop maximum positive voltage (V_0), the AC voltage waveform will follow the cosine function as shown above. If we were to add another set of stationary (“stator”) coils to the generator’s frame perpendicular to this set, that second set of coils would generate an AC voltage following the *sine* function. As it is, there is no second set of stator coils, and so this sine-function voltage is purely an imaginary quantity whereas the cosine-function voltage is the only “real” electricity output by the generator.

The canonical form of Euler’s Relation assumes a circle with a radius of 1. In order to realistically represent the output voltage of a generator, we may include a multiplying coefficient in Euler’s Relation representing the function’s peak value:

$$Ae^{j\theta} = A \cos \theta + jA \sin \theta$$

In the case of the generator shown above, the A coefficient is the peak¹³ voltage (V_0) output by the stator coils at the precise moment in time when the shaft angle is zero. In other words, the exponential term ($e^{j\theta}$) tells us which way the vector points while the coefficient (A) tells us how long the vector is.

¹³Although A truly should represent a waveform’s peak value, and θ should be expressed in units of radians to be mathematically correct, it is more common in electrical engineering to express A in RMS (root-mean-square) units and θ in degrees. For example, a 120 volt RMS sine wave voltage at a phase angle of 30 degrees will be written by an engineer as $120e^{j30}$ even though the true phase angle of this voltage is $\frac{\pi}{6}$ radians and the actual peak value is 169.7 volts.

In electrical studies it is commonplace to use various “shorthand” notations for vectors. An instantaneous quantity (e.g. voltage or current at some moment in time) described simply in terms of real and imaginary values is called *rectangular form*, for example $0.3827 + j0.9239$ volts. An instantaneous quantity described simply in terms of vector angle and length is called *polar form*, for example 1 volt $\angle 67.5^\circ$. Euler’s Relation is the mathematical link connecting rectangular and polar forms.

	Polar	Rectangular
Shorthand notation	$A\angle\theta$	$x + jy$
Full notation	$Ae^{j\theta}$	$A \cos \theta + jA \sin \theta$

All of these vectors described by the imaginary exponential function $e^{j\theta}$ are of a special kind – vectors inhabiting a space defined by one real axis and one imaginary axis (the so-called *complex plane*). In order to avoid confusing this special quantity with real vectors used to represent quantities in actual space where every axis is real, we use the term *phasor* instead. From this point on in the book, I will use the term “phasor” to describe complex quantities as represented in Euler’s Relation, and “vector” to describe any other quantity possessing both a magnitude and a direction.

Euler's Relation may also be used to describe the behavior of the rotating generator shaft and its corresponding AC voltage output as a function of *time*. Instead of specifying a shaft position (angle θ) we now specify a point in time (t) and an angular velocity (rotational speed) represented by the lower-case Greek letter "omega" (ω). Since angular velocity is customarily given in units of radians per second, and time is customarily specified in seconds, the product of those two will be an angle in radians ($\omega t = \theta$). Substituting ωt for θ in Euler's Relation gives us a function describing the state of the generator's output voltage in terms of time:

$$Ae^{j\omega t} = A \cos \omega t + jA \sin \omega t$$

Where,

A = Peak amplitude of generator voltage, in volts

e = Euler's number (approximately equal to 2.718281828)

ω = Angular velocity, in radians per second

t = Time, in seconds

$\cos \omega t$ = Horizontal projection of phasor (along a real number line) at time t

j = Imaginary "operator" equal to $\sqrt{-1}$, represented by i or j

$j \sin \omega t$ = Vertical projection of phasor (along an imaginary number line) at time t

With just a little bit of imagination we may visualize both the exponential and complex sides of Euler's Relation in physical form. The exponential ($e^{j\omega t}$) is a phasor rotating counter-clockwise about a central point over time. The complex ($\cos \omega t + j \sin \omega t$) is a pair of cosine and sine waves oscillating along an axis of time. Both of these manifestations may be joined into one visual image if you imagine a corkscrew where the centerline of the corkscrew is the time axis, and the evolution of this function over time follows the curve of the corkscrew down its length. Viewed from one end (looking along the centerline axis), the path appears to be a circle, just as a rotating phasor traces a circle with its tip. Viewed perpendicular to that axis, the path appears to trace a sinusoid, just as the sine or cosine function goes up and down as it progresses from left to right on a rectangular graph. A compression-type coil spring serves just as well as a corkscrew for a visual aid: viewed from the end the spring looks like a circle, but viewed from the side it looks like a sine wave.

To view a flip-book animation of this three-dimensional visualization, turn to Appendix [A.3](#) beginning on page [2865](#).

5.5.2 Phasor expressions of phase shifts


As previously discussed, Euler's Relation provides a mathematical tool to express AC circuit quantities such as voltage, which change over time and therefore are more challenging to represent than the static quantities found in DC circuits. At its core, Euler's Relation expresses the mathematical equivalence between an imaginary exponential function (e raised to some imaginary power) and a complex number (the sum of a real number and an imaginary number):

$$e^{j\theta} = \cos \theta + j \sin \theta$$

If we wish to represent AC quantities having magnitudes other than 1, we may modify Euler's Relation to include a multiplying coefficient specifying the function's peak value:

$$Ae^{j\theta} = A(\cos \theta + j \sin \theta) = A \cos \theta + jA \sin \theta$$


Phasor notation proves extremely useful to compare or combine AC quantities at the same frequency that are out-of-phase with each other. Consider the following example, showing two AC voltage waveforms of equal magnitude (5 volts peak) that are a constant 60 degrees ($\frac{\pi}{3}$ radians) out of step with each other:



Recalling that the assumed direction of rotation for phasors is counter-clockwise, we see here that phasor B leads phasor A (i.e. phasor B is further counter-clockwise than phasor A). This is

also evident from the rectangular graph¹⁴, where we see sinusoid *A* lags behind sinusoid *B* by $\frac{\pi}{3}$ radians. If we were to connect a dual-channel oscilloscope to these two voltage sources, we would see the same phase shift on its display.

Suppose we wished to calculate the amount of voltage *between* points B and A, sensing the *difference* in potential between these two AC voltage sources. Phasor notation makes this calculation rather easy to do, since all it entails is calculating the distance between the two phasor tips. We will designate this new measurement V_{BA} representing the potential at point B with reference to point A¹⁵:




Simply sketching a new phasor stretching between the tips of phasors B and A yields a graphical solution for V_{BA} . The length of phasor V_{BA} represents the magnitude of that voltage, while its angle from horizontal represents the phase shift of that voltage (compared to phasor A at 0°).

¹⁴The fact that this graph shows the vertical (imaginary) projections of both phasors rather than the horizontal (real) projections is irrelevant to phase shift. Either way, the voltage waveform of source B will still lead the voltage waveform of source A by 60° .

¹⁵One way to think of this is to imagine an AC voltage-measuring instrument having red and black test leads just like a regular voltmeter. To measure V_{BA} you would connect the red test lead to the first point (B) and the black test lead to the second point (A).

If we wished to plot the sine wave represented by this new phasor, all we would have to do is move the V_{BA} phasor until its tail rested at the center of the circular graph (being careful to preserve its angle), then project its tip to the rectangular graph and plot the resulting wave through the new phasor's complete rotation:



We see clearly now how phasor V_{BA} is phase-shifted 120 ahead of (leading) phasor A, and 60° ahead of (leading) phasor B, possessing the same magnitude as both A and B. We may express each of these voltages in polar form:

$$V_A = 5 \text{ volts } \angle 0^\circ$$

$$V_B = 5 \text{ volts } \angle 60^\circ$$


$$V_{BA} = 5 \text{ volts } \angle 120^\circ$$

We may mathematically derive this result by subtracting V_A from V_B to arrive at V_{BA} :

$$V_{BA} = V_B - V_A$$

$$(5 \text{ volts } \angle 60^\circ) - (5 \text{ volts } \angle 0^\circ) = 5 \text{ volts } \angle 120^\circ$$

When using phasors to describe differential voltages (i.e. voltages between two non-grounded points), it is important to specify which point is the reference. Suppose, for example, we took this same two-source system where each source has a peak voltage of 5 volts with reference to ground, but we connected our voltage-measuring instrument backwards from before such that the black (reference) lead touches point B and the red (measurement) lead touches point A (sensing V_{AB} instead of V_{BA}). Our new phasor V_{AB} must then be sketched with its tail (reference) at point B and its tip (measurement) at point A:



The dashed-line phasor is the same angle as the solid-line phasor reaching from point B to point A, just moved so that its tail rests at the circle's center in order to plot its vertical projection on the rectangular graph to the right of the circle. As you can see, the new phasor V_{AB} is 180° shifted from the old phasor V_{BA} . This makes sense, as a reversal of the test leads on our voltage-sensing instrument must reverse the polarity of that voltage at every point in time.

We may mathematically derive this result by subtracting V_B from V_A to arrive at V_{AB} :

$$V_{AB} = V_A - V_B$$

$$(5 \text{ volts } \angle 0^\circ) - (5 \text{ volts } \angle 60^\circ) = 5 \text{ volts } \angle -60^\circ$$

Thus, the new phasor V_{AB} actually *lags* behind both V_A (by 60 degrees) and V_B (by 120 degrees). If the voltage-sensing instrument is nothing more than a voltmeter, there will be no difference in its reading of V_{BA} versus V_{AB} : in both cases it will simply register 5 volts AC peak. However, if the voltage-sensing instrument is an oscilloscope with its sweep triggered by one of the other voltage signals, this difference in phase will be clearly evident.


It is important to remind ourselves that voltage and current phasors, like the AC waveforms they represent, are actually in constant motion. Just as the generator shaft in an actual AC power system spins, the phasor representing that generator's voltage and the phasor representing that generator's current must also spin. The only way a phasor possessing a *fixed* angle makes any sense in the context of an AC circuit is if that phasor's angle represents a relative shift compared to some other phasor at the same frequency¹⁶. Here, we see Euler's Relation written in fixed-angle and rotating-angle forms:

Fixed-angle phasor	$Ae^{j\theta} = A \cos \theta + jA \sin \theta$
Rotating phasor	$Ae^{j\omega t} = A \cos \omega t + jA \sin \omega t$

This concept can be very confusing, representing voltage and current phasors as possessing fixed angles when in reality they are continuously spinning at the system's frequency. In order to make better sense of this idea, we will explore an analysis of the same two-source AC system using a special instrument designed just for the purpose.

¹⁶The necessity of a shared frequency is easily understood if one considers a case of two sine waves at different frequencies: their respective phasors would spin at different speeds. Given two phasors spinning at different speeds, the angle separating those two phasors would be constantly changing. It is only when two phasors spin around at precisely the same speed that we can sensibly talk about there being a fixed angular displacement between them. Fortunately this is the usual case in AC circuit analysis, where all voltages and currents share the same frequency.

Imagine building a small synchronous electric motor with an arrow-shaped magnetic rotor. We will wind the stator coils and magnetize the rotor such that the arrow points toward the 0° mark whenever the red lead is positive and the black lead is negative (connected to a DC test source). When connected to an AC source, this “phasometer” will spin in sync¹⁷ with the AC generator powering the circuit, the rotating arrow keeping pace with the generator’s shaft¹⁸. The intent of this arrow is to serve as a mechanical analogue of an electrical phasor, pointing in the same direction as the phasor points for any sinusoidal voltage or current signal connected to the phasometer:




For a system frequency of 60 Hz as is standard for AC power systems in North America, the rotational speed of our phasometer’s rotor will be 3600 revolutions per minute – too fast for the rotating arrow to be anything but a blur of motion to an unaided human eye. Therefore, we will need to add one more component to the phasometer to make it practical: a *strobe light* connected to an AC voltage in the system to act as a synchronizing pulse. This strobe light will emit a flash of light just once per cycle of the AC waveform, and always at the same point in time (angle) within every cycle. Just as a strobe light (also called a “stroboscope”) makes a moving machine part¹⁹ appear to “freeze” in time, this strobe light will visually “freeze” the arrow so we will be able to read its position with our eyes.

¹⁷An important detail is that our phasometer must always spin counter-clockwise in order to maintain proper phasor convention. We can ensure this will happen by including a pair of *shading coils* (small copper rings wrapped around one corner of each magnetic pole) in the stator structure. For a more detailed discussion of shading coils, refer to the section on AC induction motors (10.4.1) starting on page 726.

¹⁸This, of course, assumes the generator powering the system is also a two-pole machine like the phasometer. If the generator has more poles, the shaft speed will not match the phasometer’s rotor speed even though the phasometer will still faithfully represent the generator’s *cosine wave* rotation.

¹⁹Automobile mechanics may be familiar with a tool called a *timing light*, consisting of a strobe light connected to the engine in such a way that the light flashes every time the #1 cylinder spark plug fires. By viewing the marks etched into the engine’s crankshaft with this strobe light, the mechanic is able to check the ignition timing of the engine.


Returning to our two-source AC system, we may perform a “thought experiment” to see what multiple phasometers would register if connected to various points in the system, using a single strobe light connected to source A (pulsing briefly every time its sensed voltage reaches the positive peak). Imagine this one strobe light is bright enough to clearly illuminate the faces of all phasometers simultaneously, visually “freezing” their arrows at the exact same point in every cycle:



The phasometer connected to source A must register 0° because the strobe light flashes whenever source A hits its positive peak (at 0° on its cosine waveform), and the positive peak of a waveform is the precise point at which the phasometer arrow is designed to orient itself toward the 0° mark. The phasometer connected to source B will register 60° because that source is 60 degrees ahead of (leading) source A, having passed its positive peak already and now headed downward toward zero volts at the point in time when the strobe flashes. The upper phasometer registers -60° at that same time because the voltage it senses at point A with respect to point B is lagging behind source A by that much, the V_{AB} cosine wave heading toward its positive peak.

If we were to substitute a constant illumination source for the strobe light, we would see all three phasometers' arrows spinning counter-clockwise, revealing the true dynamic nature of the voltage phasors over time. Remember that these phasors are all *continuously moving* quantities, because the voltages they represent are sinusoidal functions of time. Only when we use a strobe light keyed to source A's positive peak do we see fixed-angle readings of $V_A = 0^\circ$, $V_B = 60^\circ$, and $V_{AB} = -60^\circ$.

You may wonder what might happen if we keep the three phasometers connected to the same points, but change the strobe light's reference point. The answer to this question is that the new reference source voltage will now be our zero-degree definition, with every phasor's angle changing to match this new reference. We may see this effect here, using the same two sources and three phasometers, but moving the strobe light from source A to source B:



Having the strobe light keyed to source B makes it flash 60 degrees earlier than it did when it was keyed to source A, which in turn “freezes” the arrows on the faces of all three phasometers -60° from where they used to be when source A was the time-reference. The following table compares the phasometer readings with both strobe light reference points:

Strobe reference	V_A	V_B	V_{AB}
Source A	0°	60°	-60°
Source B	-60°	0°	-120°

It is even possible to imagine connecting the strobe light to a special electronic timing circuit pulsing the light at 60 Hz (the AC system's frequency), synchronized to an atomic clock so as to keep ultra-accurate time. Such an arrangement would permit phasor angle measurements based on an absolute time reference (i.e. the atomic clock) rather than a relative time reference (i.e. one of the AC voltage sources). So long as the two AC sources maintained the same frequency and phase shift, the three phasometers would still be displaced 60° from each other, although it would only be by blind luck that any of them would point toward 0° (i.e. that any one voltage would happen to be precisely in-phase with the electronic timer circuit's 60 Hz pulse).

Interestingly, this technique of measuring AC power system phasors against an absolute time reference is a real practice and not just a textbook thought experiment. The electrical power industry refers to this as a *synchrophasor* measurement. Special instruments called *Phasor Measurement Units* or *PMUs* connect to various points within the power system to acquire real-time voltage and current measurements, each PMU also connected to a GPS (Global Positioning System) radio receiver to obtain an absolute time reference with sub-microsecond uncertainty. Each PMU uses the ultra-precise time signal from the GPS receiver to synchronize a “standard” cosine wave at the power system frequency such that this reference waveform is at its positive peak (0°) at the top of every second in time.

Synchrophasor technology makes it possible to perform simultaneous comparisons of phasor angles throughout a power system, which is useful for such tasks as analyzing frequency stability in a power grid or detecting “islanding” conditions where tripped circuit breakers segment a power grid and allow distributed generators to begin drifting out of sync with each other.

5.5.3 Phasor expressions of impedance

The ultimate purpose of phasors is to simplify AC circuit analysis, so this is what we will explore now. Consider the problem of defining electrical opposition to current in an AC circuit. In DC (direct-current) circuits, resistance (R) is defined by Ohm's Law as being the ratio between voltage (V) and current (I):

$$R = \frac{V}{I}$$

There are some electrical components, though, which do not obey Ohm's Law. *Capacitors* and *inductors* are two outstanding examples. The fundamental reason why these two components do not follow Ohm's Law is because they do not dissipate energy like resistances do. Rather than dissipate energy (in the form of heat and/or light), capacitors and inductors *store* and *release* energy from and to the circuit in which they are connected. The contrast between resistors and these components is remarkably similar to the contrast between *friction* and *inertia* in mechanical systems. Whether pushing a flat-bottom box across a floor or pushing a heavy wheeled cart across a floor, work is required to get the object moving. However, the flat-bottom box will immediately stop when you stop pushing it due to the energy loss inherent to friction, while the wheeled cart will continue to coast because it has kinetic energy stored in it. When a resistor disconnected from an electrical source, both voltage and current immediately cease. Capacitors and inductors, however, may store a "charge" of energy when disconnected from a source (capacitors retaining voltage and inductors retaining current).


The relationships between voltage and current for capacitors (C) and inductors (L) are as follows:

$$I = C \frac{dV}{dt} \qquad V = L \frac{dI}{dt}$$

Expressed verbally, capacitors pass electric current proportional to how quickly the voltage across them *changes* over time. Conversely, inductors produce a voltage drop proportional to how quickly current through them *changes* over time. The symmetry here is beautiful: capacitors, which store energy in an electric field that is proportional to the applied voltage, oppose changes in voltage. Inductors, which store energy in a magnetic field that is proportional to applied current, oppose changes in current. The manner in which a capacitor or an inductor reacts to changes imposed upon it is a direct consequence of the *Law of Energy Conservation*: since energy can neither appear from nothing nor simply vanish, an *exchange* of energy must take place in order to alter the amount of energy stored within a capacitor or an inductor. The rate at which a capacitor's voltage may change is directly related to the rate at which electric charge (current) enters or exits the capacitor. The rate at which an inductor's current may change is directly related to the amount of electromotive force (voltage) impressed across the inductor.

When either type of component is placed in an AC circuit and subjected to oscillating signals, it will pass a finite amount of alternating current. Even though the mechanism of a capacitor's or inductor's opposition to current (called *reactance*) is fundamentally different from that of a resistor (called *resistance*), just like inertia differs in its fundamental nature from friction, it is still convenient to express the amount of electrical opposition in a common unit of measurement: the *ohm* (Ω). To do this, we will have to figure out a way to take the above equations and manipulate them to express each component's behavior as a ratio of $\frac{V}{I}$.

Let's start with capacitors. Suppose we impress a 1 volt peak AC voltage across a capacitor, representing that voltage as the exponential $e^{j\omega t}$ where ω is the angular velocity (frequency) of the signal and t is time:



We will begin by writing the current/voltage relationship for a capacitor, along with the imaginary exponential function for the impressed voltage:

$$I = C \frac{dV}{dt} \quad V = e^{j\omega t}$$

Substituting $e^{j\omega t}$ for V in the capacitor formula, we see we must apply the calculus function of *differentiation* to the time-based voltage function:

$$I = C \frac{d}{dt} (e^{j\omega t})$$

Fortunately for us, differentiation is a very simple²⁰ process with exponential functions:

$$I = j\omega C e^{j\omega t}$$

Remember that our goal here is to solve for the ratio of voltage over current for a capacitor. So far all we have is a function for current (I) in terms of time (t). If we take this function for current and divide that into our original function for voltage, however, we see that ratio simplify quite nicely:

$$\frac{V}{I} = \frac{e^{j\omega t}}{j\omega C e^{j\omega t}}$$

$$\frac{V}{I} = \frac{1}{j\omega C} = -j \frac{1}{\omega C}$$

Note²¹ how the exponential term completely drops out of the equation, leaving us with a clean ratio strictly in terms of capacitance (C), angular velocity (ω), and of course j .

²⁰Recall from calculus that the derivative of the function e^x with respect to x is simply e^x . That is, the value of an exponential function's slope is equal to the value of the original exponential function! If the exponent contains any constants multiplied by the independent variable, those constants become multiplying coefficients after differentiation. Thus, the derivative of e^{kx} with respect to x is simply ke^{kx} . Likewise, the derivative of $e^{j\omega t}$ with respect to t is $j\omega e^{j\omega t}$.

²¹Note also one of the interesting properties of the imaginary operator: $\frac{1}{j} = -j$. The proof of this is quite simple: $\frac{1}{j} = \frac{j}{j^2} = \frac{j}{-1} = -j$.

Next, we will apply this same analysis to inductors. Recall that voltage across an inductor and current through an inductor are related as follows:

$$V = L \frac{dI}{dt}$$

If we describe the AC current²² through an inductor using the familiar imaginary exponential expression $I = e^{j\omega t}$ (representing a 1 amp peak AC current at frequency ω), we may substitute this expression for current into the inductor's characteristic equation to solve for the inductor's voltage as a function of time:

$$V = L \frac{dI}{dt} \quad I = e^{j\omega t}$$

$$V = L \frac{d}{dt} (e^{j\omega t})$$

$$V = j\omega L e^{j\omega t}$$

Now that we have the inductor's voltage expressed as a time-based function, we may include the original current function and calculate the ratio of V over I :

$$\frac{V}{I} = \frac{j\omega L e^{j\omega t}}{e^{j\omega t}}$$

$$\frac{V}{I} = j\omega L$$

In summary, we may express the impedance (voltage-to-current ratio) of capacitors and inductors by the following equations:


$$Z_L = j\omega L \quad Z_C = \frac{1}{j\omega C} \text{ or } -j\frac{1}{\omega C}$$

Most students familiar with electronics from an algebraic perspective (rather than calculus) find the expressions $X_L = 2\pi fL$ and $X_C = \frac{1}{2\pi fC}$ easier to grasp. Just remember that angular velocity (ω) is really "shorthand" notation for $2\pi f$, so these familiar expressions may be alternatively written as $X_L = \omega L$ and $X_C = \frac{1}{\omega C}$.

²²Note that we begin this analysis with an exponential expression of the *current* waveform rather than the *voltage* waveform as we did at the beginning of the capacitor analysis. It is possible to begin with voltage as a function of time and use calculus to determine current through the inductor, but unfortunately that would necessitate *integration* rather than *differentiation*. Differentiation is a simpler process, which is why this approach was chosen. If $e^{j\omega t} = L \frac{dI}{dt}$ then $e^{j\omega t} dt = L dI$. Integrating both sides of the equation yields $\int e^{j\omega t} dt = L \int dI$. Solving for I yields $\frac{e^{j\omega t}}{j\omega L}$ plus a constant of integration representing a DC component of current that may or may not be zero depending on where the impressed voltage sinusoid begins in time. Solving for $Z = V/I$ finally gives the result we're looking for: $j\omega L$. Ugly, no?

Furthermore, recall that reactance (X) is a *scalar quantity*, having magnitude but no direction. Impedance (Z), on the other hand, possesses both magnitude *and* direction (phase), which is why the imaginary operator j must appear in the impedance expressions to make them complete. The impedance offered by pure inductors and capacitors alike are nothing more than their reactance values (X) scaled along the imaginary (j) axis (phase-shifted 90°).

These different representations of opposition to electrical current are shown here for components exhibiting 50 ohms, resistances (R) and reactances (X) shown as scalar quantities near the component symbols, and impedances (Z) as phasor quantities on the complex plane:



5.5.4 Phasor arithmetic

Another detail of phasor math that is both beautiful and practical is the famous expression of Euler's Relation, the one all math teachers love because it directly relates several fundamental constants in one elegant equation (remember that i and j mean the same thing, just different notational conventions for different disciplines):

$$e^{i\pi} = -1$$

This equation is actually a special case of Euler's Relation, relating imaginary exponents of e to sine and cosine functions:

$$e^{i\theta} = \cos \theta + i \sin \theta$$

What this equation says is really quite amazing: if we raise e to an imaginary exponent of some angle (θ), it is equivalent to the real cosine of that same angle plus the imaginary sine of that same angle. Thus, Euler's Relation expresses an equivalence between exponential (e^x) and trigonometric ($\sin x$, $\cos x$) functions. Specifically, the angle θ describes which way a phasor points on a complex plane, the real and imaginary coordinates of a unit phasor's²³ tip being equal to the cosine and sine values for that angle, respectively.

If we set the angle θ to a value equal to π , we see the general form of Euler's relation transform into $e^{i\pi} = -1$:

$$e^{i\theta} = \cos \theta + i \sin \theta$$

$$e^{i\pi} = \cos \pi + i \sin \pi$$

$$e^{i\pi} = -1 + i0$$

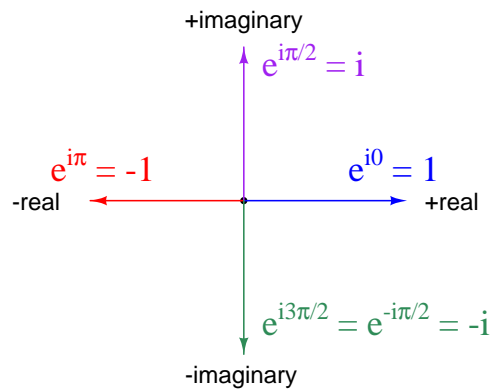
$$e^{i\pi} = -1$$

²³A "unit" phasor is one having a length of 1.

After seeing this, the natural question to ask is what happens when we set θ equal to other, common angles such as 0 , $\frac{\pi}{2}$, or $\frac{3\pi}{2}$ (also known as $-\frac{\pi}{2}$)? The following examples explore these angles:

Angle (θ)	Exponential	Trigonometric	Rectangular	Polar
0 radians = 0°	e^{i0}	$\cos 0 + i \sin 0$	$1 + i0 = 1$	$1 \angle 0^\circ$
$\pi/2$ radians = 90°	$e^{i\pi/2}$	$\cos 90^\circ + i \sin 90^\circ$	$0 + i1 = i$	$1 \angle 90^\circ$
π radians = 180°	$e^{i\pi}$	$\cos 180^\circ + i \sin 180^\circ$	$-1 + i0 = -1$	$1 \angle 180^\circ$
$-\pi/2$ radians = -90°	$e^{i-\pi/2}$	$\cos -90^\circ + i \sin -90^\circ$	$0 - i1 = -i$	$1 \angle -90^\circ$

We may show all the equivalences on the complex plane, as unit phasors:



As we saw previously, the amount of opposition to electrical current offered by reactive components (i.e. inductors and capacitors) – a quantity known as *impedance* – may be expressed as functions of j and ω :

$$Z_L = j\omega L$$

$$Z_C = -j \frac{1}{\omega C}$$

Knowing that j is equal to $e^{j\pi/2}$ and that $-j = e^{-j\pi/2}$, we may re-write the above expressions for inductive and capacitive impedance as functions of an angle:

$$Z_L = \omega L e^{j\pi/2}$$

$$Z_C = \frac{1}{\omega C} e^{-j\pi/2}$$

Using polar notation as a “shorthand” for the exponential term, the impedances for inductors and capacitors are seen to have fixed angles²⁴:

$$Z_L = \omega L e^{j\pi/2} = \omega L \angle \frac{\pi}{2} \text{ radians} = \omega L \angle 90^\circ$$

$$Z_C = \frac{1}{\omega C} e^{-j\pi/2} = \frac{1}{\omega C} \angle -\frac{\pi}{2} \text{ radians} = \frac{1}{\omega C} \angle -90^\circ$$

Beginning electronics students will likely find the following expressions of inductive and capacitive impedance more familiar, $2\pi f$ being synonymous with ω :

$$Z_L = (2\pi f L) \angle 90^\circ$$

$$Z_C = \left(\frac{1}{2\pi f C} \right) \angle -90^\circ$$

²⁴The fact that these impedance phasor quantities have fixed angles in AC circuits where the voltage and current phasors are in constant motion is not a contradiction. Since impedance represents the relationship *between* voltage and current for a component ($Z = V/I$), this fixed angle represents a relative phase shift between voltage and current. In other words, the fixed angle of an impedance phasor tells us the voltage and current waveforms will always remain that much out of step with each other despite the fact that the voltage and current phasors themselves are continuously rotating at the system frequency (ω).

The beauty of complex numbers in AC circuits is that they make AC circuit analysis equivalent to DC circuit analysis. If we represent every voltage and every current and every impedance quantity in an AC circuit as a complex number, all the same²⁵ laws and rules we know from DC circuit analysis will apply to the AC circuit. This means we need to be able to add, subtract, multiply, and divide complex numbers in order to apply Ohm's Law and Kirchhoff's Laws to AC circuits.

The basic rules of phasor arithmetic are listed here, with phasors having magnitudes of A and B , and angles of M and N , respectively:

$$Ae^{jM} + Be^{jN} = (A \cos M + B \cos N) + j(A \sin M + B \sin N)$$

$$Ae^{jM} - Be^{jN} = (A \cos M - B \cos N) + j(A \sin M - B \sin N)$$

$$Ae^{jM} \times Be^{jN} = AB e^{j(M+N)}$$

$$Ae^{jM} \div Be^{jN} = \frac{A}{B} [e^{j(M-N)}]$$

Addition and subtraction lend themselves readily to the *rectangular* form of phasor expression, where the real (cosine) and imaginary (sine) terms simply add or subtract. Multiplication and division lend themselves readily to the *polar* form of phasor expression, where magnitudes multiply or divide and angles add or subtract.

To summarize:


- Voltages and currents in AC circuits may be mathematically represented as *phasors*, which are imaginary exponential functions (e raised to imaginary powers)
- Phasors are typically written in either rectangular form (real + imaginary) or polar form (magnitude @ angle)
- Ohm's Law and Kirchhoff's Laws still apply in AC circuits as long as all quantities are in phasor notation
- Addition is best done in rectangular form: *add the real parts, and add the imaginary parts*
- Subtraction is best done in rectangular form: *subtract the real parts, and subtract the imaginary parts*
- Multiplication is best done in polar form: *multiply the magnitudes, and add the angles*
- Division is best done in polar form: *divide the magnitudes, and subtract the angles*

It should be noted that many electronic calculators possess the ability to perform all these arithmetic functions in complex-number form. If you have access to such a calculator, it will *greatly* simplify any AC circuit analysis performed with complex numbers!

²⁵With one notable exception: Joule's Law ($P = IV$, $P = V^2/Z$, $P = I^2Z$) for calculating power does not apply in AC circuits because *power* is not a phasor quantity like voltage and current.

5.5.5 Phasors and circuit measurements


A *phasor* is a special form of vector (a quantity possessing both magnitude and direction) lying in a complex plane. Phasors relate circular motion to simple harmonic (sinusoidal) motion as shown in the following diagram. In AC electrical circuits, this is the relationship between an electromechanical generator's shaft angle²⁶ and the sinusoidal voltage it outputs:



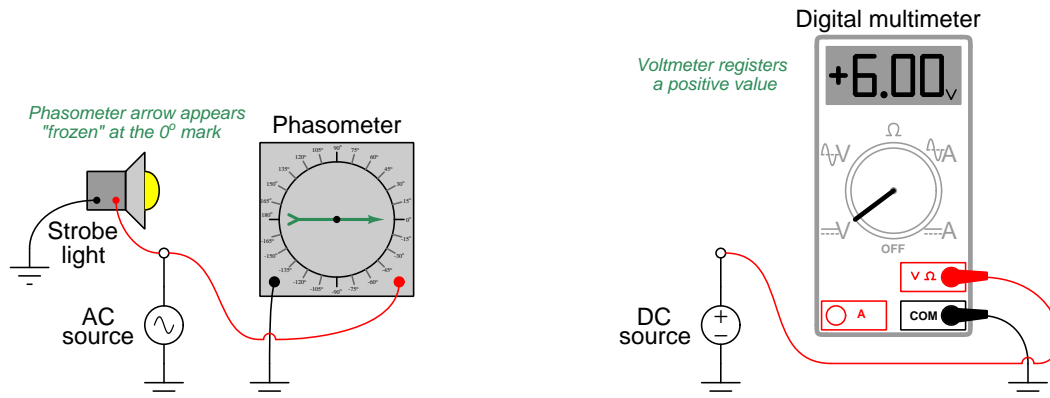
In any operating AC electrical circuit the phasors, just like the sinusoidal waveforms, never rest: they are in continuous motion. When we speak of a phasor as having a fixed angle, what we really mean is that the phasor is either leading or lagging with respect to some other “reference” phasor in the system, not that the phasor itself is stationary. We explored this concept previously, where we used an invented instrument called a “phasometer” to represent the direction of a measured phasor in real time, and then used a strobe light connected to some point in the same system to visually “freeze” the phasometer arrows so we could see which direction each arrow was pointed at the moment in time when our chosen reference waveform reached its positive peak value (i.e. 0° on a cosine wave). The fixed angle represented by each “strobed” phasometer arrow therefore represented the amount of *relative phase shift* between each respective phasor and the reference phasor.

²⁶ Assuming a two-pole generator, where each period of the sinusoidal waveform corresponds exactly to one revolution of the generator shaft.

Phasor angles are to AC quantities what *arithmetic signs* are to DC quantities. If a phasometer registers an angle of 180 degrees, it means the red lead is fully negative and the black lead is fully positive at the moment in time when the strobe light flashes. If a DC voltmeter registers a negative voltage, it means the red lead is negative and the black lead is positive at the time of the measurement:



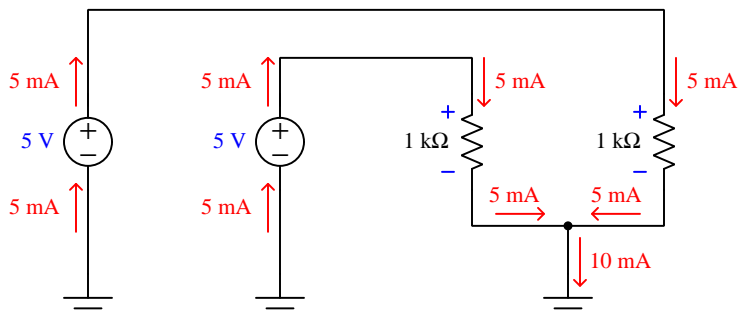
Reversing each instrument's test lead connections to the circuit will reverse its indication: flipping red and black leads on the phasometer will cause its indication to be 0° instead of 180° ; flipping red and black leads on the DC voltmeter will cause it to register a positive voltage instead of a negative voltage:



From these demonstrations we can see that the indication given by a measuring instrument depends as much on how that instrument connects to the circuit as it does on the circuit quantity itself. Likewise, a voltage or current value obtained in the course of analyzing a circuit depends as much on how we label the assumed polarity (or direction) of that quantity on the diagram as it does on the quantity itself.

To illustrate the importance of signs, assumed directions, and phasor angles in circuit analysis, we will consider two multi-source circuits: one DC and one AC. In each case we will explore how Kirchhoff's Current Law relates to currents at a similar node in each circuit, relating the current arrow directions to signs and phasor angles.

First, the DC circuit:




With each 1 kΩ resistor powered by its own 5 volt DC source, the current through each of these resistors will be 5 milliamps in accordance with Ohm's Law ($I = \frac{V}{R}$). The direction of each current is easily predicted by examining the polarity of each source (with current represented in conventional flow notation *exiting* the positive terminal and *entering* the negative terminal of each source) and/or by examining the polarity of the voltage drop across each resistor (with conventional-flow current *entering* the positive terminal and *exiting* the negative terminal of each load). At the node below the two resistors, these two currents join together to form a larger (10 milliamp) current headed toward ground. Kirchhoff's Current Law declares that the sum of all currents entering a node must equal the sum of all currents exiting that node. Here we see this is true: 5 milliamps in plus 5 milliamps in equals 10 milliamps out.

The voltage polarities and current directions in this DC circuit are all clear and unambiguous because the quantities are *constant* over time. Each power supply acts as an energy source and each resistor acts as an energy load, all the time. Our application of Kirchhoff's Current Law at the node is so obvious it hardly requires explanation: the flow of charge carriers (current) in and out must be in equilibrium.

With AC, however, we know things will not be so simple. AC circuit quantities are not constant over time as they are in DC circuits. Reactive (energy-storing) components such as inductors and capacitors play alternating roles as sources and loads, complicating their analysis. The good news in all of this is that phasor representations of voltage and current allow us to apply all the same fundamental principles we are accustomed to using for DC circuits (e.g. Ohm's Law, Kirchhoff's Voltage Law, Kirchhoff's Current Law, network theorems, etc.). In order to use these principles to calculate AC quantities, though, we must be careful in how we label them in the circuit.

Here, we will modify the circuit to include two AC power sources (phase-shifted by 60°), and replace one of the $1\text{ k}\Omega$ resistors with a capacitor exhibiting $1\text{ k}\Omega$ of reactance at the system frequency. As in the DC circuit, each of the two loads sees the full 5 volts of its respective source. Unlike the DC circuit, we must represent each of the voltage and impedance quantities in complex (phasor) form in order to apply Ohm's Law to calculate load currents.

One more thing we will do to this circuit before beginning any phasor arithmetic is to label each voltage and each current as we did in the DC circuit. Placing $+$ and $-$ symbols at the terminals of AC voltage sources, and drawing conventional-flow notation arrows showing *alternating current* through loads at first may seem absurd, since we know this is an AC circuit and so by definition these quantities lack fixed polarities and directions. However, these symbols will become very important to us because they serve to define what 0° means for each of the respective phasors. In other words, these polarities and arrows merely show us which way the voltages and currents are oriented *when each of those phasors is at its positive peak* – i.e. *when each phasor angle comes around to its own 0° mark*. As such, the polarities and arrows we draw do *not necessarily* represent simultaneous conditions. We will rely on the calculated phasor angles to tell us how far these quantities will actually be phase-shifted from each other at any given point in time:



First, calculating current through the resistor²⁷, recalling that the impedance of a resistor has a 0 degree phase angle (i.e. no phase shift between voltage and current):

$$I_R = \frac{V_R}{Z_R} = \frac{5\text{ V} \angle 0^\circ}{1\text{ k}\Omega \angle 0^\circ} = 5\text{ mA} \angle 0^\circ$$

Next, calculating current through the capacitor²⁸, recalling that the impedance for a capacitor has a -90 degree phase angle because voltage across a capacitor lags 90 degrees behind current through a capacitor:

$$I_C = \frac{V_C}{Z_C} = \frac{5\text{ V} \angle 60^\circ}{1\text{ k}\Omega \angle -90^\circ} = 5\text{ mA} \angle 150^\circ$$

²⁷When dividing two phasors in polar form, the arithmetic is as follows: divide the numerator's magnitude by the denominator's magnitude, then subtract the denominator's angle from the numerator's angle. The result in this case is 5 milliamps (5 volts divided by 1000 ohms) at an angle of 0 degrees (0 minus 0).

²⁸The same arithmetic applies to this quotient as well: the current's magnitude is 5 volts divided by 1000 ohms, while the current's phase angle is 60 degrees minus a negative 90 degrees (150 degrees).

From the arrows sketched at the node we can see the total current headed toward ground must be equal to the sum of the two currents headed into the node, just as with the DC circuit. The difference here is that we must perform the addition using phasor quantities instead of scalar quantities in order to account for phase shift between the two currents:

$$I_{total} = I_R + I_C$$

$$I_{total} = 5 \text{ mA} \angle 0^\circ + 5 \text{ mA} \angle 150^\circ$$

Representing these current phasors in rectangular mode so we may sum their real and imaginary parts:


$$I_{total} = (5 + j0 \text{ mA}) + (-4.33 + j2.5 \text{ mA})$$

$$I_{total} = 0.67 + j2.5 \text{ mA}$$


Converting the rectangular form into polar form:

$$I_{total} = 2.59 \text{ mA} \angle 75^\circ$$


This result is highly non-intuitive. When we look at the circuit and see two 5 milliamp currents entering a node, we naturally expect a sum total of 10 milliamps to exit that node. However, that will only be true if those 5 mA quantities are *simultaneous*: i.e. if the two currents are 5 mA in magnitude *at the same point in time*. Since we happen to know these two currents are phase-shifted from each other by 150 degrees (nearly opposed to each other), they never reach their full strength at the same point in time, and so their sum is considerably less than 10 milliamps:



We may make more sense of this result by adding current-sensing phasometers to the circuit, their red and black test leads connected in such a way as to match the arrows' directions (current entering the red lead and exiting the black lead, treating the phasometer as a proper electrical *load*) as though these were DC ammeters being connected into a DC circuit:




A very common and practical example of using “DC notation” to represent AC voltages and currents is in *three-phase* circuits where each power supply consists of three AC voltage sources phase-shifted 120 degrees from each other. Here, we see a “wye” connected generator supplying power to a “wye” connected load. Each of the three generator stator windings outputs 277 volts, with the entirety of that voltage dropped across each of the load resistances:



Note the directions of the three currents (red arrows) in relation to the node at the center of the generator’s “wye” winding configuration, and also in relation to the node at the center of the resistive load’s “wye” configuration. At first this seems to be a direct violation of Kirchhoff’s Current Law: how is it possible to have three currents *all exiting a node with none entering* or to have three currents *all entering a node with none exiting*? Indeed, this would be impossible if the currents were *simultaneously* moving in those directions, but what we must remember is that each of the current arrows simply shows which way each current will be moving *when its respective phasor comes around to the 0° mark*. In other words, the arrows simply define what zero degrees means for each current. Similarly, the voltage polarities would suggest to anyone familiar with Kirchhoff’s Voltage Law in DC circuits that the voltage existing between any two of the power conductors between generator and load should be 0 volts, reading the series-opposed sum of two 277 volt sources. However, once again we must remind ourselves that the + and – symbols do not actually represent polarities at the same point in time, but rather serve to define the orientation of each voltage when its phasor happens to point toward 0°.

If we were to connect three current-sensing phasometers to measure the phase angle of each line current in this system (keying the strobe light to the voltage of the 0° generator winding), we would see the true phase relationships relative to that time reference. The directions of the current arrows and the orientation of the $+$ and $-$ polarity marks serve as references for how we shall connect the phasometers to the circuit ($+$ on red, $-$ on black; current entering red, exiting black):



We may use the three arrows at the load's center node to set up a Kirchhoff's Current Law equation (i.e. the sum of all currents at a node must equal zero), and confirm that the out-of-phase currents all "entering" that node do indeed amount to a sum of zero:

$$I_{node} = 2.77 \text{ A} \angle 0^\circ + 2.77 \text{ A} \angle 120^\circ + 2.77 \text{ A} \angle 240^\circ$$

Converting polar expressions into rectangular so we may see how they add together:

$$I_{node} = (2.77 + j0 \text{ A}) + (-1.385 + j2.399 \text{ A}) + (-1.385 - j2.399 \text{ A})$$

Combining all the real terms and combining all the imaginary terms:

$$I_{node} = (2.77 - 1.385 - 1.385) + j(0 + 2.399 - 2.399) \text{ A}$$

$$I_{node} = 0 + j0 \text{ A}$$

Looking closer at these results, we may determine which way the three currents were actually flowing at the time of the strobe's flash (when the upper-right generator winding is at its peak positive voltage). The first phasor has a real value of +2.77 amps at that instant in time, which represents 2.77 amps of actual current flowing in the actual direction of our sketched current arrow. The second and third phasors each have real values of -1.385 amps at that same instant in time, representing 1.385 amps each flowing *against* the direction of our sketched arrows (i.e. leaving that node). Thus, we can see that the "snapshot" view of currents at the time of the strobe's flash makes complete sense: one current of 2.77 amps entering the node and two currents of 1.385 amps each exiting the node. Viewed at any instant in time, the principles of DC circuits hold true. It is only when we sketch polarity marks and draw arrows representing voltages and currents *reaching their positive peak values at different times* that the annotations seem to violate basic principles of circuit analysis.

5.6 The s variable

A powerful mathematical concept useful for analyzing practically any physical system – electrical circuits included – is something called the s variable. The s variable is closely related to Euler's Relation and phasor expressions of waveforms, which is why a discussion of it is included here.

5.6.1 Meaning of the s variable

As we saw previously, Euler's Relation allows us to express rotating phasors as imaginary exponents of e . For example, $Ae^{j\theta}$ represents a phasor of length A at an angle of θ radians. $Ae^{j\omega t}$ represents a phasor of length A rotating at a velocity of ω radians per second at a particular instant in time t . This happens to be an incredibly useful mathematical "trick" for representing sinusoidal waves in physical systems. For example, if we wished to mathematically express a sinusoidal AC voltage as a function of time with a peak voltage value of 10 volts and a frequency of 60 hertz (377 radians per second, since $\omega = 2\pi f$), we could do so like this:

$$V(t) = 10e^{j377t}$$

Exponential functions aren't just useful for expressing sinusoidal waves, however. They also work well for expressing rates of *growth* and *decay*, as is the case with RC and L/R time-delay circuits where exponential functions describe the charging and discharging of capacitors and inductors. Here, the exponent is a real number rather than an imaginary number: the expression $e^{-t/\tau}$ approaching zero as time (t) increases. The Greek letter "tau" (τ) represents the *time constant* of the circuit, which for capacitive circuits is the product of R and C , and for inductive circuits is the quotient of L and R . For example, if we wished to mathematically express the decaying voltage across a 33 μF capacitor initially charged to 10 volts as it dissipates its stored energy through a 27 $\text{k}\Omega$ resistor (the circuit having a time constant of 0.891 seconds, since $\tau = RC$), we could do so like this:

$$V(t) = 10e^{-(t/0.891)}$$

The sign of the exponential term here is very important: in this example we see it is a negative number. This tells us the function *decays* (approaches zero) over time, since larger positive values of t result in larger negative values of t/τ (recall from algebra that a negative exponent is the equivalent of reciprocating the expression, so that $e^{-x} = \frac{1}{e^x}$). If the exponent were a real positive number, it would represent some quantity *growing* exponentially over time. If the exponent were zero, it would represent a *constant* quantity. We expect a discharging resistor-capacitor circuit to exhibit decaying voltage and current values, and so the negative exponent sign shown here makes sense.

If imaginary exponents of e represent phasors, and real exponents of e represent growth or decay, then a *complex* exponent of e (having both real and imaginary parts) must represent a phasor that grows or decays in magnitude over time. Engineers use the lower-case Greek letter "omega" (ω) along with the imaginary operator j to represent the imaginary portion, and the lower-case Greek letter "sigma"²⁹ (σ) to represent the real portion. For example, if we wished to mathematically express a sine wave AC voltage with a frequency of 60 hertz (377 radians per second) and an amplitude beginning at 10 volts but decaying with a time constant (τ) of 25 milliseconds ($\sigma = 1/\tau = 40$ time constants per second), we could do so like this:

$$V(t) = 10e^{-40t + j377t}$$

We may factor time from the exponential terms in this expression, since t appears both in the real and imaginary parts:

$$V(t) = 10e^{(-40 + j377)t}$$

²⁹ σ is equal to the reciprocal of the signal's time constant τ . In other words, $\sigma = 1/\tau$.

With t factored out, the remaining terms $-40 + j377$ completely describe the sinusoidal wave's characteristics. The wave's decay rate is described by the real term ($\sigma = -40$ time constants per second), while the wave's phase is described by the imaginary term ($j\omega = 377$ radians per second). Engineers use a single variable s to represent the complex quantity $\sigma + j\omega$, such that any growing or decaying sinusoid may be expressed very succinctly as follows:

$$Ae^{st} = Ae^{(\sigma+j\omega)t} = Ae^{\sigma t}e^{j\omega t}$$


Where,

- A = Initial amplitude of the sinusoid (e.g. volts, amps) at time $t = 0$ (arbitrary units)
- s = Complex growth/decay rate and frequency (sec^{-1})
- $\sigma = \frac{1}{\tau}$ = Real growth/decay rate (time constants per second, or sec^{-1})
- $j\omega$ = Imaginary frequency (radians per second, or sec^{-1})
- t = Time (seconds)


Separating the expression $Ae^{\sigma t}e^{j\omega t}$ into three parts - A , $e^{\sigma t}$, and $e^{j\omega t}$ - we get a complete description of a rotating phasor:

- A = Initial amplitude of the phasor ($t = 0$)
- $e^{\sigma t}$ = How much the phasor's magnitude has grown ($\sigma > 0$) or decayed ($\sigma < 0$) at time t
- $e^{j\omega t}$ = Unit phasor (length = 1) at time t

If we set ω at some constant value and experiment with different values of σ , we can see the effect σ has on the shape of the wave over time:



If we set σ at zero and experiment with different values³⁰ of ω , we can see the effect ω has on the shape of the wave over time:



As we will soon see, characterizing a sinusoidal response using the complex variable s allows us to mathematically describe a great many things. Not only may we describe voltage waveforms using s as shown in these simple examples, but we may also describe the response of entire physical systems including electrical circuits, machines, feedback control systems, and even chemical reactions. In fact, it is possible to map the essential characteristics of *any* linear system in terms of how exponentially growing, decaying, or steady sinusoidal waves affect it, and that mapping takes the form of mathematical functions of s .

When engineers or technicians speak of a *resonant* system, they mean a circuit containing inductive and capacitive elements tending to sustain oscillations of a particular frequency (ω). A lossless resonant system (e.g. a superconducting tank circuit, a frictionless pendulum) may be expressed by setting the real portion of s equal to zero ($\sigma = 0$; no growth or decay) and

³⁰One value of ω not shown in this three-panel graphic comparison is a *negative* frequency. This is actually not as profound as it may seem at first. All a negative value of ω will do is ensure that the phasor will rotate in the opposite direction (clockwise, instead of counter-clockwise as phasor rotation is conventionally defined). The real portion of the sinusoid will be identical, tracing the same cosine-wave plot over time. Only the imaginary portion of the sinusoid will be different, as $j \sin -\theta = -j \sin \theta$.

letting the imaginary portion represent the resonant frequency ($j\omega = j2\pi f$). Real-life resonant systems inevitably dissipate some energy, and so a real resonant system's expression will have both an imaginary portion to describe resonant frequency and a negative real portion to describe the oscillation's rate of decay over time.

Systems exhibiting a positive σ value are especially interesting because they represent *instability*: unrestrained oscillatory growth over time. A feedback control loop with excessive gain programmed into the loop controller is a simple example of a system where $\sigma > 1$. This situation, of course, is highly undesirable for any control system where the goal is to maintain the process variable at a steady setpoint.

5.6.2 Impedance expressed using the s variable

Previously, we saw how the impedance of inductors and capacitors could be calculated using $j\omega$ to represent the frequency of the applied signal. Doing so greatly simplified the mathematics by eliminating the need to manipulate trigonometric functions such as sine and cosine. Here, we will discover that s works just as nicely for the same task, with the added benefit of showing how inductors and capacitors react to exponentially growing or decaying signals.

First, let's begin with capacitors. We know that voltage across a capacitor and current "through" a capacitor are related as follows:

$$I = C \frac{dV}{dt}$$

Next, we substitute an expression³¹ for voltage in terms of s and then use calculus to differentiate it with respect to time:

$$I = C \frac{d}{dt} (e^{st})$$

$$I = sC e^{st}$$

The ratio of $\frac{V}{I}$ (the definition of impedance) will then be:

$$Z_C = \frac{V}{I} = \frac{e^{st}}{sC e^{st}}$$

$$Z_C = \frac{1}{sC}$$

Instead of the common scalar expression for capacitive impedance ($Z_C = \frac{1}{2\pi fC}$) which only tells us the magnitude of the impedance (in ohms) but not the phase shift, we have a complex expression for capacitive impedance ($Z_C = \frac{1}{sC}$) describing magnitude, phase shift, and its reaction to the growth or decay of the signal.

³¹The expression used here to represent voltage is simply e^{st} . I could have used a more complete expression such as Ae^{st} (where A is the initial amplitude of the signal), but as it so happens this amplitude is irrelevant because there will be an "A" term in both the numerator and denominator of the impedance quotient. Therefore, A cancels out and is of no consequence.

Likewise, we may do the same for inductors. Recall that voltage across an inductor and current through an inductor are related as follows:

$$V = L \frac{dI}{dt}$$

Substituting an expression for current in terms of s and using calculus to differentiate it with respect to time:

$$V = L \frac{d}{dt} (e^{st})$$

$$V = sLe^{st}$$

The ratio of $\frac{V}{I}$ (the definition of impedance) will then be:

$$Z_L = \frac{V}{I} = \frac{sLe^{st}}{e^{st}}$$

$$Z_L = sL$$

As with capacitors, we now have a complex expression for inductive impedance describing magnitude, phase shift, and its reaction to signal growth or decay ($Z_L = sL$) instead of merely having a scalar expression for inductive impedance ($Z_L = 2\pi fL$).

Resistors directly oppose current by dropping voltage, with no regard to rates of change. Therefore, there are no derivatives in the relationship between voltage across a resistor and current through a resistor:

$$V = IR$$

If we substitute e^{st} for current into this formula, we will see that voltage must equal Re^{st} . Solving for the ratio of voltage over current to define impedance:

$$Z_R = \frac{V}{I} = \frac{Re^{st}}{e^{st}}$$

$$Z_R = R$$

Not surprisingly, all traces of s cancel out for a pure resistor: its impedance is exactly equal to its DC resistance.

In summary:

Inductive impedance (Z_L)	Capacitive impedance (Z_C)	Resistive impedance (Z_R)
sL	$1/sC$	R

Now let's explore these definitions of impedance using real numerical values. First, let's consider a $22 \mu\text{F}$ capacitor exposed to a steady AC signal with a frequency of 500 Hz. Since the signal in this case is steady (neither growing nor decaying in magnitude), the value of σ will be equal to zero. ω is equal to $2\pi f$, and so a frequency of 500 Hz is equal to 3141.6 radians per second. Calculating impedance is as simple as substituting these values for s and computing $1/sC$:

$$Z_C = \frac{1}{sC} = \frac{1}{(\sigma + j\omega)C}$$

$$Z_C = \frac{1}{(0 + j3141.6 \text{ sec}^{-1})(22 \times 10^{-6} \text{ F})}$$

$$Z_C = \frac{1}{j0.0691}$$

$$Z_C = \frac{-j}{0.0691}$$

$$Z_C = 0 - j14.469 \Omega \quad (\text{rectangular notation})$$

$$Z_C = 14.469 \Omega \angle -90^\circ \quad (\text{polar notation})$$

Thus, the impedance of this capacitor will be 14.469 ohms at a phase angle of -90° . The purely imaginary nature of this impedance (its orthogonal phase shift between voltage and current) tells us there is no net power dissipated by the capacitor. Rather, the capacitor spends its time alternately absorbing and releasing energy to and from the circuit.

Next, we will consider the case of a 150 mH inductor exposed to an exponentially rising DC signal with a time constant (τ) of 5 seconds. 5 seconds per time constant (τ) is equal to 0.2 time constants per second (σ). Since the signal in this case is DC and not AC, the value of ω will be equal to zero. Calculating impedance, once again, is as simple as substituting these values for s and computing sL :

$$Z_L = sL = (\sigma + j\omega)L$$

$$Z_L = (0.2 + j0 \text{ sec}^{-1})(150 \times 10^{-3} \text{ H})$$

$$Z_L = 0.03 + j0 \Omega \quad (\text{rectangular notation})$$

$$Z_L = 0.03 \Omega \angle 0^\circ \quad (\text{polar notation})$$

Thus, the impedance of this inductor will be 0.03 ohms at a phase angle of 0° . The purely real nature of this impedance (i.e. no phase shift between voltage and current) tells us energy will be continually absorbed by the inductor, and for this reason it will be seen by the rest of the circuit as though it were a resistor dissipating energy for however long the signal continues to exponentially grow.

A phase shift of 0 degrees for a reactive component such as an inductor may come as a surprise to students accustomed to thinking of inductive impedances *always* having 90 degree phase shifts! However, the application of the complex variable s to impedance mathematically demonstrates we can indeed have conditions of no phase shift given just the right circumstances. This makes conceptual sense as well if we consider how inductors store energy: if the current through an inductor increases exponentially over time, never reversing direction, it means the inductor's magnetic field will always be growing and therefore absorbing more energy from the rest of the circuit.

We see something even more interesting happen when we subject a reactive component to a *decaying* DC signal. Take for example a 33,000 μF capacitor exposed to a decaying DC signal with a time constant of 65 milliseconds. 65 milliseconds per time constant (τ) is equal to 15.38 time constants per second (σ). Once again ω will be zero because this is a non-oscillating signal. Calculating capacitive impedance:

$$Z_C = \frac{1}{sC} = \frac{1}{(\sigma + j\omega)C}$$

$$Z_C = \frac{1}{(-15.38 + j0 \text{ sec}^{-1})(33000 \times 10^{-6} \text{ F})}$$

$$Z_C = \frac{1}{-0.508}$$

$$Z_C = -1.970 + j0 \Omega \quad (\text{rectangular notation})$$

$$Z_C = 1.970 \Omega \angle 180^\circ \quad (\text{polar notation})$$

A *negative* real impedance figure represents a phase shift of 180° between voltage and current. Once again, this may surprise students of electronics who are accustomed to thinking of capacitive impedances *always* having phase shifts of -90 degrees. What a 180 degree phase shift means is the direction of current with respect to voltage polarity has the capacitor functioning as an energy *source* rather than as a *load*. If we consider what happens to a capacitor when it discharges, the 180 degree phase shift makes sense: current flowing in this direction depletes the capacitor's plates of stored charge, which means the electric field within the capacitor weakens over time as it releases that energy to the rest of the circuit.

5.7 Transfer function analysis

An extremely important topic in engineering is that of *transfer functions*. Simply defined, a transfer function is the ratio of output to input for any physical system, usually with both the output and input being mathematical functions of s . In other words, we express both the output of a system and the corresponding input of that system in terms of exponentially growing/decaying sinusoidal waves and then solve for the ratio of those two expressions.

Unfortunately, the teaching of transfer functions and their relation to real-world phenomena is often obfuscated by a heavy emphasis on mathematics. The intent of this section is to introduce this concept in a way that is very “gentle” and continually referenced to real-world applications. If I can do anything to help lift the veil of mystery surrounding concepts such as transfer functions, the s variable, and pole-zero plots, then technicians as well as engineers will be able to appreciate the power of this analytical technique and be able to exchange more ideas in the same “language”.

A simple example of a transfer function is the *gain* of an electronic amplifier. As all students of electronics learn, “gain” is the ratio of output signal to input signal for a circuit. Beginning students learn to represent circuit gains as scalar values (e.g. “The amplifier has a voltage gain of 24”), first as plain ratios and later as decibel figures (e.g. “The amplifier has a voltage gain of 27.6 dB”). One limitation of this approach is that it oversimplifies the situation when the gain of the circuit in question varies with the frequency and/or the growth/decay rate of the signal, which is quite often the case. If we take the engineering approach of expressing output and input signals as functions of s , we obtain a more complete picture of that circuit’s behavior over a wide range of conditions.

Another simple example of a transfer function is what we have just seen in this book: the *impedance* of a reactive electrical component such as a capacitor or an inductor. Here, the ratio in question is between voltage and current. If we consider current through the component to be the “input” signal and voltage across the component to be the “output” signal – both expressed in terms of s – then impedance $Z(s) = \frac{V(s)}{I(s)}$ is the transfer function for that component. This raises an important point about transfer functions: what we define as the “input” and the “output” of the system is quite arbitrary, so long as there is an actual relationship between the two signals.

If we write generalized output/input transfer functions of s for an AC circuit, we may mathematically analyze that transfer function to gain insight into the behavior and characteristics of that circuit. Features of transfer functions of interest to us include:


- **Zeros:** any value(s) of s resulting in a zero value for the transfer function (i.e. zero gain)
- **Poles:** any value(s) of s resulting in an infinite value for the transfer function (i.e. maximum gain)

An AC circuit’s *zeros* tell us where the circuit is unresponsive to input stimuli. An AC circuit’s *poles* tell us where the circuit is able to generate an output signal with no input stimulus (i.e. its natural or *un-driven* mode(s) of response).

In order to clearly understand the concept of transfer functions, practical examples are very helpful. Here we will explore some very simple AC circuits in order to grasp what transfer functions are and how they benefit system analysis.

5.7.1 Example: LR low-pass filter circuit

First, let's begin with a simple low-pass filter circuit comprised of an inductor and a resistor connected in series:



The impedance of each component as a function of s is shown in the diagram: the inductor's impedance is sL while the resistor's impedance is simply R . It should be clear to any student of electronics that these two components will function as a *voltage divider*, with the output voltage being some fraction of the input voltage. Knowing this, we may write a transfer function for this circuit based on the voltage divider formula, which tells us the ratio of output voltage to input voltage is the same as the ratio of output impedance to total impedance:

$$\text{Transfer function} = \frac{V_{out}(s)}{V_{in}(s)} = \frac{R}{R + sL} = \frac{R}{R + (\sigma + j\omega)L}$$

This transfer function allows us to calculate the “gain” of the system for any given value of s , which brings us to the next step of our analysis. At this point we will ask ourselves three questions³²:

1. How does this system respond when $s = 0$?
2. What value(s) of s make the transfer function approach a value of zero?
3. What value(s) of s make the transfer function approach a value of infinity?

The first of these questions refers to a condition where we apply a steady DC signal to the input of the system. If $s = 0$ then both σ and ω must each be equal to zero. A value of zero for σ means the signal is neither growing nor decaying over time, but remains at some unchanging value. A value of zero for ω means the signal is not oscillating. These two conditions can only refer to a steady DC signal applied to the circuit. Substituting zero for s we get:

$$\frac{R}{R + 0L}$$

$$\frac{R}{R} = 1$$

³²What we are really doing here is applying a problem-solving technique I like to call *limiting cases*. This is where we simplify the analysis of some system by considering scenarios where the mathematical quantities are easy to compute.

Therefore the transfer function of this circuit is unity (1) under DC conditions. This is precisely what we would expect given an inductor connected in series with a resistor, with output voltage taken across the resistor. If there is no change in the applied signal, then the inductor's magnetic field will be unchanging as well, which means it will drop zero voltage (assuming a pure inductor with no wire resistance) leaving the entire input voltage dropped across the resistor.

The second question refers to a condition where the output signal of this circuit is zero. Any values of s resulting in zero output from the system are called the *zeros* of the transfer function. Examining the transfer function for this particular low-pass LR filter circuit, we see that this can only be true if s becomes infinitely large, because s is located in the denominator of the fraction:

$$\frac{R}{R \pm \infty L} = 0$$

This is consistent with the behavior of a low-pass filter: as frequency (ω) increases, the filter's output signal diminishes. The transfer function doesn't just tell us how this circuit will respond to change in frequency, however – it also tells us how the circuit will respond to growing or decaying signals too. Here, we see that infinitely large σ values also result in zero output: the inductor, which tends to oppose any current exhibiting a high rate of change, doesn't allow much voltage to develop across the resistor if the input signal is growing or decaying very rapidly.

The third question refers to a condition where either the transfer function's numerator approaches infinity or its denominator approaches zero. Any values of s having this result are called the *poles* of the transfer function. Since the numerator in this particular case is a constant (R), only a denominator value of zero could cause the transfer function to reach infinity:

$$\frac{R}{R + sL} = \infty \text{ only if } R + sL = 0$$

If the necessary condition for a “pole” is that $R + sL = 0$, then we may solve for s as follows:

$$R + sL = 0$$

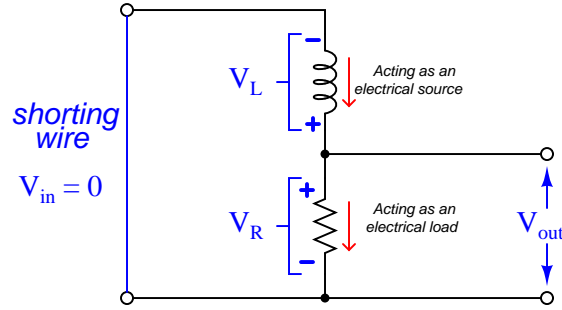
$$sL = -R$$

$$s = -\frac{R}{L}$$

Thus, this transfer function for this simple low-pass filter circuit has one pole located at $s = -R/L$. Since both R and L are real numbers (not imaginary) with positive values, then the value of s for the pole must be a real number with a negative value. In other words, the solution for s at this pole is all σ and no ω : this refers to an *exponentially decaying DC signal*.


It is important at this point to consider what this “pole” condition means in real life. The notion that a circuit is able to produce an output signal with zero input signal may sound absurd, but it makes sense if the circuit in question has the ability to store and release energy. In this particular circuit, the inductor is the energy-storing component, and it is able to produce a voltage drop across the resistor with zero input voltage in its “discharging” mode.

An illustration helps make this clear. If the “pole” condition is such that $V_{in}(s) = 0$, we may show this by short-circuiting the input of our filter circuit to ensure a zero-input condition:



Assuming the inductor has been “charged” with energy previous to the short-circuiting of the input, an output voltage will surely develop across the resistor as the inductor discharges. In other words, the inductor behaves as an electrical *source* while the resistor behaves as an electrical *load*: connected in series they must of course share the same current, but their respective voltages are equal in magnitude and opposing in polarity in accordance with Kirchhoff’s Voltage Law. Furthermore, the value of s in this “pole” condition tells us exactly how rapidly the output signal will decay: it will do so at a rate $\sigma = -R/L$. Recall that the growth/decay term of the s variable the reciprocal of the system’s time constant ($\sigma = 1/\tau$). Therefore, a σ value of R/L is equivalent to a time constant of L/R , which as all beginning students of electronics learn is how we calculate the time constant for a simple inductor-resistor circuit.

Transfer functions are easier to understand when graphically plotted as three-dimensional surfaces: the real and imaginary portions of the s variable occupying the horizontal axes, and the magnitude of the transfer function fraction displayed as height. Here is a *pole-zero plot* of this low-pass filter circuit's transfer function, with a resistor value of $R = 5 \Omega$ and an inductor value of $L = 10 \text{ H}$:




This surface plot makes the meaning of the term “pole” quite obvious: the shape of the function looks just like a rubber mat stretched up at one point by a physical pole. Here, the “pole” rises to an infinite³³ height at a value of s where $\sigma = -0.5$ time constants per second and $\omega = 0$ radians per second. The surface is seen to decrease in height at all edges of the plot, as σ and ω increase in value.

The “zero” of this transfer function is not as obvious as the pole, since the function’s value does not equal zero unless and until s becomes infinite, which of course cannot be plotted on any finite domain. Suffice to say that the zero of this transfer function lies in all horizontal directions at an infinite distance away from the plot’s origin (center), explaining why the surface slopes down to zero everywhere with increasing distance from the pole.


³³Of course, the mathematical plotting software cannot show a pole of truly infinite height, and so the pole has been truncated. This is why it appears to have a “flat” top.

One of the valuable insights provided by a three-dimensional pole-zero plot is the system's response to an input signal of constant magnitude and varying frequency. This is commonly referred to as the *frequency response* of the system, its graphical representation called a *Bode plot*. We may trace the Bode plot for this system by revealing a cross-sectional slice of the three-dimensional surface along the plane where $\sigma = 0$ (i.e. showing how the system responds to sinusoidal waves of varying frequency that don't grow or decay over time):



Only one-half of the pole-zero surface has been plotted here, in order to better reveal the cross-section along the $j\omega$ axis. The bold, red curve traces the edge of the transfer function surface as it begins at zero frequency (DC) to increasingly positive values of $j\omega$. The red trace is therefore the Bode plot for this low-pass filter, starting at a maximum value of 1 ($V_{out} = V_{in}$ for a DC input signal) and approaching zero as frequency increases.

As insightful as three-dimensional pole-zero plots are, they are laborious to plot by hand, and even with the aid of a computer may require significant³⁴ time to set up. For this reason, pole-zero plots have traditionally been drawn in a two-dimensional rather than three-dimensional format, from a “bird’s eye” view looking down at the s plane. Since this view hides any features of height, poles and zeros are instead located on the s plane by \times and \circ symbols, respectively. An example of a traditional pole-zero plot for our low-pass filter appears here:



Admittedly, this type of pole-zero plot is a lot less interesting to look at than a three-dimensional surface plotted by computer, but nevertheless contains useful information about the system. The single pole lying on the real (σ) axis tells us the system will not self-oscillate (i.e. $\omega = 0$ at the pole), and that it is inherently stable: when subjected to a pulse, its natural tendency is to decay to a stable value over time (i.e. $\sigma < 0$).


It should be noted that transfer functions and pole-zero plots apply to much more than just filter circuits. In fact, *any* physical system having the same “low-pass” characteristic as this filter circuit is describable by the same transfer function and the same pole-zero plots. Electric circuits just happen to be convenient applications because their individual component characteristics are so easy to represent as functions of s . However, if we are able to characterize the components of a different physical system in the same terms³⁵, the same mathematical tools apply.

³⁴My first pole-zero plot using the ePiX C++ mathematical visualization library took several hours to get it just right. Subsequent plots went a lot faster, of course, but they still require substantial amounts of time to adjust for a useful and aesthetically pleasing appearance.

³⁵A powerful mathematical technique known as a *Laplace Transform* does this very thing: translate any differential equation describing a physical system into functions of s , which may then be analyzed in terms of transfer functions and pole-zero plots.

5.7.2 Example: RC high-pass filter circuit

For our next exploratory example we will consider another simple filter circuit, this time comprised of a capacitor and a resistor, with the output signal taken across the resistor. As before, we may derive a transfer function by expressing V_{out}/V_{in} as a ratio of the resistor's impedance to the total series resistor-capacitor impedance (treating this as a voltage divider circuit):



$$\text{Transfer function} = \frac{V_{out}(s)}{V_{in}(s)} = \frac{R}{R + \frac{1}{sC}}$$

After writing this initial transfer function based on component impedances, we will algebraically manipulate it to eliminate compound fractions. This will aid our analysis of the circuit's DC response, zeros, and poles:

$$\begin{aligned} & \frac{R}{R + \frac{1}{sC}} \\ & \frac{R}{\frac{sRC}{sC} + \frac{1}{sC}} \\ & \frac{R}{\frac{1+sRC}{sC}} \\ & \frac{sRC}{1 + sRC} \end{aligned}$$

This transfer function allows us to calculate the “gain” of the system for any given value of s , which brings us to the next step of our analysis. Once again we will ask ourselves three questions about the transfer function:

1. How does this system respond when $s = 0$?
2. What value(s) of s make the transfer function approach a value of zero?
3. What value(s) of s make the transfer function approach a value of infinity?

In answer to the first question, we see that the transfer function is equal to zero when $s = 0$:

$$\frac{0RC}{1 + 0RC}$$

$$\frac{0}{1 + 0} = \frac{0}{1} = 0$$

Of course, a value of 0 for s means exposure to a steady DC signal: one that neither grows nor decays over time, nor oscillates. Therefore, this resistor-capacitor circuit will output zero voltage when exposed to a purely DC input signal. This makes conceptual sense when we examine the circuit itself: a DC input signal voltage means the capacitor will not experience any change in voltage over time, which means it will not pass any current along to the resistor. With no current through the resistor, there will be no output voltage. Thus, the capacitor “blocks” the DC input signal, preventing it from reaching the output. This behavior is exactly what we would expect from such a circuit, which any student of electronics should immediately recognize as being a simple *high-pass* filter: DC is a condition of zero frequency, which should be completely blocked by any filter circuit with a high-pass characteristic.

The answer to our first question is also the answer to the second question: “what value of s makes the transfer function equal to zero?” Here we see that it is only at a value of $s = 0$ that the entire transfer function’s value will be zero. Any other values for s – even infinite – yield non-zero results. In contrast to the last circuit (the resistor-inductor low-pass filter) this circuit exhibits a singular “zero” point in its transfer function: one specific location on the pole-zero plot where the function’s value diminishes to nothing.

When we consider the third question (“What value(s) of s make the transfer function approach a value of infinity?”) we proceed the same as before: by finding value(s) of s which will make the denominator of the transfer function fraction equal to zero. If we set the denominator portion equal to zero and solve for s , we will obtain the *pole* for the circuit:

$$1 + sRC = 0$$

$$sRC = -1$$

$$s = -\frac{1}{RC}$$


We know that both R and C are real numbers, not imaginary. This tells us that s will likewise be a real number at the pole. That is to say, s will be comprised of all σ and no ω . The fact that the value of σ is negative tells us the pole represents a condition of *exponential decay*, just the same as in the case of the resistor-inductor low-pass filter. As before, this means the circuit will produce an output voltage signal with no³⁶ input voltage signal when the rate of signal decay is $\sigma = -1/RC$.

Recall that the rate of decay in the s variable (σ) is nothing more than the reciprocal of the system’s time constant (τ). Thus, a rate of decay equal to $1/RC$ equates to a time constant $\tau = RC$,

³⁶As before, this counter-intuitive condition is possible only because the capacitor in this circuit has the ability to store energy. If the capacitor is charged by some previous input signal event and then allowed to discharge through the resistor, it becomes possible for this circuit to develop an output voltage even with short-circuited input terminals.


which as all electronics students know is how we calculate the time constant for any simple resistor-capacitor circuit.

Using a computer to plot a three-dimensional representation of this transfer function, we clearly see both the pole and the zero as singularities. Here I have assumed a $10\ \Omega$ resistor and a $0.2\ \text{F}$ capacitor to place the pole at the same location as with the low-pass filter circuit $s = -0.5 + j0$, for an equitable comparison:




Here we see both the pole at $s = -0.5 + j0$ and the zero at $s = 0 + j0$ quite clearly: the pole is a singular point of infinite height while the zero is a singular point of zero height. The three-dimensional surface of the transfer function looks like a rubber sheet that has been stretched to an infinite height at the pole and stretched to ground level at the zero.

As in the last example, we may re-plot the transfer function in a way that shows a cross-sectional view at $\sigma = 0$ in order to reveal the frequency response of this high-pass filter circuit:




Once again the bold, red curve traces the edge of the transfer function surface as it begins at zero frequency (DC) to increasingly positive values of $j\omega$. The red trace is therefore the Bode plot for this high-pass filter, starting at a minimum value of 0 ($V_{out} = 0$ for a DC input signal) and approaching unity (1) as frequency increases. Naturally, this is the type of response we would expect to see exhibited by a high-pass filter circuit.

A more traditional two-dimensional pole-zero plot for this circuit locates the pole with a “ \times ” symbol and the zero with a “ \circ ” symbol:



5.7.3 Example: LC “tank” circuit

Next, we will explore the transfer function for a *tank circuit*, comprised of a capacitor and an inductor. We will assume the use of pure reactances here with no electrical resistance or other energy losses of any kind, just to analyze an ideal case. The output voltage in this particular circuit will be taken across the inductor:



Writing the transfer function for this tank circuit is (once again) a matter of expressing the ratio between the output component’s impedance versus the total circuit impedance:

$$\text{Transfer function} = \frac{V_{out}(s)}{V_{in}(s)} = \frac{sL}{sL + \frac{1}{sC}}$$

Algebraically manipulating this function to eliminate compound fractions:

$$\begin{aligned} & \frac{sL}{sL + \frac{1}{sC}} \\ & \frac{sL}{\frac{sLsC}{sC} + \frac{1}{sC}} \\ & \frac{sL}{\frac{s^2LC}{sC} + \frac{1}{sC}} \\ & \frac{sL}{\frac{s^2LC+1}{sC}} \\ & \frac{sLsC}{s^2LC + 1} \\ & \frac{s^2LC}{s^2LC + 1} \end{aligned}$$

Note how this transfer function contains s^2 terms rather than s terms. This makes it a *second-order function*, which will yield very different results on the pole-zero plot than what we saw with either the resistor-inductor or resistor-capacitor filter circuits.

Asking ourselves the same three questions again:

1. How does this system respond when $s = 0$?
2. What value(s) of s make the transfer function approach a value of zero?
3. What value(s) of s make the transfer function approach a value of infinity?

In answer to the first question, we see that the transfer function is equal to zero when $s = 0$:

$$\frac{s^2 LC}{s^2 LC + 1}$$

$$\frac{0}{0 + 1} = \frac{0}{1} = 0$$

As with the RC low-pass filter, its response at DC also happens to be a “zero” for the transfer function. With a DC input signal, the output signal of this circuit will be zero volts.

In order to find poles for this transfer function, we must solve for values of s that will make the denominator term of the transfer function equal to zero:

$$s^2 LC + 1 = 0$$

$$s^2 LC = -1$$


$$s^2 = -\frac{1}{LC}$$

$$s = \sqrt{-\frac{1}{LC}}$$

$$s = \pm j\sqrt{\frac{1}{LC}}$$

Given the fact that both L and C are real, positive numbers, and therefore solving for s requires we take the square root of a negative real number, we see that the value of s must be imaginary. We also see here that there are *two* poles in this transfer function: one at $s = 0 + j\sqrt{\frac{1}{LC}}$ and another at $s = 0 - j\sqrt{\frac{1}{LC}}$.

Using a computer to plot a three-dimensional representation of this transfer function, we clearly see a single zero at $s = 0$ and two poles symmetrically positioned along the $j\omega$ axis. Here I have assumed a 0.2 F capacitor and a 5 H inductor for component values:



The two³⁷ poles located on the $j\omega$ axis (one at $s = 0 + j1$ and the other at $s = 0 - j1$) tell us the circuit is able to generate an oscillating output signal ($\omega = 1$ radian per second frequency) at constant magnitude ($\sigma = 0$) with no input signal. This is only possible because we have assumed a perfect capacitor and a perfect inductor with no energy losses whatsoever. If we charge up either or both of these components and then immediately short-circuit the input of the circuit to ensure $V_{in} = 0$, it will oscillate at its resonant frequency forever.

³⁷The *two* solutions for ω (one at +1 radian per second and the other at -1 radian per second) merely indicate the circuit is able to oscillate “forward” as well as “backward”. In other words, it is able to oscillate sinusoidally where the positive peak occurs at time $t = 0$ (+1 rad/sec) as well as oscillate sinusoidally where the negative peak occurs at time $t = 0$ (-1 rad/sec). We will find that solutions for s in general are symmetrical about the real axis, meaning if there is any solution for s requiring an imaginary number value, there will be *two* of them: one with a *positive* imaginary value and the other with a *negative* imaginary value.

Earlier we noted that the poles in this circuit were $s = 0 + j\sqrt{\frac{1}{LC}}$ and $s = 0 - j\sqrt{\frac{1}{LC}}$. In other words, its resonant frequency is $\omega = \sqrt{\frac{1}{LC}}$. Recalling that the definition for ω is radians of phasor rotation per second, and that there are 2π radians in one complete revolution (cycle), we can derive the familiar resonant frequency formula for a simple LC circuit:


$$\omega = \sqrt{\frac{1}{LC}}$$

. . . substituting $2\pi f$ for ω . . .


$$2\pi f = \sqrt{\frac{1}{LC}}$$

$$f = \frac{1}{2\pi\sqrt{LC}}$$

Taking a cross-section of this surface plot at $\sigma = 0$ to obtain a frequency response (Bode plot) of the LC tank circuit, we see the output of this circuit begin at zero when the frequency (ω) is zero, then the output peaks at the resonant frequency ($\omega = 1$ rad/sec), then the output approaches unity (1) as frequency increases past resonance:




A more traditional two-dimensional pole-zero plot for this circuit locates shows the zero and the two poles using “o” and “x” symbols:



5.7.4 Example: RLC band-pass filter circuit

For our next example circuit, we will add a resistor in series with the inductor and capacitor to explore its effects on the transfer function. Taking our output voltage across the resistor, we should expect to see *band-pass* filtering behavior from this circuit, with maximum voltage developing across the resistor at one frequency where the inductor's and capacitor's impedances cancel:



As usual, the transfer function for this circuit is the ratio between the output component's impedance (R) and the total series impedance, functioning as a voltage divider:

$$\text{Transfer function} = \frac{V_{out}(s)}{V_{in}(s)} = \frac{R}{R + sL + \frac{1}{sC}}$$

Algebraically manipulating this function to eliminate compound fractions:

$$\begin{aligned} & \frac{R}{R + sL + \frac{1}{sC}} \\ & \frac{R}{\frac{sRC}{sC} + \frac{s^2LC}{sC} + \frac{1}{sC}} \\ & \frac{R}{\frac{sRC + s^2LC + 1}{sC}} \\ & \frac{sRC}{sRC + s^2LC + 1} \end{aligned}$$

As with the pure tank circuit analyzed previously, we can see that this LC circuit exhibits a *second-order* transfer function because it contains an s^2 term. Running through our three questions again:

1. How does this system respond when $s = 0$?
2. What value(s) of s make the transfer function approach a value of zero?
3. What value(s) of s make the transfer function approach a value of infinity?

The answers to the first two questions are one and the same: the numerator of the transfer function will be zero when $s = 0$, this being the single *zero* of the function. Recalling that a condition of $s = 0$ represents a DC input signal (no growth or decay over time, and no oscillation), this makes perfect sense: the presence of the DC-blocking series capacitor in this circuit ensures the output voltage under steady-state conditions must be zero.

In answering the third question to identify any poles for this circuit, we encounter a more complicated mathematical problem than seen with previous example circuits. The denominator of the transfer function's fraction is a *second-degree polynomial* in the variable s . As you may recall from your study of algebra, any solution resulting in a polynomial having an over-all value of zero is called a *root* of that polynomial expression. Since this particular expression is found in the denominator of the transfer function where we know zero values mark poles of the system, and solutions for s resulting are roots of the polynomial, then roots of the expression $sRC + s^2LC + 1$ must mark the locations of the poles on the s plane.

A very useful algebraic tool for finding roots of a second-degree polynomial expression is the *quadratic formula*:

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

Where,

$ax^2 + bx + c$ is a polynomial expression

x is the independent variable of that polynomial expression

a is the coefficient of the second-degree (x^2) term

b is the coefficient of the first-degree (x) term

c is the coefficient of the zero-degree (constant) term

Reviewing the denominator of our transfer function again, we see that s is the independent variable, and so LC must be the “ a ” coefficient, RC must be the “ b ” coefficient, and 1 must be the “ c ” coefficient. Substituting these variables into the quadratic formula will give us a formula for computing the poles of this resistor-inductor-capacitor circuit:

$$s = \frac{-RC \pm \sqrt{(RC)^2 - 4LC}}{2LC}$$

Perhaps the most interesting part of this formula is what lies beneath the radicand (square-root) symbol: $(RC)^2 - 4LC$. This portion of the quadratic formula is called the *discriminant*, and its value determines both the number of roots as well as their real or imaginary³⁸ character. If the discriminant is equal to zero, there will be a single real root for our polynomial and therefore only one pole for our circuit. If the discriminant is greater than zero (i.e. a positive value), then there will be two real roots and therefore two poles lying on the σ axis (i.e. no imaginary $j\omega$ parts). If the discriminant is less than zero (i.e. a negative value), then there will be two complex roots for our polynomial and therefore two complex poles having both real and imaginary parts.

Let us consider for a moment what the sign of the discriminant means in practical terms. A pole that is purely real means a value for s that is all σ and no ω : representing a condition of growth

³⁸The only way to obtain a purely imaginary root for this polynomial is for the “ b ” coefficient to be equal to zero. For our example circuit, it means either R or C would have to be zero, which is impossible if both of those components are present and functioning. Thus, our RLC filter circuit will have either *real* poles or *complex* poles.

or decay but with no oscillation. This is similar to what we saw with the LR or RC low/high pass filter circuits, where the circuit in a state of discharge could generate an output signal even with its input terminals shorted to ensure no input signal.

If we have a positive value for the discriminant which yields *two* real poles, it means two different possible values for σ (rate of growth/decay) that could occur with no signal input to the circuit. This behavior is only possible with two energy-storing components in the circuit: a kind of *double time-constant* where different portions of the circuit discharge at different rates. The lack of any imaginary part within s means the circuit still will not self-oscillate.

If we have a negative value for the discriminant which yields *two* complex poles, it means two different values for s both having real and imaginary parts. Since the real part (σ) represents growth/decay while the imaginary part (ω) represents oscillation, complex poles tell us the circuit will be able to self-oscillate but not at a constant magnitude as with an ideal (lossless) tank circuit. In fact, intuition should tell us these complex poles must have negative real values representing decaying oscillations, because it would violate the Law of Energy Conservation for our circuit to self-oscillate with increasing magnitude.

Looking at the discriminant $(RC)^2 - 4LC$ we see that it is possible to push the circuit into any one of these three modes of operation merely by adjusting the value of R and leaving both L and C unchanged. If we wish to calculate the critical value of R necessary to produce a single real pole for any given values of L and C , we may set the discriminant equal to zero and algebraically solve for R as follows:

$$(RC)^2 - 4LC = 0$$

$$(RC)^2 = 4LC$$

$$RC = \sqrt{4LC}$$

$$RC = 2\sqrt{L}\sqrt{C}$$

$$R = \frac{2\sqrt{L}\sqrt{C}}{C}$$

$$R = \frac{2\sqrt{L}\sqrt{C}}{\sqrt{C}\sqrt{C}}$$

$$R = \frac{2\sqrt{L}}{\sqrt{C}} = 2\sqrt{\frac{L}{C}}$$

This critical value of R resulting in one real pole is the minimum amount of resistance necessary to prevent self-oscillation. If the circuit is operating at this point, it is said to be *critically damped*. Larger values of R will result in multiple real poles, where the circuit is said to be *over-damped*. Smaller values of R will permit some self-oscillation to occur, and the circuit is said to be *under-damped*.


Sometimes electrical engineers intentionally install resistors into circuits containing both inductance and capacitance for the express purpose of damping oscillations. In such cases, the resistor is called an *anti-resonance* resistor, because its purpose is to combat resonant oscillations that would otherwise occur as the inductive and capacitive elements of the circuit exchange energy back and forth with each other. If the engineer's intent is to install just enough resistance into the circuit to prevent oscillations without creating unnecessary time delays, then the best value of the resistor will be that which causes critical damping.

Recall that the subject of transfer functions, poles, and zeros applies to *any* linear system, not just AC circuits. Mechanical systems, feedback control loops, and many other physical systems may be characterized in the same way using the same mathematical tools. This particular subject of damping is extremely important in applications where oscillations are detrimental. Consider the design of an automobile's suspension system, where the complementary energy-storing phenomena of spring tension and vehicle mass give rise to oscillations following impact with a disturbance in the road surface. It is the job of the shock absorber to act as the "resistor" in this system and dissipate energy in order to minimize oscillations following a bump in the road. An under-sized shock absorber won't do a good enough job dissipating the energy of the disturbance, and so the vehicle's suspension will exhibit complex poles (i.e. there will be some lingering oscillations following a bump). An over-sized shock absorber will be too "stiff" and allow too much of the bump's energy to pass through to the vehicle frame and passengers. A perfectly-sized shock absorber, however, will "critically damp" the system to completely prevent oscillation while presenting the smoothest ride possible.

Likewise, oscillations following a disturbance are undesirable in a feedback control system where the goal is to maintain a process variable as close to setpoint as possible. An under-damped feedback loop will tend to oscillate excessively following a disturbance. An over-damped feedback loop won't oscillate, but it will take an excessive amount of time to converge back on setpoint which is also undesirable because this means more time spent off setpoint. A critically-damped feedback loop is the best-case compromise, where oscillations are eliminated and convergence time is held to a minimum.

In order to fully illustrate the characteristics of this circuit's transfer function, we will do so for three different resistor values: one where R yields critical damping (one real pole), one where R makes the circuit over-damped (two real poles), and one where R makes the circuit under-damped (two complex poles). We will use three-dimensional plotting to show the transfer function response in each case. To be consistent with our former tank circuit example, we will assume the same capacitor value of 0.2 Farads and the same inductor value of 5 Henrys. The resistor value will be modified in each case to create a different damping condition.

First, the critically-damped example, with a resistor value of 10 ohms:



As expected, a single zero appears at $s = 0$, and a single³⁹ pole at $s = -1 + j0$. Thus, this circuit has a decay rate of -1 time constants per second ($\tau = 1$ second) as seen when we use the quadratic formula to solve for s :

$$s = \frac{-RC \pm \sqrt{(RC)^2 - 4LC}}{2LC}$$

$$s = \frac{-(10)(0.2) \pm \sqrt{[(10)(0.2)]^2 - (4)(5)(0.2)}}{(2)(5)(0.2)}$$


$$s = \frac{-2 \pm \sqrt{0}}{2} = (-1 + j0) \text{ sec}^{-1}$$

Interestingly, only L and R determine the decay rate (σ , the real part of s) in the critically damped condition. This is clear to see if we set the discriminant to zero in the quadratic formula and look for variables to cancel:

$$s = \frac{-RC \pm \sqrt{0}}{2LC} = -\frac{R}{2L}$$

³⁹Or, one might argue there are two *repeated* poles, one at $s = -1 + j0$ and another at $s = -1 - j0$.

Next, we will plot the same transfer function with a larger resistor value (15 ohms) to ensure over-damping:



We clearly see *two* poles⁴⁰ centered along the σ axis in this plot, representing the two real roots of the transfer function's denominator. Again, we will use the quadratic formula to solve for these two values of s :

$$s = \frac{-RC \pm \sqrt{(RC)^2 - 4LC}}{2LC}$$

$$s = \frac{-(15)(0.2) \pm \sqrt{[(15)(0.2)]^2 - (4)(5)(0.2)}}{(2)(5)(0.2)}$$

$$s = \frac{-3 \pm \sqrt{(3)^2 - 4}}{2}$$


$$s = \frac{-3 + \sqrt{5}}{2} = (-0.382 + j0) \text{ sec}^{-1}$$

$$s = \frac{-3 - \sqrt{5}}{2} = (-2.618 + j0) \text{ sec}^{-1}$$

These real poles represent two different decay rates (time constants) for the over-damped circuit: a fast decay rate of $\sigma = -2.618 \text{ sec}^{-1}$ and a slow decay rate of $\sigma = -0.382 \text{ sec}^{-1}$, the slower of these two decay rates dominating the circuit's transient response over long periods of time.

⁴⁰The center of the pole farthest from the plot's origin actually lies outside the plotted area, which is why that pole appears to be vertically sliced. This plot's domain was limited to the same values (± 2) as previous plots for the sake of visual continuity, the compromise here being an incomplete mapping of one pole.

Next, we will plot the same transfer function with a smaller resistor value (5 ohms) to ensure under-damping:



We clearly see *two* poles once again, but neither of them are located on an axis. These represent two *complex* values for s describing the circuit's behavior with zero input. The imaginary ($j\omega$) part of s tells us the circuit has the ability to self-oscillate. The negative, real (σ) part of s tells us these oscillations decrease in magnitude over time. Using the quadratic formula to solve for these two poles:

$$s = \frac{-RC \pm \sqrt{(RC)^2 - 4LC}}{2LC}$$

$$s = \frac{-(5)(0.2) \pm \sqrt{[(5)(0.2)]^2 - (4)(5)(0.2)}}{(2)(5)(0.2)}$$

$$s = \frac{-1 \pm \sqrt{(1)^2 - 4}}{2}$$




$$s = \frac{-1 + \sqrt{-3}}{2} = (-0.5 + j0.866) \text{ sec}^{-1}$$

$$s = \frac{-1 - \sqrt{-3}}{2} = (-0.5 - j0.866) \text{ sec}^{-1}$$


The calculated ω value of 0.866 radians per second is slower than the 1 radian per second resonant frequency calculated for the pure tank circuit having the same L and C values, revealing that the damping resistor skews the “center” frequency of this RLC band-pass filter. The calculated σ value of -0.5 time constants per second (equivalent to a time constant of $\tau = 2$ seconds) describes the

rate at which the sinusoidal oscillations decay in magnitude. Here as well we see the under-damped decay rate ($\sigma = -0.5 \text{ sec}^{-1}$) is slower than the critically damped decay rate ($\sigma = -1 \text{ sec}^{-1}$).

If we compare two-dimensional pole-zero plots for each of the three resistor values in this RLC circuit, we may contrast the over-damped, critically-damped, and under-damped responses:



If we take cross-sectional plots of the transfer function at $\sigma = 0$ to show the frequency response of this RLC band-pass filter, we see the response become “sharper” (more selective) as the resistor’s value decreases and the poles move closer to the $j\omega$ axis. Electronic technicians relate this to the *quality factor* or Q of the band-pass filter circuit, the circuit exhibiting a higher “quality” of band-pass selection as the ratio of reactance to resistance increases:



Although each and every pole in a pole-zero plot is has the same (infinite) height, poles grow narrower when moved farther away from each other, and wider when closely spaced. As the resistance in this circuit decreases and the poles move farther away from each other and closer to the $j\omega$ axis, their widths narrow and the frequency response curve’s peak becomes narrower and steeper.

5.7.5 Summary of transfer function analysis

Here is a summary of some of the major concepts important to transfer function analysis:

- The **s variable** is an expression of growing/decaying sinusoidal waves, comprised of a real part and an imaginary part ($s = \sigma + j\omega$). The real part of s (σ) is the growth/decay rate, telling us how rapidly the signal grows or decays over time, with positive values of σ representing growth and negative values of σ representing decay. This growth/decay rate is the reciprocal of *time constant* ($\sigma = 1/\tau$), and is measured in reciprocal units of time (time constants per second, or sec^{-1}). The imaginary part of s ($j\omega$) represents the *frequency* of the sinusoidal quantity, measured in radians per second (sec^{-1}).
- An important assumption we make when analyzing any system's transfer function(s) is that the system is **linear** (i.e. its output and input magnitudes will be proportional to each other for all conditions) and **time-invariant** (i.e. the essential characteristics of the system do not change with the passage of time). If we wish to analyze a non-linear system using these tools, we must limit ourselves to ranges of operation where the system's response is approximately linear, and then accept small errors between the results of our analysis and the system's real-life response.
- For any linear time-invariant system (an "LTI" system), s is descriptive throughout the system. In other words, for a certain value of s describing the input to this system, that same value of s will also describe the output of that system.
- A **transfer function** is an expression of a systems' *gain*, measured as a ratio of output over input. In engineering, transfer functions are typically mathematical functions of s (i.e. s is the independent variable in the formula). When expressed in this way, the transfer function for a system tells us how much gain the system will have for any given value of s .
- Transfer functions are useful for analyzing the behavior of electric circuits, but they are not limited to this application. Any linear system, whether it be electrical, mechanical, chemical, or otherwise, may be characterized by transfer functions and analyzed using the same mathematical techniques. Thus, transfer functions and the s variable are general tools, not limited to electric circuit analysis.
- A **zero** is any value of s that results in the transfer function having a value of zero (i.e. zero gain, or no output for any magnitude of input). This tells us where the system will be least responsive. On a three-dimensional pole-zero plot, each zero appears as a low point where the surface touches the s plane. On a traditional two-dimensional pole-zero plot, each zero is marked with a circle symbol (\circ). We may solve for the zero(s) of a system by solving for value(s) of s that will make the *numerator* of the transfer function equal to zero, since the numerator of the transfer function represents the *output* term of the system.
- A **pole** is any value of s that results in the transfer function having an infinite value (i.e. maximum gain, yielding an output without any input). This tells us what the system is capable


of doing when it is not being “driven” by any input stimulus. Poles are typically associated with energy-storing elements in a passive system, because the only way an unpowered system could possibly generate an output with zero input is if there are energy-storing elements within that system discharging themselves to the output. On a three-dimensional pole-zero plot, each pole appears as a vertical spike on the surface reaching to infinity. On a traditional two-dimensional pole-zero plot, each pole is marked with a (\times) symbol. We may solve for the pole(s) of a system by solving for value(s) of s that will make the *denominator* of the transfer function equal to zero, since the denominator of the transfer function represents the *input* term of the system.

- **Second-order** systems are capable of self-oscillation. This is revealed by poles having imaginary values. These oscillations may be completely undamped (i.e. s is entirely imaginary, with $\sigma = 0$), in which case the system is able to oscillate forever on its own. If energy-dissipating elements are present in a second-order system, the oscillations will be damped (i.e. decay in magnitude over time).
- An **under-damped** system exhibits complex poles, with s having both imaginary ($j\omega$) frequency values and real (σ) decay values. This means the system can self-oscillate, but only with decreasing magnitude over time.
- A **critically damped** system is one having just enough dissipative behavior to completely prevent self-oscillation, exhibiting a single pole having only a real (σ) value and no imaginary ($j\omega$) value.
- An **over-damped** system is one having excessive dissipation, exhibiting multiple real poles. Each of these real poles represents a different decay rate (σ) or time constant ($\tau = 1/\sigma$) in the system. When these decay rates differ in value substantially from one another, the slowest one will dominate the behavior of the system over long periods of time.

5.8 Polyphase AC power


“Polyphase” means “many phases,” describing a form of AC electrical system where multiple sinusoidal voltages exist that are not in step with each other. The most common form of polyphase AC power in industry is *three-phase*, but all polyphase systems share similar traits. A good way to understand three-phase AC systems is to begin with an understanding of simpler, single-phase systems.

A simple *alternator* (AC generator) is nothing more than a magnetized rotor spinning between a pair of electromagnetic poles, the stationary wire coils (“stator windings”) developing AC voltage as the spinning rotor’s magnet passes by:




Note that the stator is comprised of two windings connected in series-aiding fashion, so that their respective AC voltages directly add. If each winding of this machine develops 60 volts, the series pair will develop 120 volts. This machine is properly called a *single-phase* alternator, because all its stator winding voltages are in-phase with each other.

A much more common alternator design uses three sets of stator poles, each one with its own winding pair, to generate three AC voltages phase-shifted from one another by 120° . The reason these three AC voltages are not in-phase with each other is precisely because the three stator poles are not physically aligned with each other, which means the magnetic poles of the spinning rotor will pass by each stator pole pair at different times:



Note that each pair of stator winding voltages directly add, because they are in phase with each other. In the example shown, each individual stator winding develops 60 volts, with each series-aiding pair (each “phase” of the alternator) developing 120 volts. However, the voltage appearing between different stator winding pairs is neither the simple sum ($120 + 120$) nor the simple difference ($120 - 120$) of each phase voltage. Rather, the phase-to-phase voltage is the *trigonometric sum* of two phasor quantities, spaced 120° apart. In the example shown, $120\angle 0^\circ + 120\angle 120^\circ = 207.85\angle 60^\circ$, which is approximately 208 volts. This machine is properly called a *three-phase* alternator. More specifically, this alternator is one with a *wye-connected* stator winding set, because the geometric configuration of the stator windings resembles that of the letter “Y”.


The following oscilloscope screenshot shows the output of a three-phase alternator, each channel of the oscilloscope connected across one phase of the alternator (e.g. Channel 1 across “A” and Neutral, Channel 2 across “B” and Neutral, and Channel 3 across “C” and Neutral):



In this oscillograph image we can clearly see a 120 phase shift between successive phase voltage waveforms. From this image we may also discern the *phase sequence* or *phase rotation* of the system: the sequential order in which the three phases reach their respective peak values. Phase sequence is determined by the direction of the alternator shaft’s rotation as well as the orientation of the stator phase windings. Surveying the three sine waves from left to right (the forward direction of time) on the oscillograph, we see channel 1 (Yellow) reaches its positive peak 120 degrees before channel 2 (Blue), which reaches its positive peak 120 degrees before channel 3 (Magenta). A common method of describing phase rotation is to list the phase letter labels in their order of sequence over time, in this case the triad *ABC*. It should be noted that a phase sequence of *ABC* is synonymous with *BCA* and also with *CAB*. This is easy to see if we write the letters in sequence for several rotations of the alternator, seeing that all three triads may be found within the longer sequence: *ABCABCABCABC*.

This trigonometric relationship between voltages in this “wye” connected alternator is clearly shown in the following *phasor diagram*. Solid phasors express the phase voltage for each of the three winding pairs in a three-phase alternator, while dashed lines express the line voltage between any two of the three output terminals on the alternator. The direction of each phasor expresses its phase shift in time, as the sine wave voltages produced by each of the three phase windings will be shifted apart from each other by 120 degrees. Here, each phase winding in the alternator happens to produce 120 VAC, resulting in a line voltage equal to the trigonometric sum of the 120 VAC phasors, 208 VAC:

Voltage phasor diagram for a wye-connected alternator



A simple way to calculate the side lengths of these non-right triangles is to use the *Law of Sines*, which simply states the ratio between the length of a triangle’s side and the sine of its opposite angle is constant, for *any* triangle:

$$\frac{\sin a}{A} = \frac{\sin b}{B} = \frac{\sin c}{C}$$

Applying this to the solution of the voltage between phases A and B (120 volts each, individually):

$$\frac{\sin 120^\circ}{V_{AB}} = \frac{\sin 30^\circ}{V_A} = \frac{\sin 30^\circ}{V_B}$$


$$\frac{\sin 120^\circ}{V_{AB}} = \frac{\sin 30^\circ}{120 \text{ volts}}$$

$$V_{AB} = (120 \text{ volts}) \left(\frac{\sin 120^\circ}{\sin 30^\circ} \right)$$

$$V_{AB} = 207.84 \text{ volts} \approx 208 \text{ volts}$$


The ratio $\frac{\sin 120^\circ}{\sin 30^\circ}$ is equal to the square-root of three ($\sqrt{3}$), and this factor frequently appears in three-phase electrical system calculations.

Some three-phase alternators have their phase windings connected differently, in a “delta” configuration rather than a “wye” configuration:




The phasor diagram for voltage in a delta-connected alternator is simpler, because phase voltage is exactly equal to line voltage, with each phase coil directly connected to a pair of line terminals:

Voltage phasor diagram for a delta-connected alternator



The following photograph shows the terminal connections on the rear side of a three-phase alternator intended to provide electrical power on a heavy-duty truck or boat. Note the three power terminals at the bottom of the photograph (with yellow-colored wire connectors attached), where the three-phase AC power is output. Also note the two copper “slip rings” I am pointing to with my finger, where DC power is conducted through stationary carbon “brushes,” through the copper rings on the shaft, to a winding on the spinning rotor to control its magnetization. The shaft of this alternator, normally coupled to the crankshaft of the vehicle’s engine by a V-belt, is on the far side of the alternator hidden from view of the camera:




Larger alternator units, such as those found in power plants, do not differ substantially in design from this small unit. Sets of stator windings around the circumference of the machine connected in either a wye or a delta configuration generate three-phase AC power, while a magnetized rotor spins at the center of the machine providing the changing magnetic field necessary to induce voltage in those stator windings. Permanent-magnet rotors are seldom used because they offer no way to control or regulate the machine’s output voltage during operation. Instead, the rotor has a single winding of its own, energized by DC supplied externally by a voltage regulator circuit. If more AC voltage is desired from the alternator, the regulator circuit sends more direct current to the rotor in order to strengthen its magnetic field. If less AC voltage is desired, the regulator sends less current to the rotor’s winding in order to weaken its magnetic field.

One of the advantages of three-phase AC power over single-phase AC power is a more constant delivery of electrical energy to the load over time. With three sets of stator windings at work, the combined effect is not unlike a triple bicycle with the three riders' legs staggered by 120° of rotation, or of a multi-cylinder automobile engine with the pistons staggered apart from each other: at any given time, at least one of the phases will be at or near its peak. Single-phase AC systems, by contrast, *pulsate* to a much greater extent.

In a single-phase AC circuit, energy transfer actually stops completely twice per cycle, when the current waveform passes through zero. This never happens in a polyphase system, because there are always other phases at non-zero current values when any one phase is at its zero-crossing point, owing to the fact that the phases in a polyphase power system are shifted from one another. This fact results in more efficient transfer of energy in AC power systems: a three-phase power system can actually transfer the same amount of power as a comparable single-phase power system using less metal in the power line conductors, despite the fact that a greater number of conductors is necessary (3 versus 2).

Another advantage of three-phase AC power is in applications where the AC is to be rectified into DC. The rectified output of a three-phase alternator is "smoother" than the rectified output of a single-phase alternator, with less ripple voltage to interfere with on-board electronic devices such as radios, because the phase-shifted currents overlap each other. This is why all automotive alternators are three-phase rather than single-phase machines.

A comparison of single-phase rectification versus three-phase rectification shows this clearly:




Not only is the ripple of the rectified DC voltage less in a three-phase system than in a single-phase system, but the frequency of that ripple is three times as great, making it easier to filter⁴¹ out of the DC power.

⁴¹Low-pass filter circuits are typically used to “smooth” the ripple from the output of a rectifier. The greater the frequency of this ripple voltage, the easier it is to filter from the DC (which has a frequency of zero). All other factors being equal, a low-pass filter attenuates higher-frequency components to a greater extent than lower-frequency components.


5.8.1 Delta and Wye configurations

Two basic forms of three-phase sources and loads appearing in industrial power systems are the *delta* (Δ) and the *wye* (or simply *Y*) configurations. A “wye-connected” device has its three elements joined at one common point in the middle as such:

Wye-connected
motor or generator




Wye-connected
resistive load




By contrast, a “delta-connected” device has its three elements joined as the sides of a triangle:

Delta-connected
motor or generator

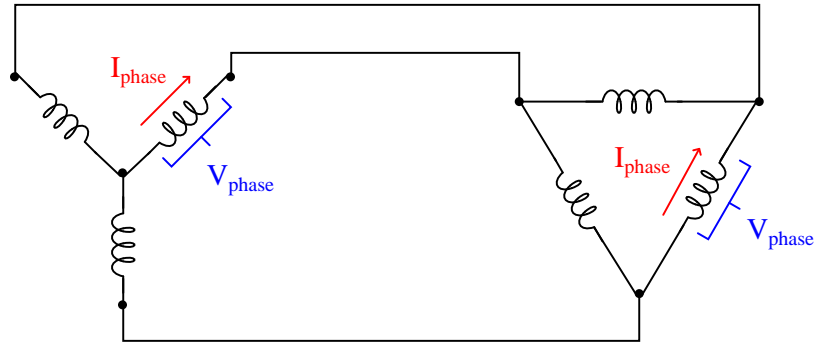


Delta-connected
resistive load




Each configuration has its own unique advantages and disadvantages in the larger context of a three-phase electrical power system. Either source type may connect to either load type (e.g. delta to wye, delta to delta, wye to delta, wye to wye) so long as the voltage and current ratings of all components are compatible.

The voltage appearing across the terminals of each element in a polyphase device is called the *phase voltage*, and the current through each element in a polyphase device is called the *phase current*:



Voltage appearing between any two of the connecting conductors (power lines) is called the *line voltage* of the polyphase system, and current through any of the connecting conductors (power lines) is called the *line current*:




Line and phase quantities relate to each other differently between delta devices and wye devices. Line voltage for a balanced⁴² wye device exceeds phase voltage by a factor of $\sqrt{3}$, while line current for a balanced delta device exceeds phase current by the same factor:


System type	Voltage	Current
Wye (Y)	$V_{line} = \sqrt{3} \times V_{phase}$	$I_{line} = I_{phase}$
Delta (Δ)	$V_{line} = V_{phase}$	$I_{line} = \sqrt{3} \times I_{phase}$

⁴²Here, the term “balanced” refers to a condition where all phase voltages and currents are symmetrically equal. Unbalanced conditions can and do exist in real polyphase power systems, but the degree of imbalance is usually quite small except in cases of component faults.

While it may be tempting to simply memorize these mathematical relationships, it is far better to *understand* why they are so. In a wye-connected device, line current must be equal to phase current because each power line is in series with each (respective) phase element, and we know that series-connected elements must share the same current. Likewise, line voltage must be equal to phase voltage in a delta-connected device because each power line pair connects in parallel fashion to each (respective) phase element, and we know that parallel-connected elements always share the same voltage:



Phase and line voltages are unequal in wye-connected devices, as are phase and line currents in delta-connected devices. In each of these cases, though, we may see once again by visual inspection that these line and phase quantities cannot be equal because the line quantities are the result of *two* joining phase quantities. In a wye network, line voltage is the series (phasor) sum of two phase voltages. In a delta network, line current is the parallel (phasor) sum of two currents summing at a node. If we know the system in question is balanced, however, we may be assured that the multiplying factor between these line and phase quantities will be the square-root of three ($\sqrt{3}$), therefore line voltage is $\sqrt{3}$ times greater than wye phase voltage, and line current is $\sqrt{3}$ times greater than delta phase current:




As an example of phase and line voltages in a wye-connected system, we see a great many three-phase industrial power circuits in the United States of the “480/277” volt configuration. These are a wye-connected devices exhibiting having phase voltages of 277 volts (each) and a balanced line voltage of 480 volts.

In the wye-connected system we see how the two phase voltages add in series (indirectly) to form the larger line voltage. In the delta-connected system we see how the larger line current splits up in parallel branches (indirectly) to form two smaller phase currents. The key to understanding these mathematical relationships is to recognize where the rules of series and parallel connections dictate the quantities be identical or different, and then all we need to remember is that if the two are different, the line quantity will be greater by a factor of $\sqrt{3}$.

5.8.2 Power in three-phase circuits

Suppose we have a 480/277 V wye-connected alternator supplying electrical power to a delta-connected load consisting of three 200 ohm resistive heating elements:



To begin our calculation of all electrical quantities in this circuit, we will apply Ohm's Law to the calculation of phase current at the load, since we already know phase voltage (480 volts) and phase resistance (200 ohms) there:

$$I_{phase(load)} = \frac{480 \text{ V}}{200 \Omega} = 2.4 \text{ A}$$

Now that we know phase current at the delta-connected load, we may calculate line current for the whole system by multiplying by the square-root of three:

$$I_{line} = (\sqrt{3})(2.4 \text{ A}) = 4.157 \text{ A}$$

This line current must be the same as the phase current in the wye-connected alternator, since line and phase currents are equal in wye-connected devices by virtue of their series connection. We already know the phase voltage of the alternator (277 volts) because that was given to us, but we could just as well calculate it from the line voltage of 480 volts as such:

$$V_{line} = (\sqrt{3})(V_{phase(source)})$$

$$V_{phase(source)} = \frac{V_{line}}{\sqrt{3}}$$

$$V_{phase(source)} = \frac{480 \text{ V}}{\sqrt{3}} = 277.1 \text{ V} \approx 277 \text{ V}$$


Tabulating all phase voltages and currents in our balanced system with a line voltage of 480 V and a line current of 4.157 A:

Quantity	Source	Load
V_{phase}	277 V	480 V
I_{phase}	4.157 A	2.4 A

Power for each of the three source or three load elements in this balanced system is simply the product of phase voltage and phase current ($P = IV$) because voltage and current are in-phase at each of the individual resistors. Expanding our table to include the power for each phase element:

Quantity	Source	Load
V_{phase}	277 V	480 V
I_{phase}	4.157 A	2.4 A
P_{phase}	1152 W	1152 W

Total generated power at the alternator (as well as total dissipated power at the resistive heater) is the simple sum of all three phase elements: 3456 watts in each case. No “square-root-of-three” factor is required in this calculation, because power (work over time) is not a phasor quantity⁴³. The Law of Energy Conservation demands that all power be accounted for, and thus three resistors dissipating 1152 watts each must be together dissipating 3456 watts total:



In the interest of convenience, though, it is helpful to have a formula to calculate power in a balanced three-phase system knowing just the line voltage and current rather than phase voltages and currents. You will note how the simple product of line voltage (480 V) and line current (4.157 A) does *not* yield total power (3456 W) in this system. We may develop a proper formula for calculating total power from line voltage and current by beginning with the formula described in calculating total power from the power of each resistor at the delta-connected load:

$$P_{total} = (3)(P_{phase})$$

$$P_{total} = (3)(I_{phase})(V_{phase})$$

⁴³You may recall from basic physics that while force and displacement are both vector quantities (having direction as well as magnitude), *work* and *energy* are not. Since power is nothing more than the rate of work over time, and neither work nor time are vector quantities, power is not a vector quantity either. This is closely analogous to voltage, current, and power in polyphase electrical networks, where both voltage and current are phasor quantities (having phase angle “direction” as well as magnitude) but where power merely has magnitude. We call such “directionless” quantities *scalar*. Scalar arithmetic is simple, with quantities adding and subtracting directly rather than trigonometrically.

We may substitute I_{line} for I_{phase} and V_{line} for V_{phase} in this equation if we properly relate them for the delta⁴⁴ connection. While $V_{line} = V_{phase}$ in a delta configuration, $I_{phase} = \frac{I_{line}}{\sqrt{3}}$:

$$P_{total} = (3) \left(\frac{I_{line}}{\sqrt{3}} \right) (V_{line})$$

We may consolidate the two constants in this formula (3 and $\sqrt{3}$) by re-writing the number 3 as the product $\sqrt{3}\sqrt{3}$, then canceling one of these with the $\sqrt{3}$ in the denominator:

$$P_{total} = (\sqrt{3}\sqrt{3}) \left(\frac{I_{line}}{\sqrt{3}} \right) (V_{line})$$

$$P_{total} = (\sqrt{3})(I_{line})(V_{line})$$

As a test, we may check to see that this new formula accurately calculates the total power of our balanced three-phase system:

$$P_{total} = (\sqrt{3})(4.157 \text{ A})(480 \text{ V})$$


$$P_{total} = 3456 \text{ W}$$

⁴⁴We end up with the same final result if we substitute line quantities in a wye-connected system, too. Instead of $V_{line} = V_{phase}$ and $I_{phase} = \frac{I_{line}}{\sqrt{3}}$ in the delta connection we have $I_{line} = I_{phase}$ and $V_{phase} = \frac{V_{line}}{\sqrt{3}}$ in the wye connection. The end-result is still $P_{total} = (\sqrt{3})(I_{line})(V_{line})$ based on line quantities.

5.8.3 Grounded three-phase circuits


So far all the three-phase configurations shown have been *ungrounded*: that is, none of the terminals or conductors have any direct connection to earth ground. While it is possible (and practical in many cases) to use polyphase power without an explicit earth ground connection, it is not always the safest. Establishing a firm connection to earth limits the voltage which may develop between any non-grounded (“hot”) conductor and earth ground. This is especially important in power systems with overhead lines, where lightning strikes may dramatically elevate common-mode voltages in the system.

In wye-connected systems, the natural point to ground is the center of the “Y”, like this:



The three “hot” terminals of the source are typically labeled “A”, “B”, and “C”, while the grounded point is referred to as the “Neutral” (N). The voltage measured between any two of the “hot” terminals (A to B, B to C, or A to C) will be $\sqrt{3}$ times more than the voltage measured between any “hot” terminal and the neutral (A to N, B to N, or C to N). Common voltages for 4-wire wye-connected systems include 208/120 and 480/277.


The existence of dual voltage levels in a center-grounded wye system enables the use of loads with different voltage ratings. For example, in a 208/120 wye system, a three-phase motor with windings rated for 208 volts would draw power from the three “hot” conductors directly, while light bulbs rated for 120 volts would connect between any “hot” conductor and the neutral:



A good practice in such systems is to equally spread the 120 volt loads among the three phases, so that (ideally) the phase loading on the source will be nicely balanced when all loads are operating. If the loading is perfectly balanced, in fact, the neutral conductor will carry no current at the point where it connects to the center of the wye.

Note the use of fuses on all the “hot” load conductors, but no fuses on any “neutral” (grounded) conductor. The reason for this is so that in the event of a fuse blowing, only the hot conductor(s) will be disconnected from the load, leaving the neutral conductor connected and thereby maintaining the lowest possible potential at the load relative to earth ground.

In delta-connected systems, there is no “natural” point to connect to earth ground as there is in wye-connected systems. The most common grounding configuration for a delta-connected source is shown here, sometimes called the “high-leg” connection. Here, one of the source’s phase coils is center-tapped to provide a connection point for earth ground:



This configuration yields *three different* voltages available to power loads. If the phase voltage of the delta-connected source is 240 volts, the three available voltages are:

- 240 volts (between A-B, B-C, and A-C)
- 120 volts (between B-N and C-N)
- 208 volts (between A-N)⁴⁵

A disadvantage of this configuration is that the lower-voltage loads cannot be balanced among all three phase coils of the source as they can in wye-connected systems. Any single-phase loads (those connected between any one “hot” conductor and the neutral conductor) inevitably place more burden on the B-C phase coil than the other two phase coils. However, this imbalance is often negligible in industrial settings where three-phase loads are predominant and single-phase loads are few (and low-wattage).


⁴⁵A colorful term for this odd voltage is *bastard voltage*.

5.8.4 Symmetrical components


Balanced three-phase networks are relatively easy to calculate quantities in, as each line and phase variable is symmetrical. For instance, if we happen to know one of the phase voltages in a balanced Wye-connected power system component is 2400 volts, then we may safely conclude the other two phase voltages are 2400 volts as well. The voltage between any two lines in this same system is guaranteed to be $\sqrt{3}$ larger than any phase voltage: $2400\sqrt{3} \approx 4160$ volts.

Calculations become much more complex, however, in *unbalanced* three-phase networks. One of the assumptions we must discard in an unbalanced network is the simple factor of $\sqrt{3}$ relating phase and line quantities: while $V_{line} = V_{phase}\sqrt{3}$ in a balanced Wye-connected system, it is not necessarily true in an *unbalanced* Wye-connected system. Compare these phasor diagrams of a balanced versus unbalanced 2400/4160 volt Wye systems, the unbalanced system being representative of a three-phase generator with a fault in one of the phase windings causing that phase voltage to be substantially less than it ought to be:

Balanced Wye-connected system



Unbalanced Wye-connected system
(phase B voltage $1/10$ of other phases)



In fact, the existence of faults in three-phase power systems is the primary reason for considering unbalanced systems, since the vast majority of three-phase electrical components are expressly designed to be balanced. If power system engineers and technicians are to analyze faults, they must have some means of quantifying unbalanced system conditions.

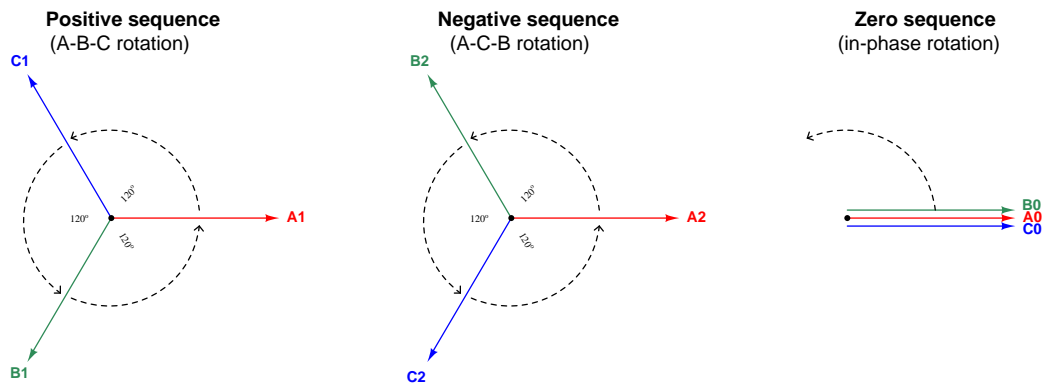
A breakthrough in mathematical analysis for unbalanced three-phase electrical circuits came in 1913 from a man named Charles Legeyt Fortescue, who presented his discovery in a paper entitled, “Method of Symmetrical Co-ordinates Applied to the Solution of Polyphase Networks” after doing research on AC induction motors operating under unbalanced conditions. His discovery condenses to the fact that any set of phasors describing conditions of voltage or current in a three-phase network, no matter how unbalanced and asymmetrical they may be, are mathematically equivalent to the sum of three unique phasor sets with different rotations. Thus, it is possible to mathematically decompose an unbalanced phasor triad into a multiple sets of balanced phasor triads, each of those

balanced sets being relatively easy to analyze on its own because the simple rules of symmetrical networks (e.g. the $\sqrt{3}$ factor between phase and line quantities) still apply.

Fortescue's breakthrough is reminiscent of Jean Baptiste Joseph Fourier's discovery roughly 100 years prior that any periodic waveform, no matter its shape, is mathematically equivalent to a summation of pure sinusoidal waveforms of harmonic frequencies (called a Fourier Series). In both cases, we see there is a mathematical equivalence between one entity that is ugly and asymmetrical, and a set of pure and symmetrical entities that are easier to deal with on mathematical terms.

In Fortescue's model, which is widely known under the title of *symmetrical components*, any set of three phasors describing voltage or current in a three-phase system is equivalent to the summation of three different three-phasor sets:


- One set of three phasors rotating in the normal A-B-C direction of the power system, called the *positive sequence*. By convention, this sequence is designated by the number 1.
- One set of three phasors rotating in the reverse direction (A-C-B)⁴⁶, called the *negative sequence*. By convention, this sequence is designated by the number 2.
- One set of three phasors all pointed in the same direction, having no sequence at all, called the *zero sequence*. By convention, this sequence is designated by the number 0.




Fortescue's discovery was that we may synthesize any possible three-phasor set simply by superimposing these positive, negative, and/or zero sequence phasor sets at the appropriate magnitudes and phase shifts. Each positive sequence, negative sequence, and zero sequence set is perfectly balanced, although they must usually differ in magnitude from each other in order for the summation of all three to equal the real-world unbalanced set.

⁴⁶If you are having difficulty seeing the A-B-C or A-C-B rotations of the positive and negative sequences, you may be incorrectly visualizing them. Remember that the phasors (arrows) themselves are rotating about the center point, and you (the observer) are stationary. If you imagine yourself standing where the tip of each "A" phasor now points, then imagine all the phasor arrows rotating counter-clockwise, you will see each phasor tip pass by your vantage point in the correct order.

Another way of stating this is to say that the actual voltages and currents in a three-phase power system, no matter how unbalanced they may be from each other, are really equivalent to multiple sets of voltages and currents existing in the circuit simultaneously, each with its own rotational sequence and perfectly balanced magnitude. Our hypothetical three-phase generator with the faulted phase winding, therefore, is really equivalent to *three* healthy generators with their windings wired together to be series-aiding: one generator spinning in the normal direction, the next generator spinning the wrong direction, and the last generator a single-phase unit with three in-phase windings. The combined sum of these three generators' outputs would create the exact same patterns of voltages and currents as the one faulted-winding generator:



Let us explore this in more detail by means of a concrete example. Taking the faulted-winding generator whose phasor voltage diagram was shown earlier (phase voltages V_A and V_C at 2400 volts each, but phase voltage V_B at only 240 volts), our task is to see how to decompose this unbalanced set of voltage phasors into three balanced sets of phasors (positive-, negative-, and zero-sequence):



Somehow, we must determine the specific pattern of balanced positive-, negative-, and zero-sequence phasor sets that will be equivalent to the unbalanced phasor set describing our faulted generator. Once we determine these symmetrical components, we will then sum them together to demonstrate how the respective phasors do indeed add up to the original (unbalanced) phasor diagram.

Thankfully, symmetrical component analysis provides us with a set of equations for calculating the “A” phasor of each sequence⁴⁷. Please note that each variable in each of the following equations is a *phasor* quantity, possessing both a magnitude and a phase angle. Also note that these equations apply just as well to calculations of *current* as they do to voltage:

$$V_{a1} = \frac{1}{3}(V_a + V_b\angle + 120^\circ + V_c\angle + 240^\circ) \quad \text{Positive-sequence phasor } A$$

$$V_{a2} = \frac{1}{3}(V_a + V_b\angle + 240^\circ + V_c\angle + 120^\circ) \quad \text{Negative-sequence phasor } A$$

$$V_{a0} = \frac{1}{3}(V_a + V_b + V_c) \quad \text{Zero-sequence phasor } A$$

The “+120°” and “+240°” references deserve further explanation⁴⁸. What we are saying here is that the calculations involve shifting the angle of the unbalanced system’s phasor by either +120 degrees or +240 degrees when calculating the phasors for the positive and negative sequences. Performing the calculations:

Given phasors in unbalanced system:

$$V_a = 2400\angle 0^\circ$$

$$V_b = 240\angle 240^\circ$$

$$V_c = 2400\angle 120^\circ$$

$$V_{a1} = \frac{1}{3}[2400\angle 0^\circ + 240\angle(240^\circ + 120^\circ) + 2400\angle(120^\circ + 240^\circ)] = 1680\angle 0^\circ$$

$$V_{a2} = \frac{1}{3}[2400\angle 0^\circ + 240\angle(240^\circ + 240^\circ) + 2400\angle(120^\circ + 120^\circ)] = 720\angle -60^\circ$$




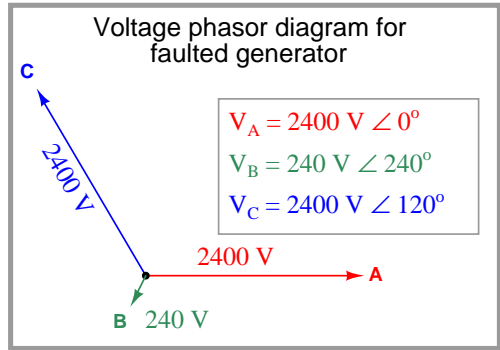
$$V_{a0} = \frac{1}{3}[2400\angle 0^\circ + 240\angle 240^\circ + 2400\angle 120^\circ] = 720\angle 60^\circ$$

Recalling that our faulted generator is mathematically equivalent to three healthy generators with their respective phase windings connected in series, these solutions tell us that the equivalence is one generator spinning the correct direction at 1680 volts per phase, connected to another generator spinning the wrong direction as well as phase-shifted -60 degrees at 720 volts per phase, connected to a single-phase generator phase-shifted $+60$ degrees at 720 volts per phase. These three healthy generators, each one producing a symmetrical (balanced) set of voltages, will together behave the same as our one 2400 volt phase generator with the faulted “B” winding.

⁴⁷It is good to remember that each of the symmetrical components is perfectly balanced (i.e. the “b” and “c” phasors each have exactly the same magnitude as the “a” phasor in each sequential set), and as such each of the phasors for each symmetrical set will have exactly the same magnitude. It is common to denote the calculated phasors simply as V_1 , V_2 , and V_0 rather than V_{a1} , V_{a2} , and V_{a0} , the “a” phasor implied as the representative of each symmetrical component.


⁴⁸A “shorthand” notation commonly seen in symmetrical component analysis is the use of a *unit phasor* called a , equal to $1\angle 120^\circ$. Multiplying any phasor quantity by a shifts that phasor’s phase angle by +120 degrees while leaving its magnitude unaffected. Multiplying any phasor quantity by a^2 shifts that phasor’s phase angle by +240 degrees while leaving its magnitude unaffected. An example of this “ a ” notation is seen in the following formula for calculating the positive sequence voltage phasor: $V_{a1} = \frac{1}{3}(V_a + aV_b + a^2V_c)$

Graphically, the decomposition of the original unbalanced phasor diagram into symmetrical components is as follows:




The graphical proof that these three symmetrical component phasor sets do indeed add up to be equivalent to the original (unbalanced) phasor diagram of the faulted generator is as follows:

Summation of symmetrical phasor sets (+, -, and 0 sequences)



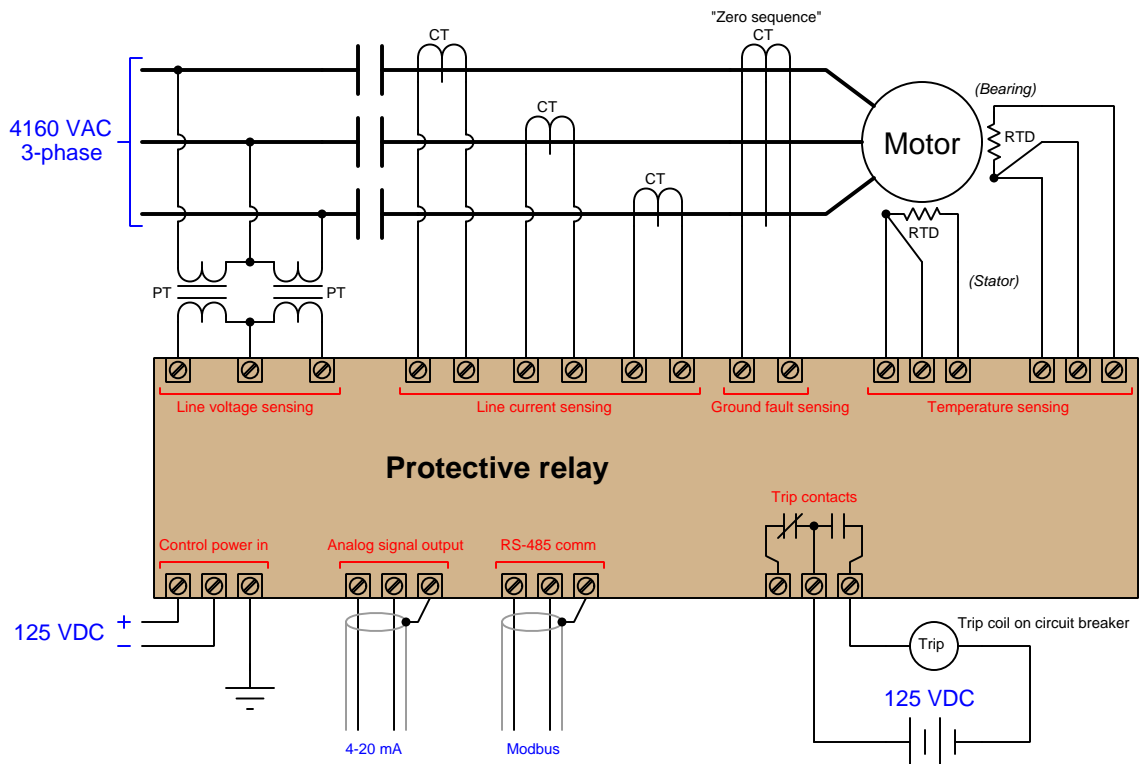
Phasor diagram for (original) faulted generator



In a perfectly balanced three-phase power system, the only sequence in existence is the positive sequence (1). Any and all measurements of voltage and current, therefore, represent the positive sequence of the system. Since the magnitudes of the negative- and zero-sequence components are zero. If an imbalance occurs for any reason – from unbalanced single-phase loads on a three-phase power system to actual faults (opens, shorts) – negative-sequence and/or zero-sequence components arise.

While negative-sequence quantities cannot be measured directly by any voltmeter or ammeter, zero-sequence components can. In Y-connected systems, zero-sequence current manifests itself as current through the neutral conductor, while zero-sequence voltage manifests itself as voltage between the center of an ungrounded Wye and ground. Zero-sequence current manifests in Delta-connected systems as current *circulating* in the phase elements of the source or load.

The following illustration shows a medium-voltage (4160 volt) electric motor circuit equipped with a *protective relay* to halt the motor in the event its electrical or thermal parameters indicate the potential for damage:




Note the "zero sequence" current transformer (CT) encircling all three lines conducting power to the motor. This large CT magnetically senses the instantaneous sum of currents into and out of the motor, which should be zero under normal operating conditions. If, however, a ground fault develops within the motor, the instantaneous sum of currents will become non-zero, causing the "zero sequence CT" to register some current. If this zero-sequence current becomes large enough,

the protective relay will command the contactor feeding power to the motor to trip (open) in order to protect the motor against further damage.

This exact same principle is used in household “GFCI” (Ground Fault Current Interruptor) receptacles and circuit breakers required by the National Electrical Code to be used in “wet” areas of a residence such as bathrooms: a small current transformer senses the net sum of current through the “hot” and “neutral” conductors. If this net current measurement is anything but zero, the interruptor contacts trip, ceasing the flow of power to the load and thereby protecting the person using that load from injury due to ground-fault current passing through their body:

Ground-Fault Current Interruptor (GFCI)




Any current leaking from the load to earth ground is sensed as an imbalance of current by the CT

5.9 Phasor analysis of transformer circuits

Three-phase AC is the dominant mode of electric power generation and distribution in modern societies. Analyzing voltages and currents in these systems can be quite complex, but phasor representation is a tool useful for simplifying that analysis by representing voltage and current quantities in graphical form. This section of the book illustrates the use of phasors as an analytical tool through examples. This section begins with single-phase phasor analysis and progresses to three-phase phasor analysis.

Extensive use of an imaginary instrument I call the “phasometer” helps relate phasor arrow direction to real circuit connections. For a detailed introduction to this instrument, see section 5.5.2 beginning on page 382. Suffice it to say, a “phasometer” is nothing more than a small synchronous AC motor with an arrow-shaped rotor, which when coupled with a strobe light set to flash whenever some reference voltage reaches its positive peak, graphically shows the relative phase of the measured quantity with respect to that reference voltage. In other words, a phasometer/strobe combination shows a “snapshot” view of the waveform’s position at that point in time when the reference wave is at its positive peak. This is precisely how phasors are drawn: arrows with angles displaced from 0° by an amount proportional to the phase shift between that quantity and some reference waveform.

Here is what a triplet of phasometers would indicate if connected to a three-phase power bus, with phase A used as the reference and given an ABC phase sequence (also called “phase rotation”):




In an ABC phase sequence, phase B’s voltage lags behind phase A’s voltage by 120° , while phase C’s voltage leads ahead of phase A’s voltage by 120° (or, one could say phase C lags a full 240° behind phase A).

5.9.1 Phasors in single-phase transformer circuits


Phasors are to AC circuit quantities as polarity is to DC circuit quantities, and for this reason they are particularly useful in analyzing transformer circuits because transformers by their very nature allow polarities to be mixed and matched.

In the “Basic principles” discussion of the “Transformers” section of this book (subsection 5.4.1 beginning on page 362), we used standard DC polarity marks (+ and −) to denote the polarity of voltage at an instant in time where the sinusoidal waves for each winding reached their peaks. Although it is often important to show the polarity of transformer windings relative to each other, the use of DC notation for this purpose can be misleading. An alternative symbology places a dot, or square, or some other distinguishing mark at one end of each transformer winding to show common polarity at any instant in time. The marked terminal of a transformer winding is called the “polarity” terminal, while the unmarked terminal is called the “nonpolarity” terminal:

Transformer polarity symbols




The marks should be interpreted in terms of *voltage polarity*, not current. To illustrate using a “test circuit⁴⁹” feeding a momentary pulse of DC to a transformer:



If the battery polarity were reversed, the “polarity” (dot) terminal of each winding would be *negative* with respect to the “nonpolarity” terminal at the moment in time when the switch is pressed. That is to say, the “polarity” marks merely show which terminals will share the same instantaneous voltage polarity at any given point in time – not which terminal is always positive.


⁴⁹The battery-and-switch test circuit shown here is not just hypothetical, but may actually be used to test the polarity of an unmarked transformer. Simply connect a DC voltmeter to the secondary winding while pressing and releasing the pushbutton switch: the voltmeter’s polarity indicated while the button is pressed will indicate the relative phasing of the two windings. Note that the voltmeter’s polarity will reverse when the pushbutton switch is released and the magnetic field collapses in the transformer coil, so be sure to pay attention to the voltmeter’s indication *only* during the instant the switch closes!

Consider these examples, where the polarity of a transformer's secondary winding is changed to illustrate the effect on AC phase shift from primary to secondary. In the left-hand example, the "polarity" terminal of the transformer's secondary winding exhibits a phase angle of 0° (with respect to ground) at the same point in time when the primary winding's voltage is at 0° : the two winding voltages are in-phase with each other. In the right-hand example, the secondary winding polarity has been reversed, the effect being a 180° phase shift in voltage. A pair of phasometers in each example, both synchronized to the positive peak of the primary voltage, shows the phase relationships clearly:




A useful way to analyze phase shifts in transformer circuits is to sketch separate phasor diagrams for the primary and secondary windings, and use the polarity marks on each winding to reference each phasor's position on the diagram. We begin this analysis by first sketching a phasor for any voltage of which the phase angle is given to us. Since transformer windings share common phase shift by virtue of those windings sharing the same magnetic field, the phasor representing the other winding's voltage must possess the same angle. The only distinction is the length of each phasor (representing primary and secondary voltage magnitudes) and its origin (starting point on the diagram).

Applying this analysis to the in-phase transformer circuit example. We will regard the secondary winding's voltage as a phasor that is half the length of the primary winding's voltage given the 2:1 turns ratio, but oriented at the same angle:



The result is a pair of phasors both at 0° angles. Note the placement of dots on each phasor to mark the relationship of the phasor to each winding polarity. The center of the diagram where the real and imaginary axes intersect is the point representing zero potential (ground). The length of each phasor tells us the magnitude of each waveform, in this case 120 volts and 60 volts AC, respectively. The “nonpolarity” end of each phasor line must be located at the origin because the nonpolarity terminal of each transformer winding is connected to ground.


Now consider the same analysis applied to the out-of-phase transformer circuit:



The result of this transformer connection is a pair of horizontal phasors, angled 180° apart from each other. Note how the secondary phasor still has its polarity mark on the right-hand end, but now that end of the phasor is anchored at the diagram's origin because now the "polarity" (dot) terminal of the transformer's secondary winding is connected to ground. Following our rule of keeping primary and secondary phasors at equal angles, the only difference made by grounding the opposite terminal of the secondary winding is the reference point of the secondary phasor. The secondary phasor must remain locked at the same angle as the primary phasor, but its position on the diagram is determined by where the transformer winding connects to ground.

The following analyses illustrate the relationships between phase angles, transformer polarity, circuit connections, and phasor diagrams.


Here we have an “autotransformer⁵⁰” connected to *boost* the primary voltage. The V_P and V_S phasors representing voltage across the primary and secondary windings must have common angles (i.e. both horizontal phasors), but the secondary phasor is “stacked” on the end of the primary phasor to create a longer phasor representing the voltage V_{out} with reference to ground. We can tell the phasors will be stacked polarity-end-to-nonpolarity-end because that is how the two transformer windings are electrically connected to each other:



The result of this “boost” configuration is an output voltage that is the direct sum of the primary and secondary winding voltages: 180 volts with a 2:1 transformer ratio and 120 volt source. If the transformer ratio were 4:1 instead of 2:1, the output would be 150 volts ($120 \text{ volts} + \frac{120}{4} \text{ volts}$).

⁵⁰An *autotransformer* is any transformer configuration where the primary and secondary windings are connected rather than galvanically isolated from each other as is typical.

Here we have another autotransformer, this time connected to *buck* the primary voltage. The V_P and V_S phasors still share common angles, but now they are “stacked” differently: polarity-end-to-polarity-end. This is challenging to see on the second phasor diagram because the V_P and V_S overlap each other, sharing a common “polarity” (dot) end point:

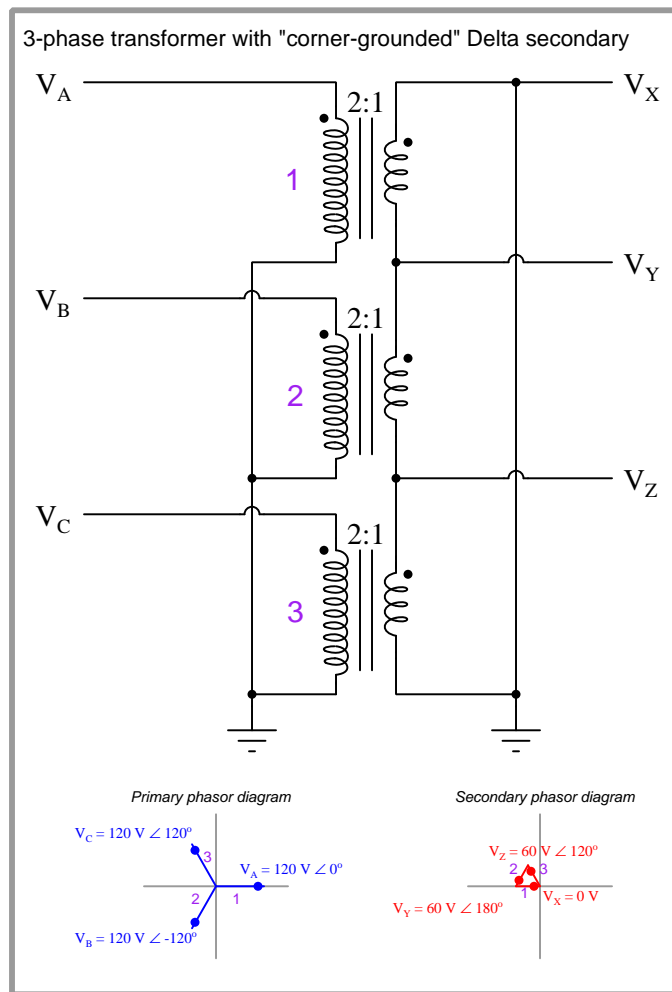


The result of this “buck” configuration is a V_{out} that is the *difference* between V_P and V_S rather than being the sum of the two as was the case with the “boost” configuration. If the transformer ratio were 4:1 instead of 2:1, the output of this buck autotransformer would be 90 volts ($120 \text{ volts} - \frac{120}{4} \text{ volts}$).

5.9.2 Phasors in three-phase transformer circuits


In the following three-phase transformer circuit, each transformer has been labeled with a number, and each pair of phasors representing primary and secondary winding voltages for each transformer has that same identifying number written next to it for convenient reference. Numbering the phasors makes it clear which phasor(s) belong to which transformer.

As usual, the primary and secondary phasors for each transformer share the same angle, but the secondary phasors are “stacked” polarity-end-to-nonpolarity-end because that is precisely how the secondary windings are electrically connected to each other. Assuming an ABC phase sequence with V_A being the reference phasor defining 0° phase angle:



5.10 Transmission lines

Many years ago, when I was first learning about electricity, I happened to discover a length of coaxial television cable with the rating “50 ohm” printed on the jacket. Having my own ohmmeter, and wanting to know what the “50 ohm” label referred to, I tried measuring the resistance of this cable to see where it would exhibit 50 ohms of resistance. First I tried measuring between the two open-ended wires (the center conductor and the braided shield conductor), where my meter registered infinite resistance (no continuity). This is what I expected a normal cable would read when connected to nothing else. Next I tried connecting my meter from one end of the cable to the other: between both ends of the center conductor my meter registered 0 ohms, and between both ends of the shield conductor my meter also read 0 ohms. In fact, this is all I was ever able to measure with my ohmmeter: either 0 ohms or infinite ohms (open), no matter what I tried. I was confounded – how could this piece of cable possibly be called a “50 ohm” cable if I could not measure 50 ohms anywhere on it?



It was not until years later that I came across a reference in an old book to something called *surge impedance*, describing how lengths of cable responded to short-duration electrical pulses, that I finally understood the “50 ohm” rating of that coaxial cable. What I failed to discover with my ohmmeter – in fact, what I never could have even seen with a plain ohmmeter – is that the cable did indeed act like a 50 ohm resistor, *but only for extremely short periods of time*. This is an aspect of electrical theory I was completely unprepared to comprehend at that time, knowing only how DC circuits functioned. This is the subject we are about to explore now.

In basic DC electrical theory, students learn that open circuits always drop the full applied voltage and cannot have electric currents anywhere in them. Students likewise learn that short

circuits drop negligible voltage while conducting full electrical current. They also learn that the effects of electricity occur instantaneously throughout the circuit: for example, if a functioning circuit suffers an “open” fault, current through that branch of the circuit halts immediately and everywhere. These basic rules of electricity are extremely useful for DC circuit analysis, indeed even essential for troubleshooting DC circuits, but they are not 100% correct. Like many of the scientific principles one initially learns, they are only approximations of reality.

The effects of electricity do *not* propagate instantaneously throughout a circuit, but rather spread at the speed of light (approximately 300,000,000 meters per second). Normally we never see any consequences of this finite speed because everything happens so fast, just as I never saw my ohmmeter register 50 ohms when I tried to test that coaxial cable many years ago. However, when we test a circuit on a time scale of nano-seconds (*billionths* of a second!), we find some very interesting phenomena during the time electricity propagates along the length of a circuit: open circuits can indeed (temporarily) exhibit current, and short circuits can indeed (temporarily) exhibit voltage drops.

When a pulse signal is applied to the beginning of a two-conductor cable, the reactive elements of that cable (i.e. capacitance between the conductors, inductance along the cable length) begin to store energy. This translates to a current drawn by the cable from the source of the pulse, as though the cable were acting as a (momentarily) resistive load. If the cable under test were infinitely long, this charging effect would never end, and the cable would indeed behave exactly like a resistor from the perspective of the signal source. However, real cables (having finite length) do stop “charging” after some time following the application of a voltage signal at one end, the length of that time being a function of the cable’s physical length and the speed of light.

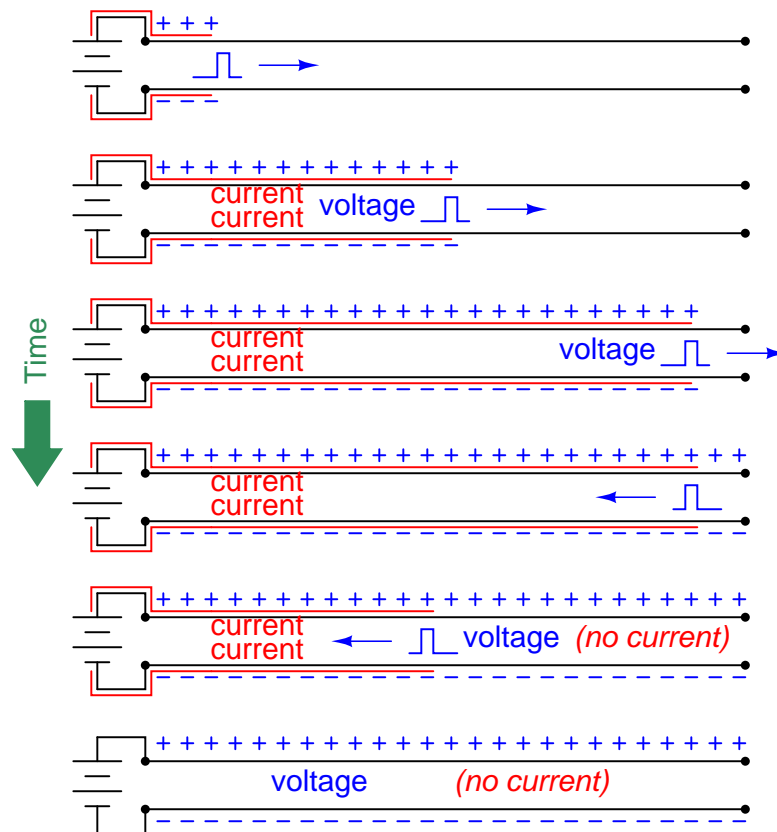
In honor of the cable’s capacity to behave as a temporary load to any impressed signal, we typically refer to it as something more than just a cable. From the perspective of an electrical pulse, measured on a time scale of nanoseconds, we refer to any cable as a *transmission line*. All electrical cables act as transmission lines, but the effects just mentioned are so brief in duration that we only notice them when the cable is exceptionally long and/or the pulse durations are exceptionally short (i.e. high-frequency signals).

During the time when a transmission line is absorbing energy from a power source – whether this is indefinitely for a transmission line of infinite length, or momentarily for a transmission line of finite length – the current it draws will be in direct proportion to the voltage applied by the source. In other words, a transmission line behaves like a resistor, at least for a moment. The amount of “resistance” presented by a transmission line is called its *characteristic impedance*, or *surge impedance*, symbolized in equations as Z_0 . Only after the pulse signal has had time to travel down the length of the transmission line and reflect back to the source does the cable stop acting as a load and begin acting as a plain pair of wires.

A transmission line’s characteristic impedance is a function of its conductor geometry (wire diameter, spacing) and the permittivity of the dielectric separating those conductors. If the line’s design is altered to increase its bulk capacitance and/or decrease its bulk inductance (e.g. decreasing the distance between conductors), the characteristic impedance will decrease. Conversely, if the transmission line is altered such that its bulk capacitance decreases and/or its bulk inductance increases, the characteristic impedance will increase. It should be noted that the length of the transmission line has absolutely no bearing on characteristic impedance. A 10-meter length of RG-58/U coaxial cable will have the exact same characteristic impedance as a 10000 kilometer length of RG-58/U coaxial cable (50 ohms, in both cases). The only difference is the *length of time* the cable will behave like a resistor to an applied voltage.


5.10.1 Open-ended transmission lines

The following sequence illustrates the propagation of a voltage pulse forward and back (reflected) on an open-ended transmission line beginning from the time the DC voltage source is first connected to the left-hand end:



The end result is a transmission line exhibiting the full source voltage, but no current. This is exactly what we would expect in an open circuit. However, during the time it took for the pulse to travel down the line's length and back, it drew current from the source equal to the source voltage divided by the cable's characteristic impedance ($I_{surge} = \frac{V_{source}}{Z_0}$). For a brief amount of time, the two-conductor transmission line acted as a *load* to the voltage source rather than an open circuit. Only after the pulse traveled down the full length of the line and back did the line finally act as a plain open circuit.

An experiment performed with a square-wave signal generator and oscilloscope⁵¹ connected to one end of a long wire pair cable (open on the far end) shows the effect of the reflected signal:



The waveform steps up for a brief time, then steps up further to full source voltage. The first step represents the voltage at the source during the time the pulse traveled along the cable's length, when the cable's characteristic impedance acted as a load to the signal generator (making its output voltage "sag" to a value less than its full potential). The next step represents the reflected pulse's return to the signal generator, when the cable's capacitance is fully charged and is no longer drawing current from the signal generator (making its output voltage "rise"). A two-step "fall" appears at the trailing edge of the pulse, when the signal generator reverses polarity and sends an opposing pulse down the cable.

The duration of the first and last "steps" on the waveform represents the time taken by the signal to propagate down the length of the cable *and* return to the source. This oscilloscope's timebase was set to 0.5 microseconds per division for this experiment, indicating a pulse round-trip travel time of approximately 0.2 microseconds. A cable of this type has a *velocity factor* of approximately 0.7, which means electrical impulses travel at only 70% the speed of light through the cable, and so the round-trip distance calculates to be approximately 42 meters which makes the cable 21 meters in length:


$$\text{Distance} = \text{Velocity} \times \text{Time} = (70\%)(3 \times 10^8 \text{ m/s})(0.2 \mu\text{s}) = 42 \text{ m}$$

$$\text{Cable length} = \frac{1}{2}(\text{Distance traveled by pulse}) = (0.5)(42 \text{ m}) = 21 \text{ m}$$

⁵¹The signal generator was set to a frequency of approximately 240 kHz with a Thévenin resistance of 118 ohms to closely match the cable's characteristic impedance of 120 ohms. The signal amplitude was just over 6 volts peak-to-peak.


5.10.2 Shorted transmission lines

The following sequence illustrates the propagation of a voltage pulse forward and back (reflected) on a shorted-end transmission line beginning from the time the DC voltage source is first connected to the left-hand end:



The end result is a transmission line exhibiting the full current of the source ($I_{max} = \frac{V_{source}}{R_{wire}}$), but no voltage. This is exactly what we would expect in a short circuit. However, during the time it took for the pulse to travel down the line's length and back, it drew current from the source equal to the source voltage divided by the cable's characteristic impedance ($I_{surge} = \frac{V_{source}}{Z_0}$). For a brief amount of time, the two-conductor transmission line acted as a moderate *load* to the voltage source rather than a direct short. Only after the pulse traveled down the full length of the line and back did the line finally act as a plain short-circuit.

An experiment performed with the same signal generator and oscilloscope connected to one end of the same long wire pair cable (shorted on the far end) shows the effect of the reflected signal:



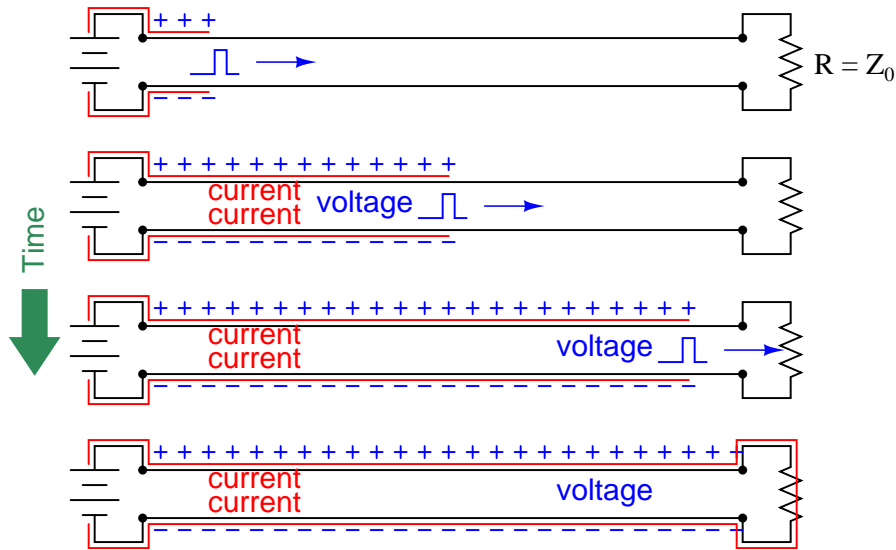
Here, the waveform steps up for a brief time, then steps down toward zero. As before, the first step represents the voltage at the source during the time the pulse traveled along the cable's length, when the cable's characteristic impedance acted as a load to the signal generator (making its output voltage "sag" to a value less than its full potential). The step down represents the (inverted) reflected pulse's return to the signal generator, nearly canceling the incident voltage and causing the signal to fall toward zero. A similar pattern appears at the trailing edge of the pulse, when the signal generator reverses polarity and sends an opposing pulse down the cable.

Note the duration of the pulse on this waveform, compared to the first and last "steps" on the open-circuited waveform. This pulse width represents the time taken by the signal to propagate down the length of the cable *and* return to the source. This oscilloscope's timebase remained at 0.5 microseconds per division for this experiment as well, indicating the same pulse round-trip travel time of approximately 0.2 microseconds. This stands to reason, as the cable length was not altered between tests; only the type of termination (short versus open).

5.10.3 Properly terminated transmission lines


Proper “termination” of a transmission line consists of connecting a resistance to the end(s) of the line so that the pulse “sees” the exact same amount of impedance at the end as it did while propagating along the line’s length. The purpose of the termination resistor is to completely dissipate the pulse’s energy in order that none of it will be reflected back to the source.

The following sequence illustrates the propagation of a voltage pulse on a transmission line with proper “termination” (i.e. a resistor matching the line’s surge impedance, connected to the far end) beginning from the time the DC voltage source is first connected to the left-hand end:



From the perspective of the pulse source, this properly terminated transmission line “looks” the same as an unterminated line of infinite length. There is no reflected pulse, and the DC voltage source “sees” an unchanging load resistance the entire time.

An experiment performed with a termination resistor in place shows the near-elimination of reflected pulses:



The pulse looks much more like the square wave it should be, now that the cable has been properly terminated⁵². With the termination resistor in place, a transmission line *always* presents the same impedance to the source, no matter what the signal level or the time of signal application. Another way to think of this is from the perspective of cable length. With the proper size of termination resistor in place, *the cable appears infinitely long* from the perspective of the power source because it never reflects any signals back to the source and it always consumes power from the source.

Data communication cables for digital instruments behave as transmission lines, and must be terminated at both ends to prevent signal reflections. Reflected signals (or “echoes”) may cause errors in received data in a communications network, which is why proper termination can be so important. For point-to-point networks (networks formed by exactly two electronic devices, one at either end of a single cable), the proper termination resistance is often designed into the transmission and receiving circuitry, and so no external resistors need be connected. For “multi-drop” networks where multiple electronic devices tap into the same electrical cable, excessive signal loading would occur if each and every device had its own built-in termination resistance, and so the devices are built with no internal termination, and the installer must place two termination resistors in the network (one at each far end of the cable).

⁵²The termination shown here is imperfect, as evidenced by the irregular amplitude of the square wave. The cable used for this experiment was a length of twin-lead speaker cable, with a characteristic impedance of approximately 120 ohms. I used a 120 ohm ($\pm 5\%$) resistor to terminate the cable, which apparently was not close enough to eliminate all reflections.

5.10.4 Discontinuities

A transmission line's characteristic impedance will be constant throughout its length so long as its conductor geometry and dielectric properties are consistent throughout its length. Abrupt changes in either of these parameters, however, will create a *discontinuity* in the cable capable of producing signal reflections. This is why transmission lines must never be sharply bent, crimped, pinched, twisted, or otherwise deformed.

The probe for a guided-wave radar (GWR) liquid level transmitter is another example of a transmission line, one where the vapor/liquid interface creates a discontinuity: there will be an abrupt change in characteristic impedance between the transmission line in vapor space versus the transmission line submerged in a liquid due to the differing dielectric permittivities of the two substances. This sudden change in characteristic impedance sends a reflected signal back to the transmitter. The time delay measured between the signal's transmission and the signal's reception by the transmitter represents the vapor space distance, or *ullage*.

For more detail on the theory and function of radar level measurement, see section [20.5.2](#) beginning on page [1473](#).

5.10.5 Velocity factor

The speed at which an electrical signal propagates down a transmission line is never as fast as the speed of light in a vacuum, owing to the permittivity of the line's electrical insulation being greater than that of a vacuum. A value called the *velocity factor* expresses the propagation velocity as a ratio to light, and its value is always less than one:

$$\text{Velocity factor} = \frac{v}{c}$$

Where,

v = Propagation velocity of signal traveling along the transmission line

c = Speed of light in a vacuum ($\approx 3.0 \times 10^8$ meters per second)

Velocity factor is a function of dielectric constant, but not conductor geometry. A greater permittivity value results in a slower velocity (lesser velocity factor).

5.10.6 Cable losses

Ideally, a transmission line is a perfectly loss-less conduit for electrical energy. That is, every watt of signal power entering the transmission line is available at the end where the load is connected. In reality, though, this is never the case. Conductor resistance, as well as losses within the dielectric (insulating) materials of the cable, rob the signal of energy.


For transmission lines, power loss is typically expressed in units of decibels per 100 feet or per 100 meters. A “decibel,” as you may recall, is ten times the logarithm of a power ratio:

$$\text{Gain or Loss in dB} = 10 \log \left(\frac{P}{P_{ref}} \right)$$

Thus, if a transmission line receives 25 milliwatts of signal power at one end, but only conveys 18 milliwatts to the far end, it has suffered a 1.427 dB loss ($10 \log \frac{0.018}{0.025} = -1.427$ dB) from end to end. Power loss in cables is strongly dependent on frequency: the greater the signal frequency, the more severe the power loss per unit cable length.


5.11 Antennas

Capacitors store energy in electric fields, proportional to the square of voltage. Inductors store energy in magnetic fields, proportional to the square of current. If capacitors and inductors are connected together, their complementary energy storage modes create a condition where electrical energy transfers back and forth between the capacitance and the inductance: voltage and current both oscillating sinusoidally. We refer to this cyclic exchange of energy as *resonance*. The simplest resonant circuit possible is the so-called *tank circuit*, comprised of a single inductor connected to a single capacitor:




The natural frequency at which a tank circuit oscillates is given by the formula $f_r = \frac{1}{2\pi\sqrt{LC}}$, where f_r is the resonant frequency in Hertz, C is the capacitance in Farads, and L is the inductance in Henrys.

A perfect tank circuit – devoid of electrical resistance and any other energy-dissipating characteristics (e.g. capacitor dielectric losses, inductor hysteresis losses) – would continue to oscillate forever once initially stimulated. That is, if initially “charged” with an impulse of DC voltage, the tank circuit would continue to produce AC voltage and current oscillations at its resonant frequency, at constant peak amplitude, forever:



Since real capacitors and inductors are not lossless, real tank circuits exhibit decaying-amplitude oscillations after initial “charging,” until no energy is stored in either the capacitor or the inductor:



Capacitive losses take the form of heat loss in the dielectric substance separating the capacitor plates. The electric field between the capacitor’s plates imparts forces on any polar⁵³ molecules within that substance, thereby doing “work” on those molecules by forcing them back and forth as the electric field alternates. Though these forces and motions are extremely small, they are nevertheless capable of draining considerable energy from the capacitor, dissipating it in the form of heat.

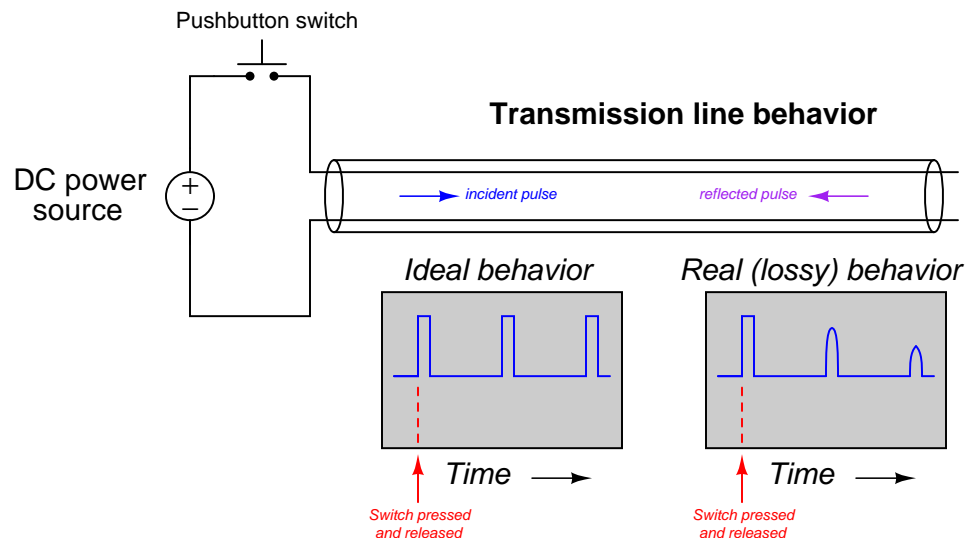
Inductive losses are similar, but in the form of work done on ferromagnetic molecules in the core material as the magnetic field alternates polarity. Like dielectric heating, magnetic losses also drain energy from the inductor, dissipating it in the form of heat.

Of course, both capacitors and inductors also contain ohmic resistance in the metals used to form the plates and wire coils, respectively. Resistance naturally dissipates energy in the form of heat, constituting another energy-loss mechanism for both capacitors and inductors (albeit much more significant in inductors than in capacitors!).

The combined effect of all these energy-loss mechanisms is that the oscillations of an unpowered tank circuit decay over time, until they cease completely. This is similar in principle to a pendulum gradually coming to a halt after being set in motion with a single push: if not for air resistance and other forms of friction, the pendulum should swing forever! With air friction and mechanical friction in the pendulum’s bearing, though, a pendulum’s oscillations gradually diminish in amplitude until all its energy has been lost – swing by swing – to heat.

⁵³A “polar” molecule is one where the constituent atoms are bound together in such a way that there is a definite electrical polarity from one end of the molecule to the other. Water (H₂O) is an example of a polar molecule: the positively charged hydrogen atoms are bound to the negatively charged oxygen atom in a “V” shape, so the molecule as a whole has a positive side and a negative side which allows the molecule to be influenced by external electric fields. Carbon dioxide (CO₂) is an example of a non-polar molecule whose constituent atoms lie in a straight line with no apparent electrical poles. Interestingly, microwave ovens exploit the fact of water molecules’ polarization by subjecting food containing water to a strong oscillating electric field (microwave energy in the gigahertz frequency range) which causes the water molecules to rotate as they continuously orient themselves to the changing field polarity. This oscillatory rotation manifests itself as heat within the food.


Capacitance and inductance, however, are not limited to capacitors and inductors: *any* pair of conductors separated by an insulating medium will exhibit capacitance, and *any* electrical conductor will exhibit inductance along its length. Even a two-conductor cable (a *transmission line*) has distributed capacitance and inductance capable of storing energy. If a long, unterminated (no load connected to the far end) cable is “charged” by a pulse of applied DC voltage, it will sustain a series of echoing pulses at a period dependent on the cable’s length and velocity factor:



The ability for a transmission line to support periodic signal “echoes” means it may also resonate when energized by an AC power source, just like a tank circuit. Unlike a tank circuit, however, a transmission line is able to resonate at more than one frequency: a “fundamental” frequency, or any whole-number multiple of that fundamental frequency called a *harmonic frequency*. For example, an unterminated transmission line with a length of 1 kilometer and a velocity factor of 0.7 has a round-trip echo time (period) of 9.53 microseconds, equivalent to a resonant frequency of 104.9 kHz. However, it will resonate equally well at exactly twice this fundamental frequency (209.8 kHz – the *second harmonic* of 104.9 kHz) as well as three times the fundamental frequency (314.8 kHz – the *third harmonic* of 104.9 kHz), etc. A simple LC tank circuit, by contrast, will only resonate at a single frequency.

This “poly-resonant” behavior of transmission lines has an analogue in the world of music. “Wind” instruments such as trombones, flutes, trumpets, saxophones, clarinets, recorders, etc., are really nothing more than tubes with at least one open end. These tubes will acoustically resonate at specific frequencies when “excited” by turbulent air flow at one end. The lowest frequency such a tube will resonate at is its “fundamental” frequency, but increasing the turbulence of the incoming air flow will cause the tone to “jump” to some harmonic of that fundamental frequency. The fundamental frequency of the tube may be altered by changing the length of the tube (e.g. as in a trombone or trumpet) or by opening ports along the tube’s length to alter its effective length (flute, saxophone, clarinet, recorder, etc.).

If we were to alter our transmission line test circuit, splaying the two conductors apart from each other rather than running them alongside one another, it would also form another type of resonant circuit, owing to the distributed capacitance and inductance along the wires' lengths:



This special case of resonant circuit has some very interesting properties. First, its resonant frequency is quite high, because the distributed inductance and capacitance values are extremely small compared to the values of discrete components such as inductors and capacitors. Second, it is a very “lossy” circuit despite having no appreciable electrical resistance to dissipate energy, no solid insulating medium to incur dielectric losses, and no ferromagnetic core to exhibit hysteresis losses. This special form of resonant circuit loses energy not to heat, but rather to *electromagnetic radiation*. In other words, the energy in the electric and magnetic fields leave the circuit and propagate through space in the form of *electro-magnetic waves*, or what we more commonly refer to now as *radio waves*: a series of electric and magnetic fields oscillating as they radiate away from the wires at the speed of light.

Tank circuits and transmission lines do not radiate energy because they are intentionally designed to *contain* their fields: capacitors are designed to fully contain their electric fields, inductors are designed to fully contain their magnetic fields, and the fields within transmission lines are likewise constrained. Our two-wire resonant circuit, by contrast, does just the opposite: its electric and magnetic fields are exposed to open space with no containment whatsoever. What we have built here is not a tank circuit nor a transmission line, but rather an *antenna*⁵⁴: a structure designed to radiate electric and magnetic fields as waves into surrounding space.

⁵⁴An older term used by radio pioneers to describe antennas is *radiator*, which I personally find very descriptive. The word “antenna” does an admirable job describing the physical appearance of the structure – like antennas on an insect – but the word “radiator” actually describes its *function*, which is a far more useful principle for our purposes.

5.11.1 Maxwell and Hertz

An interesting historical footnote is that this phenomenon of electromagnetic waves propagating through space was predicted theoretically before it was demonstrated experimentally. A Scottish physicist named James Clerk Maxwell made an astonishing theoretical prediction which he published in 1873, expressed in these four equations:

$$\oint \mathbf{E} \cdot d\mathbf{A} = \frac{Q}{\epsilon_0}$$

$$\oint \mathbf{B} \cdot d\mathbf{A} = 0$$

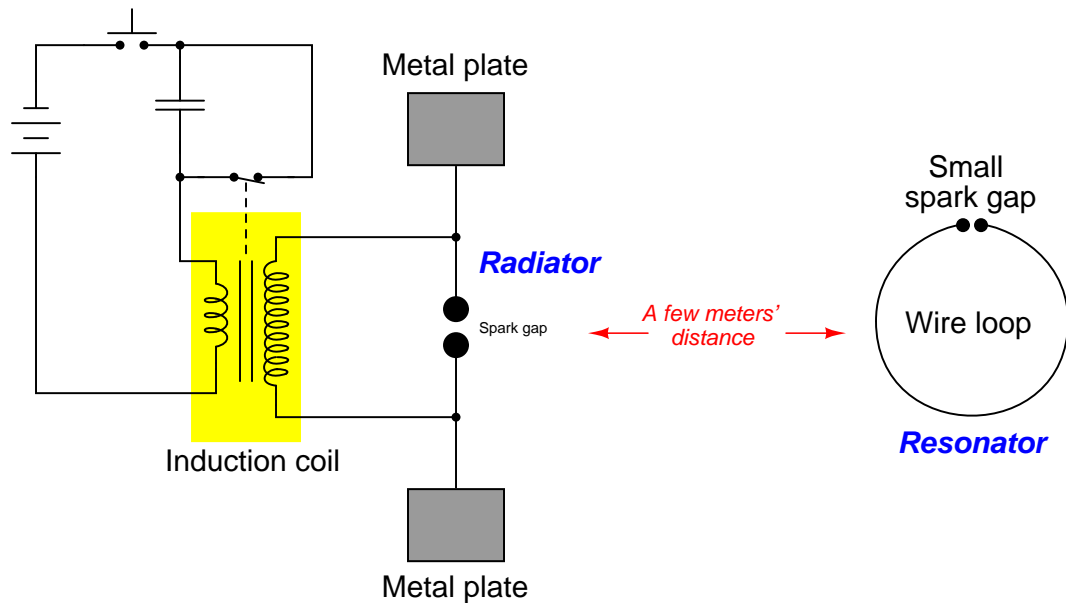
$$\oint \mathbf{E} \cdot d\mathbf{s} = -\frac{d\Phi_B}{dt}$$

$$\oint \mathbf{B} \cdot d\mathbf{s} = \mu_0 I + \mu_0 \epsilon_0 \frac{d\Phi_E}{dt}$$

The last two equations hold the most interest to us with respect to electromagnetic waves. The third equation states that an electric field (\mathbf{E}) will be produced in open space by a changing magnetic flux ($\frac{d\Phi_B}{dt}$). The fourth equation states that a magnetic field (\mathbf{B}) will be produced in open space either by an electric current (I) or by a changing electric flux ($\frac{d\Phi_E}{dt}$). Given this complementary relationship, Maxwell reasoned, it was possible for a changing electric field to create a changing magnetic field which would then create another changing electric field, and so on. This cause-and-effect cycle could continue, ad infinitum, with fast-changing electric and magnetic fields radiating off into open space without needing wires to carry or guide them. In other words, the complementary fields would be self-sustaining as they traveled.

The Prussian Academy of Science offered a reward to anyone who could experimentally validate Maxwell's prediction, and this challenge was met by Professor Heinrich Hertz at the Engineering College in Karlsruhe, Germany in 1887, eight years after Maxwell's death. Hertz constructed and tested a pair of devices: a "radiator" to produce the electromagnetic waves, and a "resonator" to receive them.

A simplified diagram showing Hertz's experimental device is shown here:



An “induction coil” (a buzzing device constructed of a self-interrupting relay and step-up transformer winding to generate a continuous pulsing waveform of high voltage) provided an extremely noisy (i.e. frequency-rich) AC signal to the radiator, while a spark gap at the resonator provided indication that the electromagnetic waves were captured and converted into voltage by the resonator wire.


Both the radiator and the resonator are what we would now call *antennas*. The purpose of the transmitting antenna (radiator) is to take high-frequency AC power and radiate that power in the form of electromagnetic waves: self-sustaining electric and magnetic fields propagating out into open space. The purpose of the receiving antenna is to capture those electromagnetic waves and convert them into an AC signal. All antennas – from historical to the most modern – behave in fundamentally the same way: energize them with high-frequency AC power, and they will radiate electromagnetic waves at that frequency; expose them to electromagnetic waves, and they will produce a very small AC signal at the same frequency as the incident radiation.

5.11.2 Antenna size

Earlier it was mentioned that antennas are fundamentally *resonant* elements: they “prefer” to electrically oscillate at a fundamental frequency (and at whole-number multiples of that fundamental frequency). An antenna will behave most efficiently – whether transmitting or receiving – when operated in a condition of resonance. The relationship between ideal frequency and antenna length is inverse: the longer the antenna, the lower its fundamental frequency, and vice-versa. This is the same mathematical relationship we see between frequency and wavelength (λ) for any moving wave:

$$v = \lambda f$$


The prototypical antenna shown earlier – with two wires oriented 180° from each other – operates best at a wavelength twice as long as the total length of wire.



A “half-wave” dipole antenna with a total length of 5 meters⁵⁵ will radiate optimally at a frequency of 30 MHz, given that the velocity of electromagnetic waves in open space is approximately 3×10^8 meters per second. This same antenna will also effectively resonate at any *harmonic* (integer multiple) of 30 MHz (e.g. 60 MHz, 90 MHz, 120 MHz, etc.) just as a transmission line is able to resonate at a fundamental frequency as well as any harmonic thereof.

⁵⁵In practice, the ideal length of a dipole antenna turns out to be just a bit shorter than theoretical, due to lumped-capacitive effects at the wire ends. Thus, a resonant 30 MHz half-wave dipole antenna should actually be about 4.75 meters in length rather than exactly 5 meters in length.

A popular variation on the theme of the half-wave dipole antenna is the so-called *quarter-wave* or “whip” antenna, which is exactly what you might think it is: one-half of a half-wave antenna. Instead of two wires pointed away from each other, we substitute an earth-ground connection for one of the wires:



Quarter-wave antennas tend to be less effective than half-wave antennas, but are usually much more convenient to construct for real applications.

5.11.3 Antenna orientation and directionality


Another factor affecting antenna performance is the orientation of the receiving antenna with respect to the transmitting antenna. Antenna conductors should be *parallel* to each other in order to maximize reception, in order that the electric and magnetic fields emanating from the wires of the transmitting antenna will “link” properly with the wires of the receiving antenna(s). If the goal is optimum communication in any direction (omnidirectionality), dipole and whip antennas should be arranged *vertically* so that all antenna conductors will be parallel to each other regardless of their geographic location.

Omnidirectionality may seem like a universally good trait for any antenna: to be able to transmit and receive electromagnetic waves equally well in any direction. However, there are good reasons for wanting directionality in an antenna. One reason is for greater security: perhaps you have an application where you do *not* wish to broadcast information in all directions, where anyone at all could receive that information. In that case, the best antennas to use would be those that work best in one direction and one direction only, with transmitting and receiving antenna pairs pointed directly at each other.


Another reason for antenna directionality is improved reception. As noted before, the AC signal received at an antenna is very small, typically on the order of *microvolts*. Since electromagnetic radiation tends to “spread” as it travels, becoming weaker with distance from the transmitting antenna, long-range radio communication benefits from increased sensitivity. A transmitting antenna that is directional will focus more of its power in one direction than in others, resulting in less “wasted” power radiated in unwanted directions. Likewise, a receiving antenna that is directional does a better job of “collecting” the electromagnetic energy it receives from that direction (as well as ignoring electromagnetic waves coming from other directions).

A simple yet highly effective antenna design for directional transmission and reception is the *Yagi*, named after its inventor. A Yagi is based on a half-wave dipole element surrounded by one or more wires longer than $\lambda/2$ to the rear (called “reflectors”) and one or more wires shorter than $\lambda/2$ to the front (called “directors”). The terms “reflector” and “director” are quite apt, describing their interaction with electromagnetic waves from the perspective of the dipole: reflectors *reflect* the waves, while directors *direct* the waves. The result is an antenna array that is much more directional than a simple dipole:

"Yagi" antenna operating in transmission mode




"Yagi" antenna operating in reception mode



An example of a Yagi antenna used as part of a SCADA system is shown in this photograph, the antenna of course being the multi-element array in the upper-right corner:



Another example of a highly directional antenna design is the *parabolic dish*, often used in microwave and satellite communications. This photograph shows two “dish” antennas, one open to the weather (right) and the other covered with a weather-proof diaphragm to keep the antenna elements protected (left):



Some dish antennas use mesh or metal-tube reflectors rather than a solid parabolic reflector, as is the case in this next photograph:



References

- Blackburn, J. Lewis and Domin, Thomas J., *Protective Relaying Principles and Applications*, Third Edition, CRC Press, Taylor & Francis Group, Boca Raton, FL, 2007.
- Boylestad, Robert L., *Introductory Circuit Analysis*, 9th Edition, Prentice Hall, Upper Saddle River, NJ, 2000.
- Dorf, Richard C., *Modern Control Systems*, Fifth Edition, Addison-Wesley Publishing Company, Reading, MA, 1989.
- Eckman, Donald P., *Automatic Process Control*, John Wiley & Sons, Inc., New York, NY, 1958.
- Field Antenna Handbook*, U.S. Marine Corps document MCRP 6-22D, 1999.
- Giancoli, Douglas C., *Physics for Scientists & Engineers*, Third Edition, Prentice Hall, Upper Saddle River, NJ, 2000.
- Harrison, Cecil A., *Transform Methods in Circuit Analysis*, Saunders College Publishing, Philadelphia, PA, 1990.
- Jenkins, John D., *Where Discovery Sparks Imagination*, American Museum of Radio and Electricity, Bellingham, WA, 2009.
- Kaplan, Wilfred, *Advanced Mathematics for Engineers*, Addison-Wesley Publishing Company, Reading, MA, 1981.
- Mileaf, Harry, *Electronics One-Seven*, Hayden Book Company, 1976.
- Nilsson, James W., *Electric Circuits*, Addison-Wesley Publishing Company, Reading, MA, 1983.
- Palm, William J., *Control Systems Engineering*, John Wiley & Sons, Inc., New York, NY, 1986.
- Smith, Steven W., *The Scientist and Engineer's Guide to Digital Signal Processing*, California Technical Publishing, San Diego, CA, 1997.
- Steinmetz, Charles P., *Theoretical Elements of Electrical Engineering*, Third Edition, McGraw-Hill Book Company, New York, NY, 1909.
- Steinmetz, Charles P., *Theory and Calculation of Alternating Current Phenomena*, Third Edition, McGraw Publishing Company, New York, NY, 1900.
- The ARRL Antenna Book*, Eleventh Edition, The American Radio Relay League, Inc., Newington, CT, 1968.
- The ARRL Handbook For Radio Amateurs*, 2001 Edition, ARRL – the national association for Amateur Radio, Newington, CT, 2001.

Chapter 6

Introduction to industrial instrumentation

Instrumentation is the science of automated measurement and control. Applications of this science abound in modern research, industry, and everyday living. From automobile engine control systems to home thermostats to aircraft autopilots to the manufacture of pharmaceutical drugs, automation surrounds us. This chapter explains some of the fundamental principles of industrial instrumentation.

The first step, naturally, is measurement. If we can't measure something, it is really pointless to try to control it. This "something" usually takes one of the following forms in industry:


- Fluid pressure
- Fluid flow rate
- The temperature of an object
- Fluid volume stored in a vessel
- Chemical concentration
- Machine position, motion, or acceleration
- Physical dimension(s) of an object
- Count (inventory) of objects
- Electrical voltage, current, or resistance

Once we measure the quantity we are interested in, we usually transmit a signal representing this quantity to an indicating or computing device where either human or automated action then takes place. If the controlling action is automated, the computer sends a signal to a final controlling device which then influences the quantity being measured.

This final control device usually takes one of the following forms:

- Control valve (for throttling the flow rate of a fluid)
- Electric motor
- Electric heater

Both the measurement device and the final control device connect to some physical system which we call the *process*. To show this as a general block diagram:



The common home thermostat is an example of a measurement and control system, with the home's internal air temperature being the "process" under control. In this example, the thermostat usually serves two functions: sensing and control, while the home's heater adds heat to the home to increase temperature, and/or the home's air conditioner extracts heat from the home to decrease temperature. The job of this control system is to maintain air temperature at some comfortable level, with the heater or air conditioner taking action to correct temperature if it strays too far from the desired value (called the *setpoint*).

Industrial measurement and control systems have their own unique terms and standards, which is the primary focus of this lesson. Here are some common instrumentation terms and their definitions:

Process: The physical system we are attempting to control or measure. *Examples: water filtration system, molten metal casting system, steam boiler, oil refinery unit, power generation unit.*

Process Variable, or PV: The specific quantity we are measuring in a process. *Examples: pressure, level, temperature, flow, electrical conductivity, pH, position, speed, vibration.*

Setpoint, or **SP**: The value at which we desire the process variable to be maintained at. In other words, the “target” value for the process variable.

Primary Sensing Element, or **PSE**: A device directly sensing the process variable and translating that sensed quantity into an analog representation (electrical voltage, current, resistance; mechanical force, motion, etc.). *Examples: thermocouple, thermistor, bourdon tube, microphone, potentiometer, electrochemical cell, accelerometer.*

Transducer: A device converting one standardized instrumentation signal into another standardized instrumentation signal, and/or performing some sort of processing on that signal. Often referred to as a *converter* and sometimes as a “relay.” *Examples: I/P converter (converts 4-20 mA electric signal into 3-15 PSI pneumatic signal), P/I converter (converts 3-15 PSI pneumatic signal into 4-20 mA electric signal), square-root extractor (calculates the square root of the input signal).*

Note: in general science parlance, a “transducer” is any device converting one form of energy into another, such as a microphone or a thermocouple. In industrial instrumentation, however, we generally use “primary sensing element” to describe this concept and reserve the word “transducer” to specifically refer to a conversion device for standardized instrumentation signals.

Transmitter: A device translating the signal produced by a primary sensing element (PSE) into a *standardized* instrumentation signal such as 3-15 PSI air pressure, 4-20 mA DC electric current, Fieldbus digital signal packet, etc., which may then be conveyed to an indicating device, a controlling device, or both.

Lower- and Upper-range values, abbreviated **LRV** and **URV**, respectively: the values of process measurement deemed to be 0% and 100% of a transmitter’s calibrated range. For example, if a temperature transmitter is calibrated to measure a range of temperature starting at 300 degrees Celsius and ending at 500 degrees Celsius, its LRV would be 300 °C and its URV would be 500 °C.

Zero and **Span**: alternative descriptions to LRV and URV for the 0% and 100% points of an instrument’s calibrated range. “Zero” refers to the beginning-point of an instrument’s range (equivalent to LRV), while “span” refers to the width of its range (URV – LRV). For example, if a temperature transmitter is calibrated to measure a range of temperature starting at 300 degrees Celsius and ending at 500 degrees Celsius, its zero would be 300 °C and its span would be 200 °C.

Controller: A device receiving a process variable (PV) signal from a primary sensing element (PSE) or transmitter, comparing that signal to the desired value (called the setpoint) for that process variable, and calculating an appropriate output signal value to be sent to a final control element (FCE) such as an electric motor or control valve.

Final Control Element, or **FCE**: A device receiving the signal output by a controller to directly influence the process. *Examples: variable-speed electric motor, control valve, electric heater.*

Manipulated Variable, or **MV**: The quantity in a process we adjust or otherwise manipulate in order to influence the process variable (PV). Also used to describe the output signal generated by a controller; i.e. the signal commanding (“manipulating”) the final control element to influence the process.

Automatic mode: When the controller generates an output signal based on the relationship of process variable (PV) to the setpoint (SP).

Manual mode: When the controller's decision-making ability is bypassed to let a human operator directly determine the output signal sent to the final control element.


Now we will explore some practical examples of measurement and control systems so you can get a better idea of these fundamental concepts.

6.1 Example: boiler water level control system

Steam boilers are very common in industry, principally because steam power is so useful. Common uses for steam in industry include doing mechanical work (e.g. a steam engine moving some sort of machine), heating, producing vacuums (through the use of "steam ejectors"), and augmenting chemical processes (e.g. reforming of natural gas into hydrogen and carbon dioxide).

The process of converting water into steam is quite simple: heat up the water until it boils. Anyone who has ever boiled a pot of water for cooking knows how this process works. Making steam continuously, however, is a little more complicated. An important variable to measure and control in a continuous boiler is the level of water in the "steam drum" (the upper vessel in a water-tube boiler). In order to safely and efficiently produce a continuous flow of steam, we must ensure the steam drum never runs too low on water, or too high. If there is not enough water in the drum, the water tubes may run dry and burn through from the heat of the fire. If there is too much water in the drum, liquid water may be carried along with the flow of steam, causing problems downstream.

In this next illustration, you can see the essential elements of a water level control system, showing transmitter, controller, and control valve:



The first instrument in this control system is the *level transmitter*, or “LT”. The purpose of this device is to sense the water level in the steam drum and report (“transmit”) that measurement to the controller in the form of a signal. In this case, the type of signal is *pneumatic*: a variable air pressure sent through metal or plastic tubes. The greater the water level in the drum, the more air pressure output by the level transmitter. Since the transmitter is pneumatic, it must be supplied with a source of clean, compressed air on which to operate. This is the meaning of the “A.S.” tube (Air Supply) entering the top of the transmitter.

This pneumatic signal is sent to the next instrument in the control system, the *level indicating controller*, or “LIC”. The purpose of this instrument is to compare the level transmitter’s signal against a *setpoint* value entered by a human operator representing the desired water level in the steam drum. The controller then generates an *output* signal telling the control valve to either introduce more or less water into the boiler to maintain the steam drum water level at setpoint. As with the transmitter, the controller in this system is pneumatic, operating entirely on compressed

air. This means the output of the controller is also a variable air pressure signal, just like the signal output by the level transmitter. Naturally, the controller requires a constant supply of clean, compressed air on which to run, which explains the “A.S.” (Air Supply) tube connecting to it.

The last instrument in this control system is the control valve, operated directly by the air pressure signal output by the controller. Its purpose is to influence the flow rate of water into the boiler, “throttling” the water flow more or less as determined by controller. This particular type of control valve uses a large diaphragm and a large spring to move the valve further open with more signal pressure and further closed with less signal pressure.

When the controller is placed in the “automatic” mode, it will move the control valve to whatever position necessary to maintain a constant steam drum water level. The phrase “whatever position necessary” suggests the relationship between the controller output signal, the process variable signal (PV), and the setpoint (SP) is complex. If the controller senses a water level above setpoint, it will close off the valve as far as necessary to decrease the water level down to setpoint. Conversely, if the controller senses a water level below setpoint, it will open up the valve as far as necessary to raise the water level up to setpoint.

What this means in a practical sense is that the controller’s output signal (equating to valve position) in automatic mode is just as much a function of process load (i.e. how much steam is being used from the boiler) as it is a function of setpoint (i.e. where we wish the water level to be). Consider a situation where the steam demand from the boiler is very low. If there isn’t much steam being drawn off the boiler, this means there will be little water boiled into steam and therefore little need for additional feedwater to be pumped into the boiler. Therefore, in this situation, one would expect the control valve to hover near the fully-closed position, allowing just enough water into the boiler to keep the steam drum water level at setpoint. If, however, there is a high demand for steam from this boiler, the rate of evaporation will be much greater. This means the control system must add feedwater to the boiler at a much greater flow rate in order to maintain the steam drum water level at setpoint. In this situation we would expect to see the control valve much closer to being fully-open as the control system “works harder” to maintain a constant water level in the steam drum. Thus, we see how the controller automatically positions the control valve to react to different boiler operating conditions even when the setpoint is fixed.

A human operator supervising this boiler has the option of placing the controller into “manual” mode. In this mode the control valve position is under direct control of the human operator, with the controller essentially ignoring the signal sent from the water level transmitter. Being an indicating controller, the controller faceplate will still show how much water is in the steam drum, but it is now the human operator’s sole responsibility to move the control valve to the appropriate position to hold water level at setpoint – *in manual mode the controller takes no corrective action of its own*. Manual mode is useful to human operators during start-up and shut-down conditions. It is also useful to instrument technicians for troubleshooting misbehaving control systems. Placing a controller into manual mode is akin to disengaging the cruise control in an automobile, transferring control of engine power from the car’s computer back to the human driver. One can easily imagine an automobile mechanic needing to throttle a car’s engine “manually” (i.e. with the cruise control turned off) in order to properly diagnose an engine or drivetrain problem. This is true for industrial processes as well, where instrument technicians may need to place a controller into manual mode in order to properly diagnose transmitter or control valve problems.

As was mentioned before, this is an example of a *pneumatic* (compressed air) control system, where all the instruments operate on compressed air, and use compressed air as the signaling medium. Pneumatic instrumentation is an old technology, dating back to the early twentieth century. While most modern instruments are electronic in nature, pneumatic instruments still find application within industry. The most common industry standard for pneumatic pressure signals is 3 to 15 PSI, with 3 PSI representing low end-of-scale and 15 PSI representing high end-of-scale. Alternative pressure ranges for pneumatic signals sometimes encountered in industry include 3 to 27 PSI, and 6 to 30 PSI. The following table shows the relationship between air signal pressure and steam drum level for this boiler's 3-15 PSI level transmitter:

Transmitter air signal pressure	Steam drum water level
3 PSI	0% (Empty)
6 PSI	25%
9 PSI	50%
12 PSI	75%
15 PSI	100% (Full)

It should be noted this table assumes the transmitter measures the *full range* of water level possible in the drum. Usually, this is not the case. Instead, the transmitter will be calibrated so it only senses a narrow range of water level near the middle of the drum. Thus, 3 PSI (0%) will not represent an empty drum, and neither will 15 PSI (100%) represent a completely full drum. Calibrating the transmitter like this helps avoid the possibility of actually running the drum completely empty or completely full in the case of an operator incorrectly setting the setpoint value near either extreme end of the measurement scale.

An example table showing this kind of realistic transmitter calibration appears here:

Transmitter air signal pressure	Actual steam drum water level
3 PSI	40%
6 PSI	45%
9 PSI	50%
12 PSI	55%
15 PSI	60%

The boiler's steam drum level controller outputs a pneumatic output signal to the control valve, using the same 3 to 15 PSI standard to command different valve positions:


Controller output signal pressure	Control valve position
3 PSI	0% open (Fully shut)
6 PSI	25% open
9 PSI	50% open
12 PSI	75% open
15 PSI	100% (Fully open)

Even though the same range of air pressure (3 to 15 PSI) is used to represent water level in the steam drum *and* the position of the control valve, there is no simple correspondence between the two signals. A common misconception for students new to this topic is to assume the transmitter signal (PV) and controller output signal must be identical. *This is not true.* Typically the 3-15 PSI signal representing level will be at some value different from the 3-15 PSI signal driving the valve, because those two signals represent two entirely different variables in the boiler system. As we have seen previously, the output signal from a controller in automatic mode is just as much a function of process conditions as it is a function of the measured variable. This error is akin to thinking the road speed signal in an automobile cruise control system (the “process variable” or PV) must be the same value as the signal sent by the cruise control computer to the engine’s accelerator control (the controller’s “output” signal). Granted, these two signals are *related* to one another, but since they represent two different parameters in the controlled system we have no reason to expect their values will be *equal* except by chance.

6.2 Example: wastewater disinfection

The final step in treating wastewater before releasing it into the natural environment is to kill any harmful microorganisms (e.g. viruses, bacteria) in it. This is called *disinfection*, and chlorine gas is a very effective disinfecting agent. However, just as it is not good to mix too little chlorine in the outgoing water (effluent) because we might not disinfect the water thoroughly enough, there is also danger of injecting too much chlorine in the effluent because then we might begin poisoning animals and beneficial microorganisms in the natural environment.

To ensure the right amount of chlorine injection, we must use a dissolved chlorine analyzer to measure the chlorine concentration in the effluent, and use a controller to automatically adjust the chlorine control valve to inject the right amount of chlorine at all times. The following P&ID (Process and Instrument Diagram) shows how such a control system might look:



Chlorine gas coming through the control valve mixes with the incoming water (influent), then has time to disinfect in the contact chamber before exiting out to the environment.

The transmitter is labeled “AT” (Analytical Transmitter) because its function is to *analyze* the concentration of chlorine dissolved in the water and *transmit* this information to the control system. The “Cl₂” (chemical notation for a chlorine molecule) written near the transmitter bubble declares this to be a chlorine analyzer. The dashed line coming out of the transmitter tells us the signal is electric in nature, not pneumatic as was the case in the previous (boiler control system) example. The most common and likely standard for electronic signaling in industry is 4 to 20 milliamps DC, which represents chlorine concentration in much the same way as the 3 to 15 PSI pneumatic signal standard represented steam drum water level in the boiler:

Transmitter signal current	Chlorine concentration
4 mA	0% (no chlorine)
8 mA	25%
12 mA	50%
16 mA	75%
20 mA	100% (Full concentration)

The controller is labeled “AIC” because it is an Analytical Indicating Controller. Controllers are always designated by the process variable they are charged with controlling, in this case the chlorine analysis of the effluent. “Indicating” means there is some form of display that a human operator or technician can read showing the chlorine concentration. “SP” refers to the setpoint value entered by the operator, which the controller tries to maintain by adjusting the position of the chlorine injection valve.

A dashed line going from the controller to the valve indicates another electronic signal: a 4 to 20 mA direct current signal again. Just as with the 3 to 15 PSI pneumatic signal standard in the pneumatic boiler control system, the amount of electric current in this signal path directly relates to a certain valve position:

Controller output signal current	Control valve position
4 mA	0% open (Fully shut)
8 mA	25% open
12 mA	50% open
16 mA	75% open
20 mA	100% (Fully open)


Note: it is possible, and in some cases even preferable, to have either a transmitter or a control valve that responds in reverse fashion to an instrument signal such as 3 to 15 PSI or 4 to 20 milliamps. For example, this valve could have been set up to be wide open at 4 mA and fully shut at 20 mA. The main point to recognize here is that both the process variable sensed by the transmitter and the position of the control valve are proportionately represented by analog signals.

Just as with the 3 to 15 PSI pneumatic signals used to represent water level and control valve position in the boiler seen previously, the two 4 to 20 milliamp current signals in this system represent two different variables in the system and therefore will not be equal to each other except by coincidence. A common misconception for people first learning about analog instrumentation signals is to assume the transmitter’s signal (“Process Variable”) must be identical in value to the control valve’s signal (“Manipulated Variable” or “Output”), but this is not true.

The letter “M” inside the control valve bubble tells us this is a motor-actuated valve. Instead of using compressed air pushing against a spring-loaded diaphragm as was the case in the boiler control system, this valve is actuated by an electric motor turning a gear-reduction mechanism. The gear reduction mechanism allows slow motion of the control valve stem even though the motor spins at a fast rate. A special electronic control circuit inside the valve actuator modulates electric power to the electric motor in order to ensure the valve position accurately matches the signal sent by the controller. In effect, this is another control system in itself, controlling valve position according to a “setpoint” signal sent by another device (in this case, the AIT controller which is telling the valve what position to go to).

6.3 Example: chemical reactor temperature control

Sometimes we encounter a diversity of instrument signal standards in one control system. Such is the case with the following chemical reactor temperature control system, where three different signal standards convey information between the instruments. A P&ID (Process and Instrument Diagram) shows the inter-relationships of the process piping, vessels, and instruments:



The purpose of this control system is to ensure the chemical solution inside the reactor vessel is maintained at a constant temperature. A steam-heated “jacket” envelops the reactor vessel, transferring heat from the steam into the chemical solution inside. The control system maintains a constant temperature by measuring the temperature of the reactor vessel, and throttling steam from a boiler to the steam jacket to add more or less heat as needed.

We begin as usual with the temperature transmitter, located near the bottom of the vessel. Note the different line type used to connect the temperature transmitter (TT) with the temperature-indicating controller (TIC): hollow diamonds with lines in between. This signifies a *digital electronic instrument signal* – sometimes referred to as a *fieldbus* – rather than an analog type (such as 4 to 20 mA or 3 to 15 PSI). The transmitter in this system is actually a digital computer, and so is the controller. The transmitter reports the process variable (reactor temperature) to the controller using digital bits of information. Here there is no analog scale of 4 to 20 milliamps, but rather electric voltage/current pulses representing the 0 and 1 states of binary data.

Digital instrument signals are capable of transferring multiple data points rather than single data points as is the case with analog instrument signals. This means digital instrument signals may convey device status information (such as self-diagnostic test results) as well as the basic measurement value. In other words, the digital signal coming from this transmitter not only tells the controller how hot the reactor is, but it may also communicate to the controller how well the transmitter is functioning.

The dashed line exiting the controller shows it to be analog electronic: most likely 4 to 20 milliamps DC. This electronic signal does not go directly to the control valve, however. It passes through a device labeled “TY”, which is a *transducer* to convert the 4 to 20 mA electronic signal into a 3 to 15 PSI pneumatic signal which then actuates the valve. In essence, this signal transducer acts as an electrically-controlled air pressure regulator, taking the supply air pressure (usually 20 to 25 PSI) and regulating it down to a level commanded by the controller’s electronic output signal.

At the temperature control valve (TV) the 3 to 15 PSI pneumatic pressure signal applies a force on a diaphragm to move the valve mechanism against the restraining force of a large spring. The construction and operation of this valve is the same as for the feedwater valve in the pneumatic boiler water control system. The letters “ATO” immediately below the valve symbol mean “Air-To-Open,” referring to the direction this valve mechanism will move (wider open) as more air signal pressure is applied to its actuator.

A detail not shown on this diagram, yet critically important to the operation of the temperature control system, is the *direction of action* for the controller while in automatic mode. It is possible to configure general-purpose controllers to act either in a *direct* fashion where an increasing process variable signal automatically results in an increasing output signal, or in a *reverse* fashion where an increasing process variable signal automatically results in a decreasing output signal. An effective way to identify the proper direction of action for any process controller is to perform a “thought experiment¹” whereby we imagine the process variable increasing over time, and then determine which way the controller’s output needs to change in order to bring the process variable value back to setpoint based on the final control element’s influence within the process.

In this process, let us imagine the reactor temperature increasing for some reason, perhaps an increase in the temperature of the feed entering the reactor. With an increasing temperature, the controller must *reduce* the amount of steam applied to the heating jacket surrounding the reactor in order to correct for this temperature change. With an air-to-open (ATO) steam valve, this requires a *decreased* air pressure signal to the valve in order to close it further and reduce heat input to the reactor. Thus, if an increasing process variable signal requires a decreasing controller output signal, the controller in this case needs to be configured for *reverse* action.

We could easily imagine reasons why the temperature controller in this process might have to be configured for direct action instead of reverse action. If the piping were altered such that the control valve throttled the flow of *coolant* to the reactor rather than steam, an increasing temperature would require a further-open valve, which would only happen if the controller were configured for direct action. Alternatively, if the steam valve were air-to-close (ATC) rather than air-to-open (ATO), an increasing reactor temperature (requiring less steam be sent to the reactor) would necessitate the controller outputting an increased signal to the valve, so that more air signal pressure pushed the valve further closed.

¹For more information on conducting “thought experiments,” refer to the subsection of this book titled “Using Thought Experiments” (34.3.4) beginning on page 2755.

An example of a chemical reaction temperature control system requiring *direct* controller action is shown in the following photograph. Here, we see a jacketed stainless-steel vessel used to ferment beer at cold temperatures. The jacket surrounding this vessel is pumped full of chilled glycol solution (similar to automotive antifreeze), to draw heat away from the fermenting beer and maintain its temperature well below ambient:



If the beer becomes too warm, the controller sends an increased signal to the glycol valve sending more chilled glycol through the vessel's jacket to remove heat from the beer. Since the relationship between the controller's process variable and its output is direct (i.e. rising PV results in rising Output), the controller needs to be configured for direct action.

This is why general-purpose process controllers always provide a user-selectable option for either direct or reverse action: it makes them adaptable to the needs of *any* process, no matter the physics of the process or the behavior of the other loop instruments (e.g. transmitter and final control element).

An additional instrument connected to our hypothetical chemical reactor is a pressure transmitter (PT) on the feed line. While not a part of the temperature control loop, it is shown here to illustrate yet another type of instrumentation signaling: *digital wireless*. Here, the transmitter reports its measurement data to an indicator at the control room via radio signals, using digital codes much like fieldbus to communicate not only the basic process data but also transmitter diagnostic and radio network management data.


At the time of this writing (2011), wireless instrumentation is not recommended for mission-critical control applications, and finds its greatest use in low-priority monitoring instrumentation. The most obvious advantage of wireless instruments is that they do not require wires of any kind. Since wiring is a major capital cost when installing instruments, this fact makes wireless instrumentation relatively inexpensive to install. Freedom from wires also allows these instruments to be used in applications that would be impossible for wired instruments, such as communicating data from sensors installed in moving vehicles to stationary monitoring or control equipment. However, the elimination of wires means wireless instruments must provide for their own power requirements, usually with long-life batteries. Reliance on battery power alone places restrictions on how frequently these instrument perform their functions: less frequent data transmission results in longer battery life, but correspondingly reduces the instrument's practicality for real-time control. Potential blockage of the radio signals from moving objects such as large vehicles (cranes, lifts, etc.) also poses challenges to signal reliability. Despite these limitations, the total absence of signal or power wiring for a wireless instrument is an essential feature for certain applications. Wireless is just another tool to help us automate processes, and like any other tool it has its advantages and disadvantages.

6.4 Other types of instruments

So far we have just looked at instruments that sense, control, and influence process variables. Transmitters, controllers, and control valves are respective examples of each instrument type. However, other instruments exist to perform useful functions for us.

6.4.1 Indicators

One common “auxiliary” instrument is the *indicator*, the purpose of which is to provide a human-readable indication of an instrument signal. Quite often process transmitters are not equipped with readouts for whatever variable they measure: they just transmit a standard instrument signal (3 to 15 PSI, 4 to 20 mA, etc.) to another device. An indicator gives a human operator a convenient way of seeing what the output of the transmitter is without having to connect test equipment (pressure gauge for 3-15 PSI, ammeter for 4-20 mA) and perform conversion calculations. Moreover, indicators may be located far from their respective transmitters, providing readouts in locations more convenient than the location of the transmitter itself. An example where remote indication would be practical is shown here, in a nuclear reactor temperature measurement system:



It would be unsafe for human beings to approach the nuclear reactor when it is in full-power operation, due to the strong radiation flux it emits. The temperature transmitter is built to withstand the radiation, though, and it transmits a 4 to 20 milliamp electronic signal to an indicating recorder located on the other side of a thick concrete wall blocking the reactor’s radiation, where it is safe for human occupancy. There is nothing preventing us from connecting multiple indicators, at multiple locations, to the same 4 to 20 milliamp signal wires coming from the temperature

transmitter. This allows us to display the reactor temperature in as many locations as we desire, since there is no absolute limit on how far we may conduct a DC milliamp signal along copper wires.

A numerical-plus-bargraph indicator appears in this next photograph, mounted in the face of a metal panel inside of a control room:



This particular indicator shows the position of a flow-control gate in a wastewater treatment facility, both by numerical value (98.06%) and by the height of a bargraph (very near full open – 100%). It is directly wired in series with the same 4-20 milliamp current signal sent to the gate actuator.

A less sophisticated style of panel-mounted indicator shows only a numeric display, such as this unit shown here:




Indicators may also be used in “field” (process) areas to provide direct indication of measured variables if the transmitter device lacks a human-readable indicator of its own. The following photograph shows a field-mounted indicator, operating directly from the electrical power available in the 4-20 mA loop. The numerical display of this indicator uses LCD technology rather than red-glowing LEDs, in order to use less electrical power:




6.4.2 Recorders

Another common “auxiliary” instrument is the *recorder* (sometimes specifically referred to as a *chart recorder* or a *trend recorder*), the purpose of which is to draw a graph of process variable(s) over time. Recorders usually have indications built into them for showing the instantaneous value of the instrument signal(s) simultaneously with the historical values, and for this reason are usually designated as *indicating* recorders. A temperature indicating recorder for the nuclear reactor system shown previously would be designated as a “TIR” accordingly.

Paper chart recorders are a form of instrumentation with a long history. The following image shows an illustration of a Bristol brand recording pressure gauge found on page 562 of *Cassier's Magazine* volume 8, published in 1895. Note the circular form of the paper chart, allowing the pen to draw a trace as the circular chart slowly spins. A padlock on the front glass cover prevents tampering with the chart recording:




A typical chart from one of these recording devices is shown in this illustration, taken from page 563 of the same engineering periodical:




This particular recording is of a steam boiler's pressure over a 24-hour period, showing pressure build-up beginning at 4:00 AM and boiler shut-down at 9:30 PM. This steam boiler's pressure appears to have been operated at approximately 70 PSI. Dips and peaks in the trace reflect changes in steam demand as well as irregularities in the firing of the boiler's coal furnace.

Another design of paper chart recorder is the *strip* style, using a long strip of paper between two spools (one to play out the paper and another to take it up). Like the circular chart recorder design, the strip chart recorder also has a long history. The following illustration from page 560 of the same *Cassier's Magazine* volume:



This next image shows a practical example of a strip chart's record for a city water supply company, taken from page 566 of the same periodical:



Several interesting details² may be seen on this particular recording, including pressure fluctuations caused by periodic draws of water from a fire hydrant to fill street carts used to spray

² Many interesting points may be drawn from these two illustrations. Regarding the strip chart recording instrument itself, it is worth noting the ornate design of the metal frame (quite typical of machinery design from that era), the attractive glass dome used to shield the chart and mechanism from the environment, and the intricate mechanism used to drive the strip chart and move the pen. Unlike a circular chart, the length of a strip chart is limited only by the diameter of the paper roll, and may be made long enough to record many days' worth of pressure measurements. The label seen on the front of this instrument ("Edson's Recording and Alarm Gauge") tells us this instrument has the ability to alert a human operator of abnormal conditions, and a close inspection of the mechanism reveals a bell

the city's dirt roads with water to minimize dust. Pressure drop caused by a burst 4-inch water pipe is also seen on this recording, between 5:00 PM and 6:15 PM.

Both circular and strip chart recorder designs survive to this day. Two circular chart recorders are shown in the following photograph, used to record temperatures at a powdered milk processing facility:




Two more chart recorders appear in the next photograph, a strip chart recorder on the right and a *paperless* chart recorder on the left. The strip chart recorder uses a scroll of paper drawn slowly past one or more lateral-moving pens, while the paperless recorder does away with paper entirely by plotting graphic trend lines on a computer screen:



on the top which presumably rings under alarm conditions. Regarding the strip chart record, note the “compressed” scale, whereby successive divisions of the vertical scale become closer in spacing, reflecting some inherent nonlinearity of the pressure-sensing mechanism.


Recorders are extremely helpful for troubleshooting process control problems. This is especially true when the recorder is configured to record not just the process variable, but also the controller's setpoint and output variables as well. Here is an example of a typical "trend" showing the relationship between process variable, setpoint, and controller output in automatic mode, as graphed by a recorder:



Here, the setpoint (SP) appears as a perfectly straight red line, the process variable as a slightly erratic blue line, and the controller output as a moderately erratic purple line. We can see from this trend that the controller is doing exactly what it should: holding the process variable value close to setpoint, manipulating the final control element as far as necessary to do so. The chaotic appearance of the output signal is not really a problem, contrary to most peoples' first impression. The fact that the process variable never deviates significantly from the setpoint tells us the control system is operating quite well. What accounts for the erratic controller output, then? The answer to this question is any factor requiring the controller to change its output in order to maintain the process variable at setpoint. Variations in process load would account for this: as other variables in the process change over time, the controller is forced to compensate for these variations in order to ensure the process variable does not drift from setpoint. Referencing our previous example of a steam boiler water level control system, one of these influencing variables is steam demand. If this trend shows the steam drum water level (PV) and feedwater valve position (Output), variations in the controller's output signal could be indicative of steam demand rising and falling, the controller modulating water flow into the boiler to properly compensate for this load and maintain a fairly constant steam drum water level. A wildly-fluctuating output signal *may* indicate a problem somewhere else in the process (placing undue demands on the control system), but there is certainly no problem with the control system itself: it is doing its job perfectly well.


Recorders become powerful diagnostic tools when coupled with the controller's manual control mode. By placing a controller in "manual" mode and allowing direct human control over the final control element (valve, motor, heater), we can tell a lot about a process. Here is an example of a

trend recording for a process in manual mode, where the process variable response is seen graphed in relation to the controller output as that output is increased and decreased in steps:



Notice the time delay between when the output signal is “stepped” to a new value and when the process variable responds to the change. This delay is called *dead time*, and it is generally detrimental to control system performance. Imagine trying to steer an automobile whose front wheels respond to your input at the steering wheel only after a 5-second delay! This would be a very challenging car to drive. The same problem plagues any industrial control system with a time lag between the final control element and the transmitter. Typical causes of this problem include *transport delay* (where there is a physical delay resulting from transit time of a process medium from the point of control to the point of measurement) and mechanical problems in the final control element.

This next example shows another type of problem revealed by a trend recording during manual-mode testing:



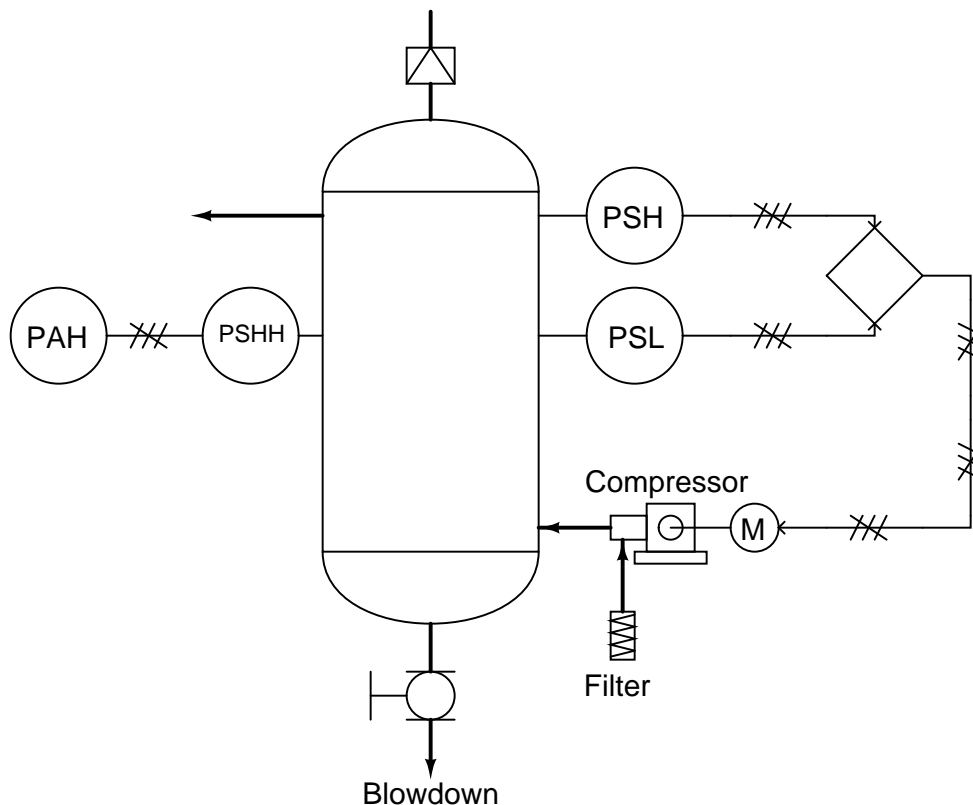
Here, we see the process quickly responding to all step-changes in controller output except for those involving a change in direction. This problem is usually caused by mechanical friction in the final control element (e.g. “sticky” valve stem packing in a pneumatically-actuated control valve), and is analogous to “loose” steering in an automobile, where the driver must turn the steering wheel a little bit extra after reversing steering direction. Anyone who has ever driven an old farm tractor knows what this phenomenon is like, and how it detrimentally affects one’s ability to steer the tractor in a straight line.

Sometimes it becomes useful to temporarily place a recorder into an instrumentation system for diagnostic purposes. On the simplest level, this might consist of a digital multimeter (DMM) connected to measure signal voltage or current, with its “minimum/maximum” capture mode engaged. On a more complex level, this might be a personal computer running data graphing software, connected to the instrumentation circuit through a data acquisition (DAQ) module converting the circuit’s analog voltage or current signals into digital values readable by the computer.

6.4.3 Process switches and alarms

Another type of instrument commonly seen in measurement and control systems is the *process switch*. The purpose of a switch is to turn on and off with varying process conditions. Usually, switches are used to activate alarms to alert human operators to take special action. In other situations, switches are directly used as control devices.

The following P&ID of a compressed air control system shows both uses of process switches:




The “PSH” (*pressure switch, high*) activates when the air pressure inside the vessel reaches its high control point. The “PSL” (*pressure switch, low*) activates when the air pressure inside the vessel drops down to its low control point. Both switches feed discrete (on/off) electrical signals to a logic control device (symbolized by the diamond) which then controls the starting and stopping of the electric motor-driven air compressor.

Another switch in this system labeled “PSHH” (*pressure switch, high-high*) activates only if the air pressure inside the vessel exceeds a level beyond the high shut-off point of the high pressure control switch (PSH). If this switch activates, something has gone wrong with the compressor control system, and the high pressure alarm (PAH, or *pressure alarm, high*) activates to notify a human operator.

All three switches in this air compressor control system are directly actuated by the air pressure in the vessel: in other words, these are direct *process-sensing* switches. It is possible, however, to build switch devices that interpret standardized instrumentation signals such as 3 to 15 PSI (pneumatic) or 4 to 20 milliamps (analog electronic), allowing us to build on/off control systems and alarms for *any* type of process variable measurable with a transmitter.

For example, the chlorine wastewater disinfection system shown earlier may be equipped with a couple of electronic alarm switches to alert an operator if the chlorine concentration ever exceeds pre-determined high or low limits:



The labels “AAL” and “AAH” refer to *analytical alarm low* and *analytical alarm high*, respectively. Note how the diagram shows these two alarm units connected to the electronic (4-20 mA) signal output by the chlorine analyzer (AT). This tells us the AAL and AAH alarm units are really just electronic circuits, alarming if the analytical transmitter’s 4-20 mA analog signal falls below (AAL) or exceeds (AAH) certain pre-set limits. As such, the AAL and AAH alarms do not *directly* sense the chlorine concentration in the water, but rather *indirectly* sense it by monitoring the chlorine analyzer’s 4-20 milliamp output signal.

Since both alarms work off the 4 to 20 milliamp electronic signal output by the chlorine analytical transmitter (AT) rather than directly sensing the process, their construction is greatly simplified. If these were process-sensing switches, each one would have to be equipped with the analytical capability of directly sensing chlorine concentration in water. In other words, each switch would have to be its own self-contained chlorine concentration analyzer, with all the attendant complexity.


An example of an electronic alarm module (triggered by a 4-20 mA current signal coming from a transmitter) is the Moore Industries model SPA (“Site Programmable Alarm”), shown here:



In addition to providing alarm capability, this SPA module also provides a digital display (a small LCD screen) to show the analog signal value for operational or diagnostic purposes.


Like all current-operated alarm modules, the Moore Industries SPA may be configured to “trip” electrical contacts when the current signal reaches a variety of different programmed thresholds. Some of the alarm types provided by this unit include high process, low process, out-of-range, and high rate-of-change.

In a similar manner, we may add pressure-actuated process alarm switches to pneumatic (3-15 PSI) signal lines coming from pneumatic transmitters to add alarm capability to a system designed for continuous measurement. For example, if high- and low-alarm capability were desired for the steam drum water level process described earlier in this chapter, one could add a pair of pressure-actuated switches to the pneumatic level transmitter's 3-15 PSI output signal line:



These two pressure-actuated switches serve as *water level* alarms, because the air pressure signal actuating them comes from the pneumatic *level* transmitter, which outputs an air pressure signal in direct proportion to water level in the boiler's steam drum. Even though the physical stimulus actuating each switch is an air *pressure*, the switches still serve the purpose of liquid *level* alarm signaling because that air pressure is an analogue (representation) of water level in the steam drum. In other words, these two alarm switches (LSL and LSH) *indirectly* sense water level by monitoring the pneumatic signal pressure output by the level transmitter (LT).

The alternative to pressure-actuated water level alarm switches would be independent level-sensing switches attached directly to the steam drum, each switch equipped with its own means³ of directly sensing water level:



It should be mentioned that the choice between using process alarm switches directly actuated by the process versus alarm switches actuated by a transmitter's analog signal is not arbitrary. In the system where the two alarm switches actuate from the transmitter's 3-15 PSI output signal, the integrity of the water level control *and* that of the high- and low-level alarms all depend on the proper function of one transmitter. If that one transmitter were to fail, all three system functions would be compromised. This elevates the importance of a single instrument, which is generally not desirable from the perspective of reliability and process safety. In the system where each level alarm switch independently senses steam drum water level, one device may fail without compromising either of the other two functions. This independence is desirable because it greatly reduces the probability of "common-cause" failures, where a single fault disables multiple system functions. The final determination should be based on a rigorous analysis of device versus system reliability, which is typically the task of a process engineer.

³These might be float-driven switches, where each switch is mechanically actuated by the buoyancy of a hollow metal float resting on the surface of the water. Another technology uses metal electrodes inserted into the water from above, sensing water level by electrical conductivity: when the water level reaches the probe's tip, an electrical circuit is closed. For more information on liquid level switches, refer to section 9.6 beginning on page 673.

Process alarm switches may be used to trigger a special type of indicator device known as an *annunciator*. An annunciator is an array of indicator lights and associated circuitry designed to secure a human operator's attention⁴ by blinking and sounding an audible buzzer when a process switch actuates into an abnormal state. The alarm state may be then "acknowledged" by an operator pushing a button, causing the alarm light to remain on (solid) rather than blink, and silencing the buzzer. The indicator light does not turn off until the actual alarm condition (the process switch) has returned to its regular state.

This photograph shows an annunciator located on a control panel for a large engine-driven pump. Each white plastic square with writing on it is a translucent pane covering a small light bulb. When an alarm condition occurs, the respective light bulb flashes, causing the translucent white plastic to glow, highlighting to the operator which alarm is active:



Note the two pushbutton switches below labeled "Test" and "Acknowledge." Pressing the "Acknowledge" button will silence the audible buzzer and also turn any blinking alarm light into a steady (solid) alarm light until the alarm condition clears, at which time the light turns off completely. Pressing the "Test" button turns all alarm lights on, to ensure all light bulbs are still functional.

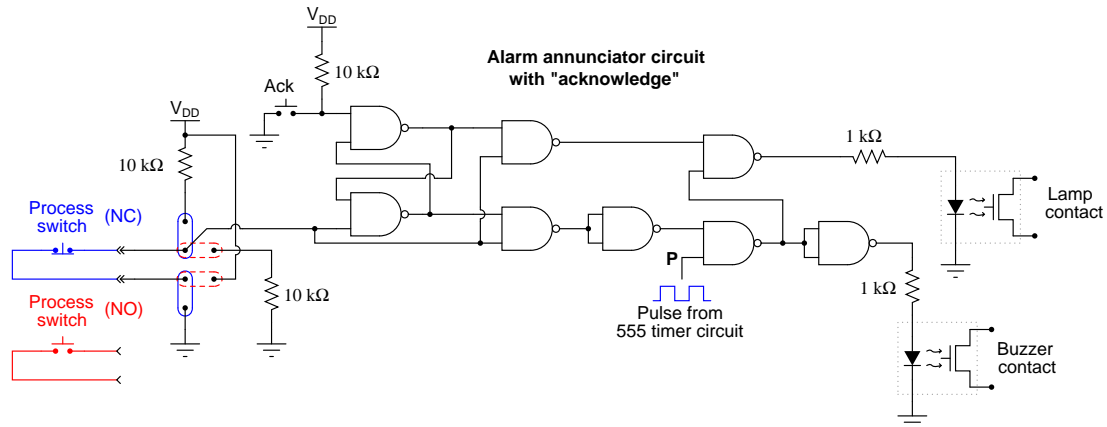
⁴D.A. Strobhar, writing in *The Instrument Engineers' Handbook* on the subject of alarm management, keenly observes that alarms are the only form of instrument "whose sole purpose is to alter the operator's behavior." Other instrument devices work to control the process, but only alarms work to *control the human operator*.

Opening the front panel of this annunciator reveals modular relay units controlling the blinking and acknowledgment latch functions, one for each alarm light:



This modular design allows each alarm channel to be serviced without necessarily interrupting the function of the other channels in the annunciator panel.

A simple logic gate circuit illustrates the acknowledgment latching feature (here implemented by an S-R latch circuit) common to all process alarm annunciators:



Panel-mounted annunciators are becoming a thing of the past, as computer-based alarm displays replace them with advanced capabilities such as time logging, first-event⁵ recording, and multiple layers of acknowledgment/access. Time logging is of particular importance in the process industries, as the sequence of events is often extremely important in investigations following an abnormal operating condition. Knowing *what happened*, and *exactly when it happened* is much more informative than simply knowing which alarms have tripped.

⁵When a complex machine or process with many shutdown sensors automatically shuts down, it may be difficult to discern after the fact *which* shutdown device was responsible. For instance, imagine an engine-powered generator automatically shutting down because one of the generator's "trip" sensors detected an under-voltage condition. Once the engine shuts down, though, *multiple* trip sensors will show abnormal conditions simply because the engine is not running anymore. The oil pressure sensor is one example of this: once the engine shuts down, there will no longer be any oil pressure, thus causing that alarm to activate. The under-voltage alarm falls into this category as well: once the engine shuts down, the generator will no longer be turning and therefore its output voltage *must* be zero. The problem for any human operator encountering the shut-down engine is that he or she cannot tell *which* of these alarms was the initiating cause of the shutdown versus *which* of these alarms simply activated after the fact once the engine shut off. An annunciator panel showing both an under-voltage and a low oil pressure light does not tell us which event happened first to shut down the generator. A "first-event" (sometimes called a "first-out") annunciator, however, shows which trip sensor was the *first* to activate, thus revealing the initiating cause of the event.

6.5 Summary

Instrument technicians maintain the safe and efficient operation of industrial measurement and control systems. This career requires a broad command of technical skill. Instrumentation is more than just physics or chemistry or mathematics or electronics or mechanics or control theory or risk analysis or troubleshooting alone. An instrument technician must know all these things to some degree, and more importantly how to synthesize and apply this knowledge to real applications.

The technical diversity of this profession is daunting. Adding to this challenge is the continued adoption of new technologies. The advent of new technologies, however, does not necessarily relegate legacy technologies to the scrap heap. It is quite common to find state-of-the-art instruments in the very same facility as decades-old instruments; digital fieldbus networks installed alongside 3 to 15 PSI pneumatic signal tubes; microprocessor-based sensors mounted right next to old mercury tilt-switches. Thus, the competent instrument technician must be comfortable working with both old and new technologies, understanding their merits, weaknesses, and especially their interactions.

This is why the most important skill for an instrument technician is the ability to teach oneself. It is impossible to fully prepare for a career like this with any amount of preparatory schooling. The profession is so broad and the responsibility so great, and the landscape so continuously subject to change, that life-long learning for the instrument technician is a matter of professional survival.

Perhaps the single greatest factor determining a person's ability to independently learn is their skill at *reading*. Being able to “digest” the written word is *the* key to learning what is difficult or impractical to directly experience. In an age where information is readily accessible, the skilled reader has the advantage of leveraging generations of experts in virtually any subject. Best of all, reading is a skill anyone can master, and everyone should.

My advice to all those desiring to become self-directed learners is to build a library of reading material on subjects that interest you (hopefully, instrumentation is one of those subjects!), and then immerse yourself in those writings. Feel free to “mark up”⁶ your books, or take notes in a separate location, so as to actively engage in your reading. Try as much as possible to approach reading as though you were having a *conversation* with the author: pose questions, challenge concepts and ideas, and do not stop doing so until you can clearly see what the author is trying to say.

I also advise *writing* about what you learn, because explaining new ideas in your own words helps you consolidate the learning, and “makes it your own” in a way few other activities do. You don't necessarily have to write your own book, but the act of expressing what you have learned is a powerful tool not only for building understanding, but also for revealing what you do not (yet) know. A method I have used with great success is to imagine myself having to explain a new concept to a precocious child: someone with enough mental capacity to grasp the concept but lacking the necessary vocabulary and experience to grasp a sophisticated presentation of it. This mental exercise forces you to explain things as simply as possible without error (because anyone can devise an explanation that is both simple and wrong!). All teachers know the power of this technique: you never learn a subject as well as when you must teach it to someone else.

⁶A fun and informative essay to read on this subject is Mortimer Adler's *How to Mark a Book*, widely disseminated on the Internet. In it, Adler argues persuasively for the habit of annotating the books you read, and gives some practical tips for doing so. He says reading a book should be a sort of *conversation with the author* where the flow of information is not just from the author to you, but also from you to yourself as you question, consider, and even argue the author's points.

6.6 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Representative signal:** using a signaling medium such as compressed air, electric current, or voltage pulses to represent some range of measured variable.
- **Common-cause failures:** when multiple functions in a system depend on a single element, failure of that element will cause all dependent functions to fail. Relevant to design of process alarm switches.
- **Negative feedback:** when the output of a system is degeneratively fed back to the input of that same system, the result is decreased (overall) gain and greater stability. Relevant to loop controller action: in order for a control system to be stable, the feedback must be negative.

References

- Adler, Mortimer, "How to Mark a Book", *The McGraw-Hill Reader*, McGraw-Hill Book Company, New York, NY, 1982.
- Hague, Charles A. "The Recording Gauge Applied to Water Pressure and Other Uses", *Cassier's Magazine* Volume 8, 1895.
- Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Software and Digital Networks*, Third Edition, CRC Press, New York, NY, 2002.

Chapter 7

Instrumentation documents

Every technical discipline has its own standardized way(s) of making descriptive diagrams, and instrumentation is no exception. The scope of instrumentation is so broad, however, that no one form of diagram is sufficient to capture all we might need to represent. This chapter will discuss three different types of instrumentation diagrams:

- Process Flow Diagrams (PFDs)
- Process and Instrument diagrams (P&IDs)
- Loop diagrams (“loop sheets”)
- Functional diagrams

At the highest level, the instrument technician is interested in the interconnections of process vessels, pipes, and flow paths of process fluids. The proper form of diagram to represent the “big picture” of a process is called a *process flow diagram*. Individual instruments are sparsely represented in a PFD, because the focus of the diagram is the process itself.

At the lowest level, the instrument technician is interested in the interconnections of individual instruments, including all the wire numbers, terminal numbers, cable types, instrument calibration ranges, etc. The proper form of diagram for this level of fine detail is called a *loop diagram*. Here, the process vessels and piping are sparsely represented, because the focus of the diagram is the instruments themselves.

Process and instrument diagrams (P&IDs) lie somewhere in the middle between process flow diagrams and loop diagrams. A P&ID shows the layout of all relevant process vessels, pipes, and machinery, but with instruments superimposed on the diagram showing what gets measured and what gets controlled. Here, one can view the flow of the process as well as the “flow” of information between instruments measuring and controlling the process.

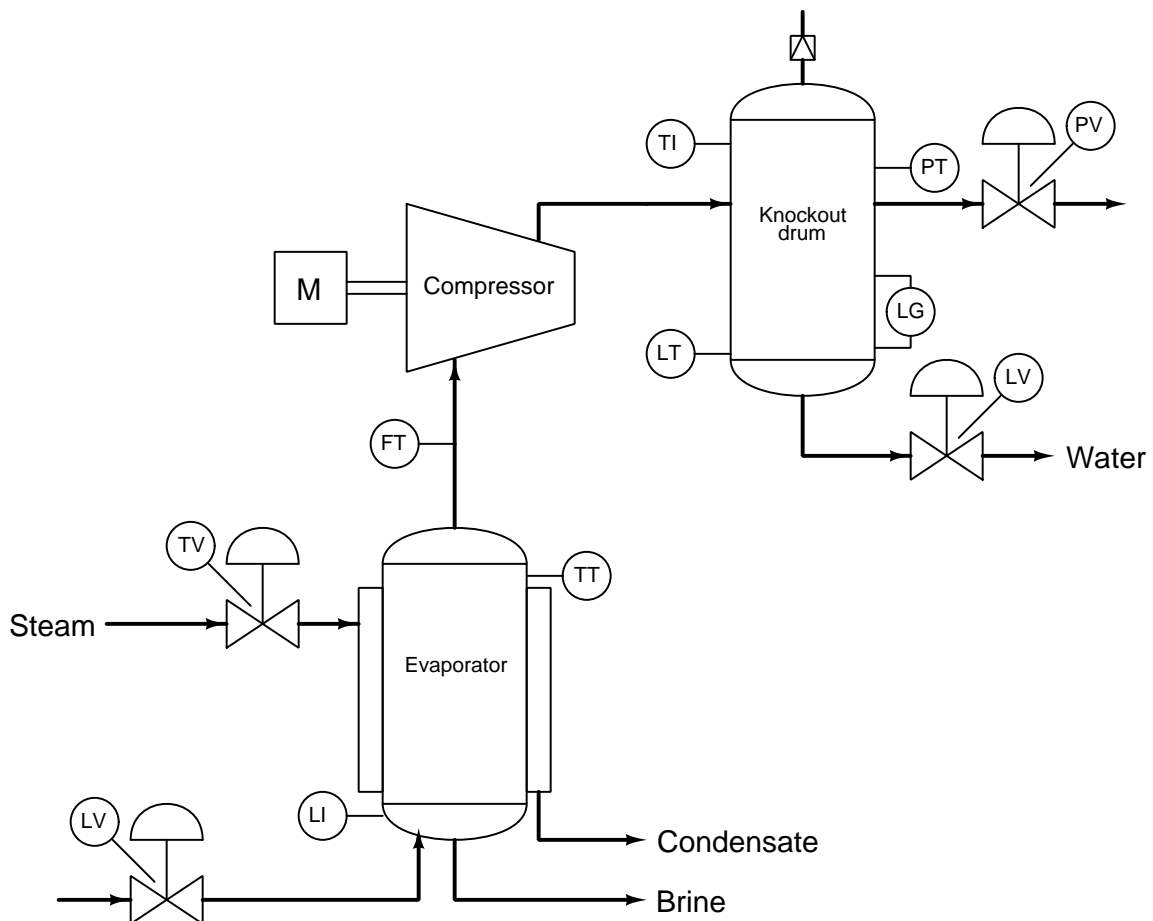
Functional diagrams are used for an entirely different purpose: to document the *strategy* of a control system. In a functional diagram, emphasis is placed on the algorithms used to control a process, as opposed to piping, wiring, or instrument connections. These diagrams are commonly found within the power generation industry, but are sometimes used in other industries as well.

An instrument technician must often switch between different diagrams when troubleshooting a complex control system. There is simply too much detail for any one diagram to show everything. Even if the page were large enough, a “show everything” diagram would be so turgid with details that it would be difficult to focus on any particular grouping of details you happened to be interested in. The narrowing of scope with the progression from PFD to loop diagram may be visualized as a process of “zooming in,” as though one were viewing a process through the lens of a microscope at different powers. First you begin with a PFD or P&ID to get an overview of the process, to see how the major components interact. Then, once you have identified which instrument “loop” you need to investigate, you go to the appropriate loop diagram to see the interconnection details of that instrument system so you know where to connect your test equipment and what signals you expect to find when you do.

Another analogy for this progression of documents is a map, or more precisely, a globe, an atlas, and a city street map. The globe gives you the “big picture” of the Earth, countries, and major cities. An atlas allows you to “zoom in” to see details of particular provinces, states, and principalities, and the routes of travel connecting them all. A city map shows you major and minor roads, canals, alleyways, and perhaps even some addresses in order for you to find your way to a particular destination. It would be impractical to have a globe large enough to show you all the details of every city! Furthermore, a globe comprehensive enough to show you all these details would have to be updated *very* frequently to keep up with all cities’ road changes. There is a certain economy inherent to the omission of fine details, both in ease of use and in ease of maintenance.

7.1 Process Flow Diagrams

To show a practical process example, let's examine three diagrams for a compressor control system, beginning with a Process Flow Diagram, or PFD. In this fictitious process, water is being evaporated from a process solution under partial vacuum (provided by the compressor). The compressor then transports the vapors to a "knockout drum" where they condense into liquid form. As a typical PFD, this diagram shows the major interconnections of process vessels and equipment, but omits details such as instrument signal lines and auxiliary instruments:




One might guess the instrument interconnections based on the instruments' labels. For instance, a good guess would be that the level transmitter (LT) on the bottom of the knockout drum might send the signal that eventually controls the level valve (LV) on the bottom of that same vessel. One might also guess that the temperature transmitter (TT) on the top of the evaporator might be part of the temperature control system that lets steam into the heating jacket of that vessel.

Based on this diagram alone, one would be hard-pressed to determine what control system, if

any, controls the compressor itself. All the PFD shows relating directly to the compressor is a flow transmitter (FT) on the suction line. This level of uncertainty is perfectly acceptable for a PFD, because its purpose is merely to show the general flow of the process itself, and only a bare minimum of control instrumentation.

7.2 Process and Instrument Diagrams

The next level of detail is the Process and Instrument Diagram¹, or P&ID. Here, we see a “zooming in” of scope from the whole evaporator process to the compressor as a unit. The evaporator and knockout vessels almost fade into the background, with their associated instruments absent from view²:



Now we see there is more instrumentation associated with the compressor than just a flow transmitter. There is also a differential pressure transmitter (PDT), a flow indicating controller (FIC), and a “recycle” control valve allowing some of the vapor coming out of the compressor’s


¹Sometimes P&ID stands for *Piping* and Instrument Diagram. Either way, it means the same thing.

²It should be noted that the “zooming in” of scope in a P&ID does not necessarily mean the scope of other areas of the process must be “zoomed out.” In fact, it is rather typical in a P&ID that the *entire* process system is shown in finer detail than in a PFD, but not all on one page. In other words, while a PFD may depict a process in its entirety on one piece of paper, a comprehensive P&ID will typically span multiple pieces of paper, each one detailing a section of the process system.

discharge line to go back around into the compressor's suction line. Additionally, we have a pair of temperature transmitters reporting suction and discharge line temperatures to an indicating recorder.

Some other noteworthy details emerge in the P&ID as well. We see that the flow transmitter, flow controller, pressure transmitter, and flow valve all bear a common number: 42. This common "loop number" indicates these four instruments are all part of the same control system. An instrument with any other loop number is part of a different control system, measuring and/or controlling some other function in the process. Examples of this include the two temperature transmitters and their respective recorders, bearing the loop numbers 41 and 43.

Please note the differences in the instrument "bubbles" as shown on this P&ID. Some of the bubbles are just open circles, where others have lines going through the middle. Each of these symbols has meaning according to the ISA (Instrumentation, Systems, and Automation society) standard:



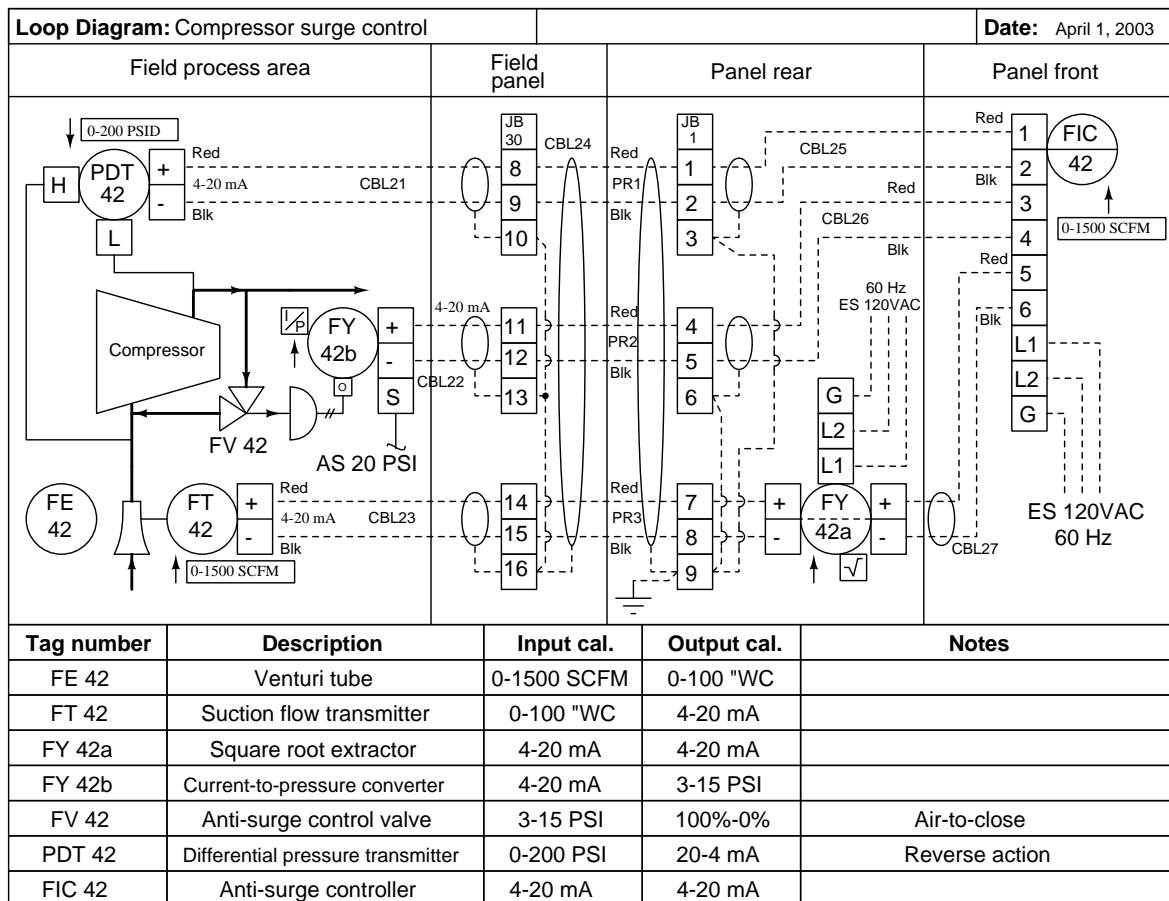
The type of "bubble" used for each instrument tells us something about its location. This, obviously, is quite important when working in a facility with many thousands of instruments scattered over acres of facility area, structures, and buildings.

The rectangular box enclosing both temperature recorders shows they are part of the same physical instrument. In other words, this indicates there is really only one temperature recorder instrument, and that it plots both suction and discharge temperatures (most likely on the same trend graph). This suggests that each bubble may not necessarily represent a discrete, physical instrument, but rather an instrument *function* that may reside in a multi-function device.

Details we do not see on this P&ID include cable types, wire numbers, terminal blocks, junction boxes, instrument calibration ranges, failure modes, power sources, and the like. To examine this level of detail, we must turn to another document called a *loop diagram*.

7.3 Loop diagrams

Finally, we arrive at the loop diagram (sometimes called a *loop sheet*) for the compressor surge control system (loop number 42):



Here we see that the P&ID didn't show us all the instruments in this control "loop." Not only do we have two transmitters, a controller, and a valve; we also have two signal transducers. Transducer 42a modifies the flow transmitter's signal before it goes into the controller, and transducer 42b converts the electronic 4 to 20 mA signal into a pneumatic 3 to 15 PSI air pressure signal. Each instrument "bubble" in a loop diagram represents an individual device, with its own terminals for connecting wires.

Note that dashed lines now represent individual copper wires instead of whole cables. Electrical terminals where these wires connect to are represented by squares with numbers in them. Fluid ports on instruments are also represented by labeled squares. Cable numbers, wire colors, junction block numbers, panel identification, and even grounding points are all shown in loop diagrams. The only type of diagram for this system more detailed than a loop diagram would be an electronic

schematic diagram for an individual instrument, which of course would only show details pertaining to that one instrument. Thus, the loop diagram is the most detailed form of diagram for a control system as a whole, and as such it must contain all details omitted by PFDs and P&IDs alike.

To the novice it may seem excessive to include such trivia as wire colors in a loop diagram. To the experienced instrument technician who has had to work on systems lacking such documented detail, this information is highly valued. The more detail you put into a loop diagram, the easier it makes the inevitable job of maintaining that system at some later date. When a loop diagram shows you exactly what wire color to expect at exactly what point in an instrumentation system, and exactly what terminal that wire should connect to, it becomes much easier to proceed with any troubleshooting, calibration, or upgrade task.


Loop diagrams are fairly constrained in their layout as per the ISA 5.1 standard. Field instruments are always placed on the left-hand side, while control-panel or control-room instruments must be located on the right-hand side. Text describing instrument tags, ranges, and notes are always placed on the bottom. Unlike PFDs and P&IDs where component layout is largely left to the whim of the designer drawing the diagram, loop sheets offer little room for creativity. This is intentional, as creativity and readability are mutually exclusive in cases where there is an immense amount of technical detail embedded in a diagram. It is simply easier to find details you're looking for when you know *exactly* where they ought to be.

An interesting detail seen on this loop diagram is an entry specifying “input calibration” and “output calibration” for each and every instrument in the system. This is actually a very important concept to keep in mind when troubleshooting a complex instrumentation system: every instrument has at least one input and at least one output, with some sort of mathematical relationship between the two. Diagnosing where a problem lies within a measurement or control system often means testing various instruments to see if their output responses appropriately match their input conditions, so it is important to document these input and output ranges.

For example, one way to test the flow transmitter in this system would be to subject it to a number of different pressures within its range (specified in the diagram as 0 to 100 inches of water column differential) and seeing whether or not the current signal output by the transmitter was consistently proportional to the applied pressure (e.g. 4 mA at 0 inches pressure, 20 mA at 100 inches pressure, 12 mA at 50 inches pressure, etc.).

Given the fact that a calibration error or malfunction in any one of these instruments can cause a problem for the control system as a whole, it is nice to know there is a way to determine which instrument is to blame and which instruments are not. This general principle holds true regardless of the instrument's type or technology. You can use the same input-versus-output test procedure to verify the proper operation of a pneumatic (3 to 15 PSI) level transmitter or an analog electronic (4 to 20 mA) flow transmitter or a digital (fieldbus) temperature transmitter alike. Each and every instrument has an input and an output, and there is always a predictable (and testable) correlation from one to the other.

Another interesting detail seen on this loop diagram is the *direction of action* of each instrument. You will notice a box and arrow (pointing either up or down) next to each instrument bubble. An “up” arrow (↑) represents a *direct-acting* instrument: one whose output signal increases as the input stimulus increases. A “down” arrow (↓) represents a *reverse-acting* instrument: one whose output signal decreases as the input stimulus increases. All the instruments in this loop are direct-acting with the exception of the pressure differential transmitter PDT-42:




Here, the “down” arrow tells us the transmitter will output a full-range signal (20 mA) when it senses zero differential pressure, and a 0% signal (4 mA) when sensing a full 200 PSI differential. While this calibration may seem confusing and unwarranted, it serves a definite purpose in this particular control system. Since the transmitter’s current signal decreases as pressure increases, and the controller must be correspondingly configured, a decreasing current signal will be interpreted by the controller as a high differential pressure. If any wire connection fails in the 4-20 mA current loop for that transmitter, the resulting 0 mA signal will be naturally “seen” by the controller as a pressure over-range condition. Excessive pressure drop across a compressor is considered dangerous because it may lead to the compressor surging³. Thus, the controller will naturally take action to prevent surge by commanding the anti-surge control valve to open, because it “thinks” the compressor is about to surge. In other words, the transmitter is intentionally calibrated to be reverse-acting such that any break in the signal wiring will naturally bring the system to its safest condition.

³Compressor “surge” is a violent and potentially self-destructing action experienced by a centrifugal compressor if the pressure drop across it becomes too high and the flow rate through it becomes too low. Surging may be prevented by opening up a “recycle” valve from the compressor’s discharge line to the suction line, ensuring adequate flow through the compressor while simultaneously unloading the high pressure differential across it.

7.4 Functional diagrams


A unique form of technical diagram for describing the abstract functions comprising a control system (e.g. PID controllers, rate limiters, manual loaders) is a *functional diagram*⁴. This form of document finds wide application in the power generation industry to document control strategies. Functional diagrams focus on the flow of information within a control system rather than on the process piping or instrument interconnections (wires, tubes, etc.). The general flow of a functional diagram is top-to-bottom, with the process sensing instrument (transmitter) located at the top and the final control element (valve or variable-speed motor) located at the bottom. No attempt is made to arrange symbols in a functional diagram to correspond with actual equipment layout: these diagrams are all about the *algorithms* used to make control decisions, and nothing more.

A sample functional diagram appears here, showing a flow transmitter (FT) sending a process variable signal to a PID controller, which then sends a manipulated variable signal to a flow control valve (FCV):



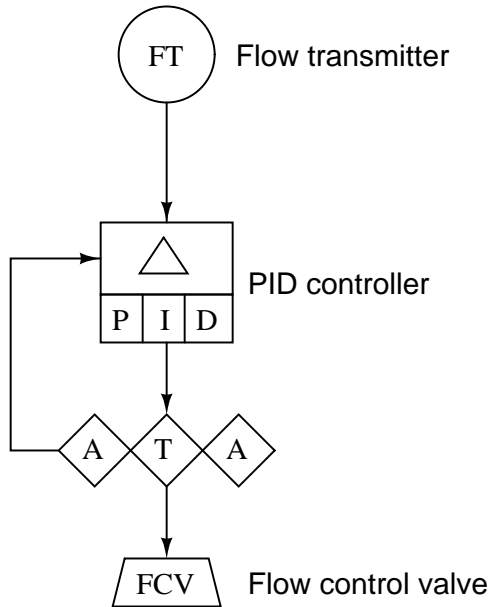
⁴Functional diagrams are sometimes referred to as *SAMA* diagrams in honor of the organization responsible for their standardization, the *Scientific Apparatus Makers Association*. This organization has been succeeded by the Measurement, Control, and Automation Association (MCAA), thus obsoleting the “SAMA” acronym.

A cascaded control system, where the output of one controller acts as the setpoint for another controller to follow, appears in functional diagram form like this:




In this case, the primary controller senses the level in a vessel, commanding the secondary (flow) controller to maintain the necessary amount of flow either in or out of the vessel as needed to maintain level at some setpoint.

Functional diagrams may show varying degrees of detail about the control strategies they document. For example, you may see the auto/manual controls represented as separate entities in a functional diagram, apart from the basic PID controller function. In the following example, we see a transfer block (T) and two manual adjustment blocks (A) providing a human operator the ability to separately adjust the controller's setpoint and output (manipulated) variables, and to transfer between automatic and manual modes:



Rectangular blocks such as the Δ , P, I, and D shown in this diagram represent automatic functions. Diamond-shaped blocks such as the A and T blocks are manual functions (i.e. set by a human operator). Showing even more detail, the following functional diagram indicates the presence of *setpoint tracking* in the controller algorithm, a feature that forces the setpoint value to equal the process variable value any time the controller is in manual mode:






















Here we see a new type of line: dashed instead of solid. This too has meaning in the world of functional diagrams. Solid lines represent analog (continuously variable) signals such as process variable, setpoint, and manipulated variable. Dashed lines represent discrete (on/off) signal paths, in this case the auto/manual state of the controller commanding the PID algorithm to get its setpoint either from the operator's input (A) or from the process variable input (the flow transmitter: FT).

7.5 Instrument and process equipment symbols


This section shows some of the many instrument symbols found in different types of technical diagrams used to document instrument systems.

7.5.1 Line types


Process flow line 	Instrument supply or process connection (impulse line) 	Waveguide 	Undefined 
Pneumatic signal (continuous) 	Pneumatic signal (discrete -- on/off) 	Capillary tube 	Hydraulic signal 
Electric signal (continuous)  (or) 	Electric signal (discrete -- on/off)  (or) 	Fieldbus network 	Data link (smart instrument) 
Mechanical link 	Data link (common system) 	Data link (independent systems) 	Radio link 
Sonic or other wave 			

Note: the single backslash signifying a “discrete” or “binary” signal type has been removed from the ISA standard as of the 2009 ANSI publication. Regular pneumatic and electrical line symbols may represent either continuous or discrete states. The “triple-slash” alternative linetype for electrical symbols is also absent from the 2009 ANSI/ISA standard.

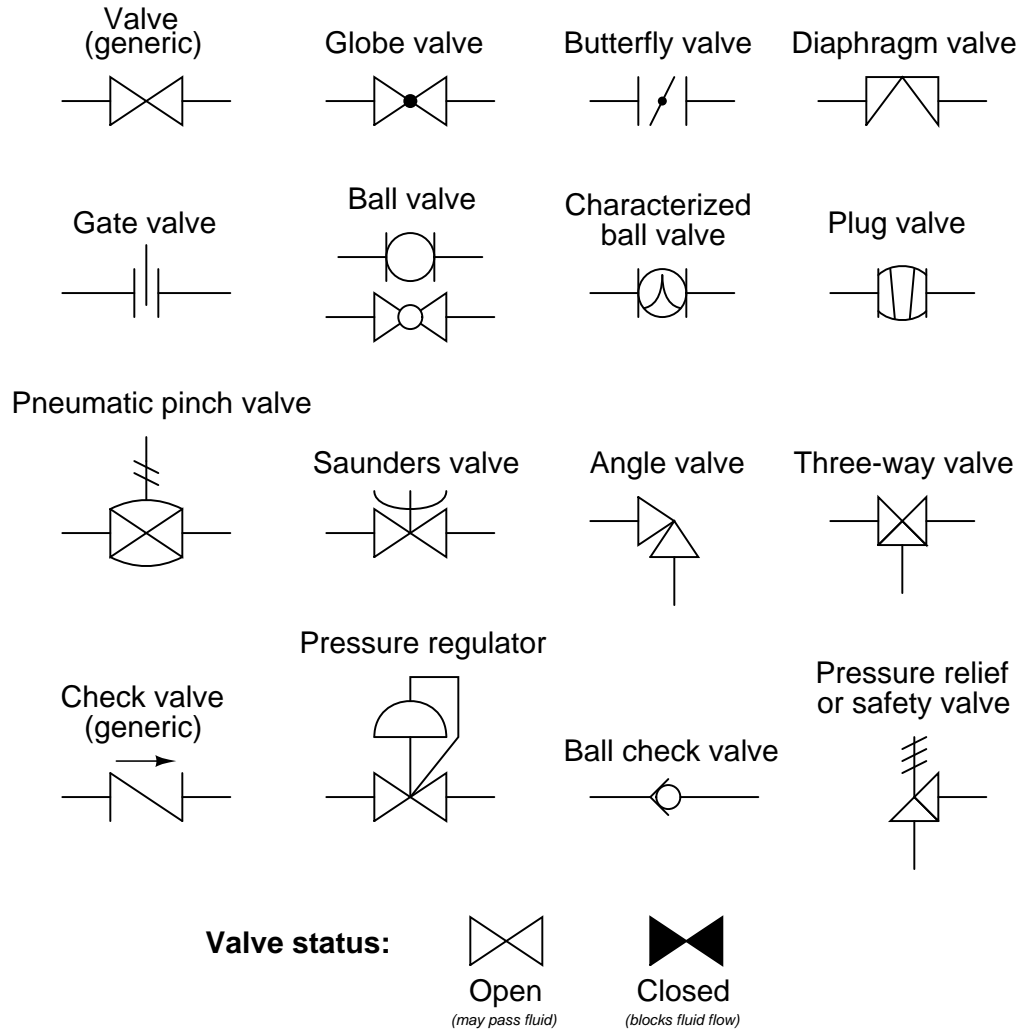
7.5.2 Process/Instrument line connections



7.5.3 Instrument bubbles













7.5.4 Process valve types









Valve status may or may not be shown in a process diagram. If you happen to see solid-colored valve symbols anywhere in a diagram, you know that status is being represented. If you see no solid-colored valves anywhere in the diagram, either all valves are shown open or else status is not represented at all.

7.5.5 Valve actuator types




7.5.6 Valve failure mode




7.5.7 Liquid level measurement devices


Bubbler (dip tube)




Capacitive




Tape-and-float




Hydrostatic




Hydrostatic (w/ seals)




Displacer




Radar (guided)




Radar (non-contact)




Ultrasonic




Laser




7.5.8 Flow measurement devices (flowing left-to-right)






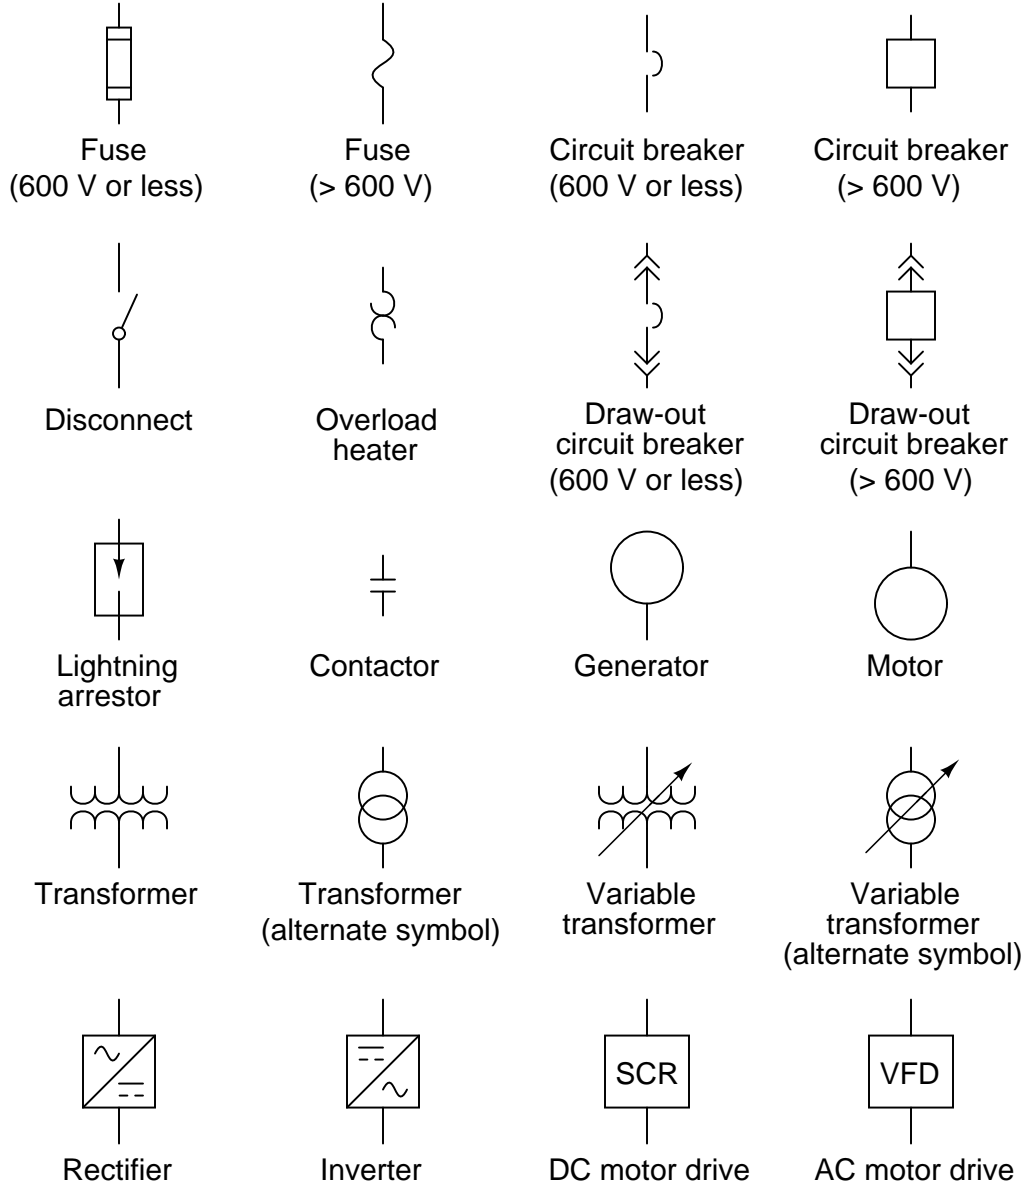
7.5.9 Process equipment















7.5.10 Functional diagram symbols












7.5.11 Single-line electrical diagram symbols




			
Voltmeter	Ammeter	Wattmeter	Frequency meter
			
Phase meter	VAR meter	Lamp	Kilowatt-hour meter
			
KiloVAR-hour meter	Current transformer (CT)	Potential transformer (PT)	Synchronization meter


7.5.12 Fluid power diagram symbols

Hydraulic pump
(fixed displacement)Hydraulic pump
(variable displacement)Hydraulic motor
(fixed displacement)Hydraulic motor
(variable displacement)Air compressor
(fixed displacement)Air compressor
(variable displacement)Air motor
(fixed displacement)Air motor
(variable displacement)Cylinder, single-acting
(ram)

Cylinder, double-acting



Cylinder, differential



Electric motor







Combustion engine



Accumulator




Filter

Fixed restriction,
laminar flowVariable restriction
laminar flowFixed restriction,
inviscid flow


Check valve



Fluid heater




Fluid cooler



Open reservoir




Closed reservoir



Various spool valve "box" symbols



Hand pump



Solenoid actuator



Pressure actuator



Lever actuator



Roller actuator




Button actuator




Return spring



Pressure relief (shunt regulator)




Pressure regulator (series)



Hydraulic line



Pneumatic line



7.6 Instrument identification tags

Up until this point, we have explored various types of instrumentation diagram, each one making reference to different instruments by lettered identifiers such as TT (Temperature Transmitter), PDT (Pressure Differential Transmitter), or FV (Flow Valve), without formally defining all the letters used to identify instruments. Part of the ISA 5.1 standard does exactly this, which is what we will now investigate.

Each instrument within an instrumented facility should have its own unique identifying *tag* consisting of a series of letters describing that instrument's *function*, as well as a number identifying the particular *loop* it belongs to. An optional numerical prefix typically designates the larger area of the facility in which the loop resides, and an optional alphabetical suffix designates multiple instances of instruments within one loop.

For example, if we were to see an instrument bearing the tag FC-135, we would know it was a *flow controller* (FC) for loop number 135. In a large manufacturing facility with multiple processing “unit” areas, a tag such as this might be preceded by another number designating the unit area. For example, our hypothetical flow controller might be labeled 12-FC-135 (flow controller for loop #135, located in unit #12). If this loop happened to contain multiple controllers, we would need to distinguish them from each other by the use of suffix letters appended to the loop number (e.g. 12-FC-135A, 12-FC-135B, 12-FC-135C).

Each and every instrument within a particular loop is first defined by the variable that loop seeks to sense or control, regardless of the physical construction of the instrument itself. Our hypothetical flow controller FC-135, for example, may be physically identical to the level controller in loop #72 (LC-72), or to the temperature controller in loop #288 (TC-288). What makes FC-135 a *flow* controller is the fact that the transmitter sensing the main process variable measures *flow*. Likewise, the identifying tag for every other instrument within that loop⁵ must begin with the letter “F” as well. This includes the final control element as well: in a level control loop, the transmitter is identified as an “LT” even if the actual sensing element works on *pressure* (because the variable that the loop strives to sense or control is actually level, despite the fact that liquid level is being inferred from pressure), the controller is identified as an “LC”, and the control valve throttling fluid *flow* is identified as an “LV”: every instrument in that level-controlling loop serves to help control *level*, and so its primary function is to be a “level” instrument.

⁵Exceptions do exist to this rule. For example, in a cascade or feedforward loop where multiple transmitters feed into one or more controllers, each transmitter is identified by the type of process variable *it* senses, and each controller's identifying tag follows suit.

Valid letters recognized by the ISA for defining the primary process variable of an instrument within a loop are shown in the following table. Please note that the use of a modifier defines a unique variable: for example, a “PT” is a transmitter measuring *pressure* at a single point in a process, whereas a “PDT” is a transmitter measuring a *pressure difference* between two points in a process. Likewise, a “TC” is a controller controlling temperature, whereas a “TKC” is a controller controlling the *rate-of-change of temperature*:

Letter	Variable	Modifier
A	Analytical (composition)	
B	Burner or Combustion	
C	<i>User-defined</i>	
D	<i>User-defined</i>	Differential
E	Voltage	
F	Flow	Ratio or Fraction
G	<i>User-defined</i>	
H	Hand (manual)	
I	Current	
J	Power	Scan
K	Time or Schedule	Time rate-of-change
L	Level	
M	<i>User-defined</i>	Momentary
N	<i>User-defined</i>	
O	<i>User-defined</i>	
P	Pressure or Vacuum	
Q	Quantity	Time-Integral or Total
R	Radiation	
S	Speed or Frequency	Safety
T	Temperature	
U	Multi-function	
V	Vibration	
W	Weight or Force	
X	<i>Unclassified</i>	X-axis
Y	Event, State, or Presence	Y-axis
Z	Position or Dimension	Z-axis

A “user-defined” letter represents a non-standard variable used multiple times in an instrumentation system. For example, an engineer designing an instrument system for measuring and controlling the *refractive index* of a liquid might choose to use the letter “C” for this variable. Thus, a refractive-index transmitter would be designated “CT” and a control valve for the refractive-index loop would be designated “CV”. The meaning of a user-defined variable need only be defined in one location (e.g. in a legend for the diagram).

An “unclassified” letter represents one or more non-standard variables, each used only once (or a very limited number of times) in an instrumentation system. The meaning of an unclassified variable is best described immediately near the instrument’s symbol rather than in a legend.

Succeeding letters in an instrument tag describe the function that instrument performs relative to the process variable. For example, a “PT” is an instrument *transmitting* a signal representing pressure, while a “PI” is an *indicator* for pressure and a “PC” is a *controller* for pressure. Many instruments have multiple functions designated by multiple letters, such as a TRC (Temperature *Recording Controller*). In such cases, the first function letter represents the “passive” function (usually provided to a human operator) while the second letter represents the “active” (automated) control function.

Letter	Passive function	Active function	Modifier
A	Alarm		
B	<i>User-defined</i>	<i>User-defined</i>	<i>User-defined</i>
C		Control	
E	Element (sensing)		
G	Glass or Viewport		
H			High
I	Indicate		
K		Control station	
L	Light		Low
M			Middle or Intermediate
N	<i>User-defined</i>	<i>User-defined</i>	<i>User-defined</i>
O	Orifice		
P	Test point		
R	Record		
S		Switch	
T		Transmit	
U	Multi-function	Multi-function	Multi-function
V		Valve, Damper, Louver	
W	Well		
X	<i>Unclassified</i>	<i>Unclassified</i>	<i>Unclassified</i>
Y		Relay, Compute, Convert	
Z		Driver, Actuator, or unclassified final control element	

A variety of other letter combinations are often used to identify details not standardized by the ISA. For example, chemical analyzer instruments often have their sample tube connections represented by the letter combination “SC,” although this does not appear anywhere in the ISA 5.1 standard.

Some examples of instrument tag letters are shown in the following list:

- **AIT** = Analytical Indicating Transmitter (*e.g. an oxygen concentration analyzer with a built-in display of oxygen percentage*)
- **ESL** = Voltage Switch, Low (*e.g. a switch used to detect an under-voltage condition in an electrical power system*)
- **FFI** = Flow Ratio Indicator (*e.g. a device indicating the ratio between air and fuel for a large industrial engine*)
- **FIC** = Flow Indicating Controller (*i.e. a controller designed to indicate flow to a human operator*)
- **HC** = Hand Controller (*i.e. a device allowing a human operator to set a control signal to some desired level, usually to operate a valve or other final control element*)
- **JQR** = Power Totalizing Recorder (*e.g. a watt-hour recorder, tracking total energy used*)
- **LSHH** = Level Switch, High-High (*e.g. a level-sensing switch designed to detect a dangerously high liquid level and initiate an automatic shutdown in that event*)
- **LT** = Level Transmitter (*i.e. a device sensing liquid level and reporting that level in some analog or digital form*)
- **PIT** = Pressure Indicating Transmitter (*e.g. a Rosemount model 3051 pressure transmitter with a built-in display of measured pressure*)
- **PDT** = Pressure Differential Transmitter (*i.e. a pressure transmitter built and installed to sense the difference of pressure between two points in a fluid system*)
- **PV** = Pressure Valve (*i.e. a control valve installed in a loop where the process variable is pressure*)
- **TE** = Temperature Element (*i.e. a sensing element used to directly detect the temperature of a process material; e.g. a thermocouple, thermistor, filled-bulb, bimetallic spring*)
- **TKAH** = Temperature Rate-of-change Alarm, High (*i.e. a device alarming when the rate of temperature change exceeds a pre-set limit*)
- **TV** = Temperature Valve (*i.e. a control valve installed in a loop where the process variable is temperature*)
- **TY** = Temperature Converter (*e.g. an I/P transducer in a temperature loop*)
- **VSH** = Vibration Switch, High (*i.e. a switch used to detect a high level of vibration on a piece of machinery*)
- **ZXI, ZYI, and ZZI** = Position Indicators for X, Y, and Z axes respectively (*e.g. indicators showing the three axis positions for a CNC machine tool*)

References

ANSI/ISA-5.1-2009, Instrumentation Symbols and Identification, Research Triangle Park, NC, 2009.

“Commonly Used Electrical Symbols”, Eaton Electrical Inc., Eaton Corporation, Moon Township, PA, 2005.

Instrumentation, Systems, and Automation Society Standards, 5.1-1984 (R1992), Instrumentation Symbols and Identification, Research Triangle Park, NC, 1984.

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Software and Digital Networks*, Third Edition, CRC Press, New York, NY, 2002.

Chapter 8

Instrument connections


All instruments connect to their respective processes and to each other by means of pipe, tube, and/or wires. Improper installation of these connective lines can make the difference between success or failure in an installation. Safety is also impacted by improper connections between instruments and the process, and from instrument to instrument.

8.1 Pipe and pipe fittings

Pipe is a hollow structure designed to provide an enclosed pathway for fluids to flow, usually manufactured from cast metal (although plastic is a common pipe material for many industrial applications). This section discusses some of the more common methods for joining pipes together (and joining pipe ends to equipment such as pressure instruments).


8.1.1 Flanged pipe fittings

In the United States of America, most large industrial pipes are joined together by *flanges*. A pipe “flange” is a ring of metal, usually welded to the end of a pipe, with holes drilled in it parallel to the pipe centerline to accept several bolts:



Flange joints are made pressure-tight by inserting a donut-shaped gasket between the flange pairs prior to tightening the bolts. Gaskets are manufactured from materials softer than the flange material. When sandwiched between a pair of flanges, the gasket will be “crushed” between them to seal all potential leak paths.

In instrument diagrams such as P&IDs, flanges are denoted by two short parallel lines, both perpendicular to the pipe. The pipe size of the flange is often written near the flange symbol, as is the case with this 8-inch flange symbol shown below:



A photograph showing a Rosemount magnetic flowmeter installed with 4-bolt flange fittings appears here:



If you examine the flanged connections closely, you can see the gap between the flange faces created by the thickness of the gasket material “sandwiched” between the flange pairs.


In this next photograph, we see a pair of large pipe flange connections on either end of a relatively short “spool” pipe section. The large number of studs holding each flange set together gives you some indication of the pressure of the fluid within, in this case upwards of 1000 PSI!



Like the flowmeter flanges shown previously, gaps between the flange ring faces reveal the space occupied by the gasket sealing those flange surfaces together to form a pressure-tight seal.

A common method of installing such a flange gasket is to first install only half of the bolts (in the holes lower than the centerline of the pipe), drop the gasket between the flanges, insert the remaining bolts, then proceed to tighten all bolts to the proper torques:

(All views shown end-wise)



Flanges differ with regard to their sealing design and required gasket type. In the United States, one of the most common flange “face” designs is the *raised-face* (RF) flange, designed to seal against a gasket by means of a set of concentric grooves machined on the face of the flange. These grooves form a sealing surface with far greater leakage path length than if the faces were smooth, thus discouraging leakage of process fluid under pressure.

Another flange face design is called *ring-type joint* (RTJ). In this design, a special metal ring sits inside a groove machined into the faces of both mating flanges, crushing and filling that groove when the flanges are properly tightened together. RTJ flanges are typically found on high-pressure applications where leakage control is more challenging. The grooves in RTJ flanges must be completely free of foreign material, and well-formed (not distorted) in order to achieve proper sealing.

In the United States, flanges are often rated according to a system of “pressure classes” defined in the ANSI (American National Standards Institute) standard 16.5. These pressure classes are designated by numerical values followed by “pound”, “lb”, or “#”. Common ANSI ratings include the 150#, 300#, 400#, 600#, 900#, 1500#, and 2500# pressure classes. It should be noted that these class numbers do *not* refer directly to pressure ratings in units of PSI, but that they do scale with pressure (i.e. a 600# flange will have a greater pressure rating than a 300# flange, all other factors being equal). Pressure ratings not only vary with the “class” of the flange, but also with operating temperature, as metals tend to weaken at elevated temperature.


Originally, the ANSI class designations were based on the ratings of these flanges in steam line service. A 250# flange, for instance, was rated such because it was designed to be used in piping service where the fluid was saturated steam at 250 PSI (and 400 degrees Fahrenheit). As metallurgy advanced, these flanges became capable of handling higher pressures at higher temperatures, but the original “pound” rating remained¹. This state of affairs is not unlike the “tonnage” rating of American light trucks: a “one-ton” truck is actually capable of hauling far more than 2000 pounds of cargo. The “one-ton” designation refers to a particular design which used to be rated for

¹EBAA Iron Sales, Inc published a two-page report in 1994 (“Connections” FL-01 2-94) summarizing the history of flange “pound” ratings, from the ASME/ANSI B16 standards.

approximately 2000 pounds, but through advances in metallurgy and manufacturing is now able to carry well over that rating.

Piping flanges and components must have matching flange ratings and sizes in order to properly function. For example, a control valve with a flanged body rated as a 4-inch ANSI class 300# can only be properly joined to another 4-inch ANSI class 300# pipe flange. The physical integrity of the piping system will be jeopardized if mis-matched pressure-class flanges are connected together. Proper gasket types must also be selected to coordinate with the pressure class of the mating flanges. Thus, each and every flanged joint must be considered a complete *system*, with integrity ensured only if all components comprising that system are designed to work together.

A very important procedure to observe when tightening the bolts holding two flanges together is to evenly distribute the bolt pressure, so that no single region of the flange receives significantly more bolt pressure than any other region. In an ideal world, you would tighten all bolts to the same torque limit *simultaneously*. However, since this is impossible with just a single wrench, the best alternative is to tighten the bolts in alternating sequence, in stages of increasing torque. An illustrative torque sequence is shown in the following diagram (the numbers indicate the order in which the bolts should be tightened):



With one wrench, you would tighten each bolt to a preliminary torque in the sequence shown. Then, you would repeat the tightening sequence with additional torque for a more cycles until all bolts had been tightened to the recommended torque value. Note how the torque sequence alternates between four quadrants of the flange, ensuring the flanges are evenly compressed together as all bolts are gradually tightened. This technique of alternating quadrants around the circle is often referred to as *cross-torquing*.

Special wrenches called *torque wrenches* exist for the purpose of measuring applied torque during the tightening process. In critical, high-pressure applications, the actual *stretch* of each flange bolt is measured as a direct indication of bolting force. A special bolt sold under the brand name of *Rotabolt* contains its own built-in strain indicator, letting the mechanic know when the bolt has been sufficiently tightened regardless of the tool used to tighten it.


Another important procedure to observe when working with flanged pipe connections is to loosen the bolts on the *far* side of the flange before loosening the bolts on the side of the flange nearest you. This is strictly a precautionary measure against the spraying of process fluid toward your face or body in the event of stored pressure inside of a flanged pipe. By reaching over the pipe to first loosen flange bolts on the far side, if any pressure happens to be inside the pipe, it should leak there first, venting the pressure in a direction away from you.

A special provision of flanged pipe connections is the ability to install a blank metal plate called a *blind* over or between flange faces, thereby preventing flow. This is useful when a pipe must be blocked in a semi-permanent fashion, for example if that section of pipe has been decommissioned, or if the section of pipe must be sealed for reasons of safety during maintenance operations.

In order to install a blind, the flange joint must first be broken, then the flanges pried apart to provide the necessary space for the blind. After installing new gaskets along with the blind, the flanged bolts may then be re-installed and torqued to specification. A photograph of a stainless-steel blind (not installed on a pipe) appears here, two welded lifting tabs being clearly seen to facilitate handling this heavy piece of hardware:



In applications where “blinding” is frequent, a permanent form of blind called a *spectacle blind* may be installed to simplify the task. A spectacle blind is comprised of a regular blind plate attached to an equal-diameter ring by a short tab, the outline of which resembles a pair of spectacles:



Since the spectacle blind’s ring is exactly the same thickness as its blind plate, the piping system may be designed and built with the blind’s thickness in mind, the flange-to-flange gap remaining constant for the “open” and “blinded” states. This is especially helpful in very large piping systems, where the force required to separate formerly mated flange faces may be very large.

A spectacle blind may be seen in this next photograph, where the blind is installed in such a way that the yellow-painted “blind” half is exposed and the “open” half is sandwiched between the pipe flanges to allow flow through that pipe:




This next photograph shows a spectacle blind installed the other way, where the “open” half is exposed and the “blind” half is blocking any fluid from moving through the pipe:




8.1.2 Tapered thread pipe fittings

For smaller pipe sizes, *threaded fittings* are more commonly used to create connections between pipes and between pipes and equipment (including some instruments). A very common design of threaded pipe fitting is the *tapered* pipe thread design. The intent of a tapered thread is to allow the pipe and fitting to “wedge” together when engaged, creating a joint that is both mechanically rugged and leak-free.

When male and female tapered pie threads are first engaged, they form a loose junction:




After tightening, however, the tapered profile of the threads acts to wedge both male and female pieces tightly together as such:



Several different standards exist for tapered-thread pipe fittings. For each standard, the angle of the thread is fixed, as is the angle of taper. Thread *pitch* (the number of threads per unit length)

varies with the diameter of the pipe fitting².

In the United States, the most common tapered thread standard for general-purpose piping is the *NPT*, or *National Pipe Taper* design. NPT threads have an angle of 60° and a taper of $1^\circ 47'$ (1.7833°):




NPT pipe threads must have some form of *sealant* applied prior to assembly to ensure pressure-tight sealing between the threads. Teflon tape and various liquid pipe “dope” compounds work well for this purpose. Sealants are necessary with NPT threads for two reasons: to lubricate the male and female pieces (to guard against galling the metal surfaces), and also to fill the spiral gap formed between the root of the female thread and the crest of the male thread (and vice-versa).

NPTF (National Pipe Thread) pipe threads are engineered with the same thread angle and pitch as NPT threads, but carefully machined to avoid the spiral leak path inherent to NPT threads. This design – at least in theory – avoids the need to use sealant with NPTF threads to achieve a pressure-tight seal between male and female pieces, which is why NPTF threads are commonly referred to as *dryseal*. However, in practice it is still recommended that some form of sealant be used (or at the very least some form of thread *lubricant*) in order to achieve reliable sealing.

ANPT (Aeronautical National Pipe Tapered) is identical to NPT, except with a greater level of precision and quality for its intended use in aerospace and military applications.

²For example, 1/8 inch NPT pipe fittings have a thread pitch of 27 threads per inch. 1/4 inch and 3/8 inch NPT fittings are 18 threads per inch, 1/2 inch and 3/4 inch NPT fittings are 14 threads per inch, and 1 inch through 2 inch NPT fittings are 11.5 threads per inch.


Another tapered-thread standard is the *BSPT*, or *British Standard Pipe Tapered*. BSPT threads have a narrower thread angle than NPT threads (55° instead of 60°) but the same taper of $1^\circ 47'$ (1.7833°):



8.1.3 Parallel thread pipe fittings


An alternative to tapered threads in pipe joints is the use of parallel threads, similar to the threads of machine screws and bolts. Since parallel threads are incapable of forming a pressure-tight seal on their own, the sealing action of a parallel thread pipe fitting must be achieved some other way. This function is usually met with an O-ring or gasket.

In the United States, a common design of parallel-thread pipe fitting is the *SAE straight thread*, named after the *Society of Automotive Engineers*:



Sealing is accomplished as the O-ring is compressed against the shoulder of the female fitting. The threads serve only to provide force (not fluid sealing), much like the threads of a fastener.

Another parallel-thread pipe standard is the *BSPP*, or *British Standard Pipe Parallel*. Like the BSPT (tapered) standard, the thread angle of BSPP is 55° . Like the SAE parallel-thread standard, sealing is accomplished by means of an O-ring which compresses against the shoulder of the matching female fitting:




8.1.4 Sanitary pipe fittings


Food processing, pharmaceuticals manufacturing, and biological research processes are naturally sensitive to the presence of micro-organisms such as bacteria, fungi, and algae. It is important in these processes to ensure the absence of harmful micro-organisms, for reasons of both human health and quality control. For this reason, the process piping and vessels in these industries is designed first and foremost to be thoroughly cleaned without the need for disassembly. Regular cleaning and sterilization cycles are planned and executed between production schedules (batches) to ensure no colonies of harmful micro-organisms can grow.

A common *Clean-In-Place* (CIP) protocol consists of draining all process piping and vessels of process liquid, then flushing them with a sequence of rinse water, detergent solution, caustic solution, and sometimes an acid solution, followed by a final water rinse. For increased sanitization, a *Steam-In-Place* (SIP) cycle may be incorporated as well, sterilizing all process pipes and vessels with hot steam to ensure the destruction of any micro-organisms.

An important design feature of any sanitary process is the elimination of any “dead ends” (often called *dead legs* in the industry), crevices, or voids where fluid may collect and stagnate. This includes any instruments contacting the process fluids. It would be unsafe, for example, to connect something as simple as a bourdon-tube pressure gauge to a pipe carrying biologically sensitive fluid(s), since the interior volume of the bourdon tube will act as a stagnant refuge for colonies of micro-organisms to grow:



Instead, any pressure gauge must use an isolating diaphragm, where the process fluid pressure is transferred to the gauge mechanism through a sterile “fill fluid” that never contacts the process fluid:



With the isolating diaphragm in place, there are no stagnant places for process fluid to collect and avoid flushing by CIP or SIP cycles.

Standard pipe fittings are problematic in sanitary systems, as tiny voids between the mating threads of male and female pipe fittings may provide refuge for micro-organisms. To avoid this problem, special *sanitary fittings* are used instead. These fittings consist of a matched pair of flanges, held together by an external clamp. An array of sanitary fittings on an instrument test bench appear in the following photograph:



The next photograph shows the installation of a pressure transmitter on an ultra-pure water line using one of these sanitary fittings. The external clamp holding the two flanges together is clearly visible in this photograph:



Sanitary pipe fittings are not limited to instrument connections, either. Here are two photographs of process equipment (a ball valve on the left, and a pump on the right) connected to process pipes using sanitary fittings:



8.2 Tube and tube fittings

Tube, like pipe, is a hollow structure designed to provide an enclosed pathway for fluids to flow. In the case of tubing, it is usually manufactured from rolled or extruded metal (although plastic is a common tube material for many industrial applications). This section discusses some of the more common methods for joining tubes together (and joining tube ends to equipment such as pressure instruments).

One of the fundamental differences between tube and pipe is that tube is *never* threaded at the end to form a connection. Instead, a device called a *tube fitting* must be used to couple a section of tube to another tube, or to a section of pipe, or to a piece of equipment (such as an instrument). Unlike pipes which are thick-walled by nature, tubes are thin-walled structures. The wall thickness of a typical tube is simply too thin to support threads.

Tubes are generally favored over pipe for small-diameter applications. The ability for skilled workers to readily cut and bend tube with simple hand tools, as well as the ability to repeatedly break and re-make tube connections without compromising the integrity of the seals, makes tube the preferred choice for connecting instruments to process piping. When used as the connecting units between an instrument and a process pipe or vessel, the tube is commonly referred to as an *impulse tube* or *impulse line*³.

³Impulse lines are alternatively called *gauge lines* or *sensing lines*.

8.2.1 Compression tube fittings

By far the most common type of tube fitting for instrument impulse lines is the *compression-style* fitting, which uses a compressible *ferrule* to perform the task of sealing fluid pressure. The essential components of a compression tube fitting are the *body*, the *ferrule*, and the *nut*. The ferrule and body parts have matching conical profiles designed to tightly fit together, forming a pressure-tight metal-to-metal seal. Some compression fitting designs use a two-piece ferrule assembly, such as this tube fitting shown here⁴ (prior to full assembly):



Just prior to assembly, we see how the nut will cover the ferrule components and push them into the conical entrance of the fitting body:




After properly tightening the nut, the ferrule(s) will *compress* onto the outside circumference of the tube, slightly crimping the tube in the process and thereby locking the ferrules in place:



⁴This happens to be a Swagelok brass instrument tube fitting being installed on a 3/8 inch copper tube.

When initially assembling compression-style tube fittings, you should always precisely follow the manufacturer's instructions to ensure correct compression. For Swagelok-brand instrument tube fittings 1 inch in size and smaller, the general procedure to "swage" a new connector to a tube is to tighten the nut 1-1/4 turns past finger-tight. Insufficient turning of the nut will fail to properly compress the ferrule around the tube, and excessive turning will over-compress the ferrule, resulting in leakage. After this initial "swaging," the connector may be separated by loosening the nut until it no longer engages with the body, then the connection may be re-made by threading the nut back on the body until finger-tight and then gently tightening with a wrench until snug (no additional 1-1/4 turns!!!).

Swagelok provides special gauges which may be used to measure proper ferrule compression during the assembly process. The design of the gauge is such that its thickness will fit between the nut and fitting shoulder if the nut is insufficiently tightened, but will not fit if it is sufficiently tightened. Thus the gauge has the ability to reveal an under-tightened fitting, but not an over-tightened fitting. These gauges fit easily in the palm of one's hand:



Such gauges are referred to in the industry as *no-go gap gauges*, because their inability to fit between the nut and body shoulder of a tube fitting indicates a properly-tightened fitting. In other words, the gauge fit will be "no-go" if the tube fitting has been properly assembled.

Photographs showing one of these gauges testing a properly-tightened fitting (left) versus an under-tightened fitting (right) appear here:



Parker is another major manufacturer⁵ of instrument tube fittings, and their product line uses a single-piece ferrule instead of the two-piece design preferred by Swagelok. Like Swagelok fittings, Parker instrument fitting sized 1/4 inch to 1 inch require 1-1/4 turns past hand tight to properly compress the ferrule around the circumference of the tube. Parker also sells gauges which may be used to precisely determine when the proper amount of ferrule compression is achieved.

⁵So is Gyrolok, Hoke, and a host of others. It is not my intent to advertise for different manufacturers in this textbook, but merely to point out some of the more common brands an industrial instrument technician might encounter on the job.

What a gap gauge will not indicate is *over-tightening*. Over-tightening of a compression fitting is just as bad as under-tightening, as the fitting cannot form a robust seal once the ferrule and tube have been deformed. An example of an over-tightened Swagelok two-piece ferrule (left) on a plastic tube appears in the following photograph, next to a properly swaged ferrule (right):



Note the lack of a substantial gap between the two ferrule pieces in the over-tightened example. Note also the steeper cone taper of the over-tightened front ferrule, as a result of being pushed too deep inside the fitting body.

Regardless of the brand, compression-style instrument tube fittings are incredibly strong and versatile. Unlike pipe fittings, tube fittings may be disconnected and reconnected with ease. No special procedures are required to “re-make” a disassembled instrument fitting connection: merely tighten the nut “snug” to maintain adequate force holding the ferrule to the fitting body, but not so tight that the ferrule compresses further around the tube than it did during initial assembly.


A very graphic illustration of the strength of a typical instrument tube fitting is shown in the following photograph, where a short section of 3/8 inch stainless steel instrument tube was exposed to high liquid pressure until it ruptured. Neither compression fitting on either side of the tube leaked during the test, despite the liquid pressure reaching a peak of 23000 PSI before rupturing the tube⁶:



⁶It should be noted that the fitting nuts became seized onto the tube due to the tube's swelling. The tube fittings may not have leaked during the test, but their constituent components are now damaged and should never be placed into service again.


8.2.2 Common tube fitting types and names

Tube fittings designed to connect a tube to pipe threads are called *connectors*. Tube fittings designed to connect one tube to another are called *unions*:




If a tube union joins together different tube sizes rather than tubes of the same size, it is called a *reducing union*.

A variation on the theme of tube connectors and unions is the *bulkhead* fitting. Bulkhead fittings are designed to fit through holes drilled in panels or enclosures to provide a way for a fluid line to pass through the wall of the panel or enclosure. In essence, the only difference between a bulkhead fitting and a normal fitting is the additional length of the fitting "barrel" and a special nut used to lock the fitting into place in the hole. The following illustration shows three types of bulkhead fittings:




Tubing *elbows* are tube connectors with a bend. These are useful for making turns in tube runs without having to bend the tubing itself. Like standard connectors, they may terminate in male pipe thread, female pipe threads, or in another tube end:




These elbows shown in the above illustration are all 90° , but this is not the only angle available. 45° elbows are also common.


Tee fittings join three fluid lines together. Tees may have one pipe end and two tube ends (*branch tees* and *run tees*), or three tube ends (*union tees*). The only difference between a branch tee and a run tee is the orientation of the pipe end with regard to the two tube ends:



Of course, branch and run tee fittings also come in female pipe thread versions as well. A variation of the theme of union tees is the *cross*, joining four tubes together:




Special tube fittings are made to terminate tube connections, so they are sealed up instead of open. A piece designed to seal off the open end of a tube fitting is called a *plug*, while a piece designed to seal off the end of an open tube is called a *cap*:



8.2.3 Bending instrument tubing

Tube bending is something of an art, especially when done with stainless steel tubing. It is truly magnificent to see a professionally-crafted array of stainless steel instrument tubes, all bends perfectly made, all terminations square, all tubes parallel when laid side by side and perfectly perpendicular when crossing.

If possible, a goal in tube bending is to eliminate as many connections as possible. Connections invite leaks, and leaks are problematic. Long runs of instrument tubing made from standard 20 foot tube sections, however, require junctions be made somewhere, usually in the form of tube *unions*. When multiple tube unions must be placed in parallel tube runs, it is advisable to offset the unions so it is easier to get a wrench around the tube nuts to turn them. The philosophy here, *as always*, is to build the tubing system with future work in mind. A photograph of several tube junctions shows one way to do this:



If an instrument tube must connect between a stationary object and a vibrating object, a straight (square) run of tube is actually not desirable, since it will not have much flexibility to absorb the vibration. Instead, a *vibration loop* should be made in the tube, giving it the necessary elasticity to tolerate the vibrational stresses. An example of a vibration loop placed in the air supply tube going to a control valve appears in this photograph:



When bending such a loop, it is helpful to use the circumference of a large pipe as a mandrel to form the tube rather than attempt to form a loop purely by hand.

8.2.4 Special tubing tools

A variety of specialized tools exist to help tubing installers work with compression-style tube fittings. One of these special devices is an electronic power tool manufactured by American Power Tool expressly for use with instrument tube fittings:



The Aeroswage SX-1 has a microprocessor-controlled electric motor programmed to rotate a tube fitting's nut to a precise angular dimension, in order to properly swage the fitting. The tool comes complete with a holding jig to engage the body of the tube fitting, in order that all tightening torque is borne by the tool and not imposed on the person operating the tool:



Not only does this feature reduce the amount of stress placed on the tube fitter's hand and wrist, but it also enables the tool to be used in the demanding environment of zero gravity, for example aboard a space station. In such an environment, torque applied to the tool operator could be disastrous, as the human operator has no weight to stabilize herself.

This next pair of photos shows how the tool is able to support itself on a piece of stiff ($\frac{1}{2}$ inch stainless steel) tubing, and indeed may even be operated hands-free:



The amount of rotation is programmable, enabling the tool to be used with different kinds of fittings. For standard industrial Swagelok compression fitting sizes ($\frac{1}{4}$ inch, $\frac{3}{8}$ inch, and $\frac{1}{2}$ inch), the recommended swaging rotation of 1-1/4 turns may be entered into the tool as a tightening angle of 450 degrees:



Being a microprocessor-controlled device, the SX-1 has the ability to digitally record all actions. This is useful in high-reliability production environments (e.g. aerospace tube installation) where individual tube fitting data are archived for safety and quality control purposes. This data may be downloaded to a personal computer through a serial port connection on the side of the tool. Here you can see the tool's digital display showing the recorded action number, tightening angle, date, and time:




For large instrument compression fittings, hydraulic swaging tools are also available to provide the force necessary to properly compress the ferrule(s) onto the tube. Instrument tube manufacturers will provide specific recommendations for the installation of non-standard tube types, sizes, and materials, and also recommend particular swaging tools to use with their fittings.

8.3 Electrical signal and control wiring

There is much to be said for neatness of assembly in electrical signal wiring. Even though the electrons don't "care" how neatly the wires are laid in place, human beings who must maintain the system certainly do. Not only are neat installations easier to navigate and troubleshoot, but they tend to inspire a similar standard of neatness when alterations are made⁷.


The following photographs illustrate excellent wiring practice. Study them carefully, and strive to emulate the same level of professionalism in your own work.

Here we see 120 volt AC power distribution wiring. Note how the hoop-shaped "jumper" wires are all cut to (nearly) the same length, and how each of the wire labels is oriented such that the printing is easy to read:



⁷No one wants to become known as the person who "messed up" someone else's neat wiring job!

This next photograph shows a great way to terminate multi-conductor signal cable to terminal blocks. Each of the pairs was twisted together using a hand drill set to very slow speed. Note how the end of the cable is wrapped in a short section of heat-shrink tubing for a neat appearance:




Beyond esthetic preferences for instrument signal wiring are several practices based on sound electrical theory. The following subsections describe and explain these wiring practices.


8.3.1 Connections and wire terminations

Many different techniques exist for connecting electrical conductors together: twisting, soldering, crimping (using compression connectors), and clamping (either by the tension of a spring or under the compression of a screw) are popular examples. Most industrial field wiring connections utilize a combination of compression-style crimp “lugs” (often referred to as *ferrules* or *compression terminals*) and screw clamps to attach wires to instruments and to other wires.

The following photograph shows a typical *terminal strip* or *terminal block* array whereby twisted-pair signal cables connect to other twisted-pair signal cables. Metal bars inside each plastic terminal section form connections horizontally, so that wires fastened to the left side are connected to wires fastened to the right side:




If you look closely at this photograph, you can see the bases of crimp-style ferrules at the ends of the wires, just where they insert into the terminal block modules. These terminal blocks use screws to apply force which holds the wires in close electrical contact with a metal bar inside each block, but metal ferrules have been crimped on the end of each wire to provide a more rugged tip for the terminal block screw to hold to. A close-up view shows what one of these ferrules looks like on the end of a wire:




Also evident in this photograph is the dual-level connection points on the left-hand side of each terminal block. Two pairs of twisted signal conductors connect on the left-hand side of each terminal block pair, where only one twisted pair of wires connects on the right-hand side. This also explains why each terminal block section has two screw holes on the left but only one screw hole on the right.

A close-up photograph of a single terminal block section shows how the screw-clamp system works. Into the right-hand side of this block a single wire (tipped with a straight compression ferrule) is clamped securely. No wire is inserted into the left-hand side:



If another wire were secured by the screw clamp on the left-hand side of this terminal block, it would be made electrically common with the wire on the right-hand side by virtue of the metal bar joining both sides.


Some terminal blocks are *screwless*, using a spring clip to make firm mechanical and electrical contact with the wire's end:




In order to extract or insert a wire end from or to a “screwless” terminal block, you must insert a narrow screwdriver into a hole in the block near the insertion point, then pivot the screwdriver (like a lever) to exert force on the spring clip. Screwless terminal blocks are generally faster to terminate and un-terminate than screw type terminal blocks, and the pushing action of the release tool is gentler on the body⁸ than the twisting action required to loosen and tighten screws.

⁸An occupational hazard for technicians performing work on screw terminations is *carpal tunnel syndrome*, where repetitive wrist motion (such as the motions required to loosen and tighten screw terminals) damages portions of the wrist where tendons pass.

Many different styles of modular terminal blocks are manufactured to suit different wiring needs. Some terminal block modules, for example, have multiple “levels” instead of just one. The following photograph shows a two-level terminal block with screwless wire clamps:




The next photograph shows a three-level terminal block with screw type clamps:



Some multi-level terminal blocks provide the option of *internal jumpers* to connect two or more levels together so they will be electrically common instead of electrically isolated. This use of a multi-level terminal block is preferable to the practice of inserting multiple wires into the same terminal, when wires need to be made common to each other.

Other modular terminal blocks include such features as LED indicator lamps, switches, fuses, and even resettable circuit breakers in their narrow width, allowing the placement of actual circuit components near connection points. The following photograph shows a swing-open fused terminal block module, in the open position:



Modular terminal blocks are useful for making connections with both solid-core and stranded metal wires. The clamping force applied to the wire's tip by the screw mechanism inside one of these blocks is direct, with no sliding or other motions involved. Some terminal blocks, however, are less sophisticated in design. This next photograph shows a pair of "isothermal" terminals designed to connect thermocouple wires together. Here you can see how the bare tip of the screw applies pressure to the wire inserted into the block:



The rotary force applied to each wire's tip by these screws necessitates the use of solid wire. Stranded wire would become frayed by this combination of forces.

Many field instruments, however, do not possess “block” style connection points at all. Instead, they are equipped with pan-head machine screws designed to compress the wire tips directly between the heads of the screws and a metal plate below.

Solid wires may be adequately joined to such a screw-head connection point by partially wrapping the bare wire end around the screw’s circumference and tightening the head on top of the wire, as is the case with the two short wire stubs terminated on this instrument:




The problem with directly compressing a wire tip beneath the head of a screw is that the tip is subjected to both compressive and shear forces. As a result, the wire’s tip tends to become mangled with repeated connections. Furthermore, tension on the wire will tend to turn the screw, potentially loosening it over time.

This termination technique is wholly unsuitable for stranded wire⁹, because the shearing forces caused by the screw head’s rotation tends to “fray” the individual strands. The best way to attach a stranded wire tip directly to a screw-style connection point is to first crimp a compression-style *terminal* to the wire. The flat metal “lug” (ferrule) portion of the terminal is then inserted underneath the screw head, where it can easily tolerate the shearing and compressive forces exerted by the head.

⁹An exception is when the screw is equipped with a square washer underneath the head, designed to compress the end of a stranded wire with no shear forces. Many industrial instruments have termination points like this, for the express purpose of convenient termination to either solid or stranded wire ends.

This next photograph shows five such stranded-copper wires connected to screw-style connection points on a field instrument using compression-style terminals:




Compression-style terminals come in two basic varieties: *fork* and *ring*. An illustration of each type is shown here:



Fork terminals are easier to install and remove, since they merely require loosening of the connector screw rather than removal of the screw. Ring terminals are more secure, since they cannot “fall off” the connection point if the screw ever accidentally loosens.


Just as direct termination underneath a screw head is wholly unsuitable for stranded wires, compression-style terminals are wholly unsuitable for solid wire. Although the initial crimp may feel secure, compression terminals lose their tension rapidly on solid wire, especially when there is any motion or vibration stressing the connection. Compression wire terminals should only be crimped to stranded wire!

Properly installing a compression-type terminal on a wire end requires the use of a special *crimping* tool. The next photograph shows one of these tools in use:




Note the different places on the crimping tool, labeled for different wire sizes (gauges). One location is used for 16 gauge to 10 gauge wire, while the location being used in the photograph is for wire gauges 22 through 18 (the wire inside of the crimped terminal happens to be 18 gauge).

This particular version of a “crimping” tool performs most of the compression on the underside of the terminal barrel, leaving the top portion undisturbed. The final crimped terminal looks like this when viewed from the top:





8.3.2 DIN rail

An industry-standard structure for attaching terminal blocks and small electrical components to flat metal panels is something called a *DIN rail*. This is a narrow channel of metal – made of bent sheet steel or extruded aluminum – with edges designed for plastic components to “clip” on. The following photograph shows terminal blocks, relay sockets, fuses, and more terminal blocks mounted to a horizontal length of DIN rail in a control system enclosure:



Two photographs of a terminal block cluster clipped onto a length of DIN rail – one from above and one from below – shows how specially-formed arms on each terminal block module fit the edges of the DIN rail for a secure attachment:



The DIN rail itself mounts on to any flat surface by means of screws inserted through the slots in its base. In most cases, the flat surface in question is the metal subpanel of an electrical enclosure to which all electrical components in that enclosure are attached.


An obvious advantage of using DIN rail to secure electrical components versus individually attaching those components to a subpanel with their own sets of screws is convenience: much less labor is required to mount and unmount a DIN rail-attached component than a component attached with its own set of dedicated screws. This convenience significantly eases the task of altering a panel's configuration. With so many different devices manufactured for DIN rail mounting, it is easy to

upgrade or alter a panel layout simply by unclipping components, sliding them to new locations on the rail, or replacing them with other types or styles of components.


This next photograph shows some of the diversity available in DIN rail mount components. From left to right we see four relays, a power supply, and three HART protocol converters, all clipped to the same extruded aluminum DIN rail:





As previously mentioned, DIN rail is available in both stamped sheet-steel and extruded aluminum forms. A comparison of the two materials is shown here, with sheet steel on the left and aluminum on the right:



The form of DIN rail shown in all photographs so far is known as “top hat” DIN rail. A variation in DIN rail design is the so-called “G” rail, with a markedly different shape:



Fortunately, many modular terminal blocks are formed with the ability to clip to either style of DIN rail, such as these two specialty blocks, the left-hand example being a terminal block with a built-in disconnect switch, and the right-hand example being a “grounding” terminal block whose termination points are electrically common to the DIN rail itself:




If you examine the bottom structure of each block, you will see formations designed to clip either to the edges of a standard (“top hat”) DIN rail or to a “G” shaped DIN rail.

Smaller DIN rail standards also exist, although they are far less common than the standard 35mm size:



A nice feature of many DIN rail type terminal blocks is the ability to attach pre-printed terminal numbers. This makes documentation of wiring much easier, with each terminal connection having its own unique identification number:



8.3.3 Cable routing

In the interest of safety and longevity, one cannot simply route electrical power and signal cables randomly between different locations. Electrical cables must be properly supported to relieve mechanical stresses on the conductors, and protected from harsh conditions such as abrasion which might degrade the insulation.

A traditional and rugged technique for cable routing is *conduit*, either metal or plastic (PVC). Conduit resembles piping used to convey fluids, except that it is much thinner-walled than fluid pipe and is not rated to withstand internal pressure as pipe is. In fact, threaded conduit uses the same thread pitch and diameter standards as NPT (National Pipe Taper) fluid pipe connections.

Metal conduit naturally forms a continuously-grounded enclosure for conductors which not only provides a measure of protection against electrical shock (all enclosures and devices attached to the conduit become safely grounded through the conduit) but also shields against electrostatic interference. This is especially important for power wiring to and from devices such as rectifiers and variable-frequency motor drive (VFD) units, which have a tendency to broadcast large amounts of electromagnetic noise.

Plastic conduit, of course, provides no electrical grounding or shielding because plastic is a non-conductor of electricity. However, it is superior to metal conduit with regard to chemical corrosion resistance, which is why it is used to route wires in areas containing water, acids, caustics, and other wet chemicals.

Thin-wall conduit is made with metal so thin that threads cannot be cut into it. Instead, special connectors are used to join “sticks” of thin-wall conduit together, and to join thin-wall conduit to electrical enclosures. Several runs of thin-wall conduit appear in this next photograph. Two of those conduit runs have been severed following a wiring change, exposing the conductors inside:



Installing cable into an electrical conduit is a task referred to as *cable pulling*, and it is something of an art. Cable “pulls” may be especially challenging if the conduit run contains many bends, and/or is close to capacity in terms of the number and size of conductors it already holds. A good practice is to always leave a length of nylon *pull string* inside each length of conduit, ready to use for pulling a new wire or cable through. When performing a wire “pull,” a new length of nylon pull string is pulled into the conduit along with the new wires, to replace the old pull string as it is pulled out of the conduit. Special lubricating “grease” formulated for electrical wiring may be applied to conductors pulled into a conduit, to reduce friction between those new conductors and the conductors already inside the conduit.

When connecting electrical conduit to end-point devices, it is common to use flexible *liquid-tight conduit* as a connector between the rigid metal (or plastic) conduit and the final device. This provides some stress relief to the conduit in the event the device moves or vibrates, and also allows more freedom in positioning the device with respect to the conduit. Here, we see a motor-operated control valve with two runs of liquid-tight conduit routing conductors to it:



Liquid-tight conduit comes in two general varieties: metallic and non-metallic. The metallic kind contains a spiraled metal sheath just underneath the plastic outer coating to provide a continuously-

grounded shield much the same way that rigid metal conduit does. Non-metallic liquid-tight conduit is nothing more than plastic hose, providing physical protection against liquid exposure and abrasion, but no electrical grounding or shielding ability.

Another technique for cable routing is the use of *cable tray*. Trays may be made of solid steel wire for light-duty applications such as instrument signal cabling or computer network cabling, or they may be made of steel or aluminum channel for heavy-duty applications such as electrical power wiring. Unlike conduit, cable trays are open, leaving the cables exposed to the environment. This often necessitates special cable insulation rated for exposure to ultraviolet light, moisture, and other environmental wear factors. A decided advantage of cable trays is ease of cable installation, especially when compared to electrical conduit.

While cable tray does provide a continuously-grounded surface for electrical safety the same as metal conduit, cable tray does *not* naturally provide shielding for the conductors because it does not completely enclose the conductors the way metal conduit does.

An example of light-duty cable tray appears here, used to support Ethernet cabling near the ceiling of a room at a college campus. The cable tray is made of solid steel wire, bent to form a “basket” to support dozens of yellow Ethernet cables:



Heavy-duty cable tray appears throughout this next photograph, supporting large-gauge power conductors for electric generators at a gas turbine power plant. Here, the cable tray has the appearance of an aluminum ladder, with extruded metal rails and rungs providing physical support for the cables:



Similar cables trays appear in the next photograph, supporting feeder cables from a stationary transformer and switchgear cabinets:



A special form of wiring often seen in industrial facilities for power distribution is *busway*, also known as *bus duct*. These are rectangular sheet-metal tubes containing pre-fabricated copper busbars for the conduction of three-phase AC power. Special junction boxes, “tees,” and tap boxes allow busways to extend and branch to other busways and/or standard conductor wiring.

Busways are used in indoor applications, often in motor control center (MCC) and power distribution center rooms to route electrical power to and from large disconnect switches, fuses, and circuit breakers. In this photograph, we see busway used to distribute power along the ceiling of an MCC room, alongside regular rigid conduit:



As useful and neat in appearance as busways are, they are definitely limited in purpose. Busways are only used for electrical power distribution; not for instrumentation, control, or signaling purposes.

Two materials useful for neatly routing power, signal, and instrumentation conductors inside an enclosure are *wire duct* and *wire loom*. Wire duct is a plastic channel with slotted sides, designed to be attached to the subpanel of an enclosure along with all electrical devices inside that enclosure. Wires pass from the devices to the duct through the slots (gaps) in the sides of the duct, and are enclosed by a removable plastic cover that snaps onto the top of the duct. A common brand name of wire duct in the industry is *Panduit* and so you will often hear people refer to wire duct as “Panduit” whether or not that particular brand is the one being used¹⁰. Wire loom is a loose spiral

¹⁰This is similar to people referring to adhesive bandages as “Band-Aids” or tongue-and-groove joint pliers as

tube made of plastic, used to hold a group of individual wires together into a neat bundle. Wire loom is frequently used when a group of conductors must periodically flex, as is the case of a wire bundle joining devices inside a panel to other devices mounted on the hinging door of that panel.

A photograph showing both wire duct and wire loom inside an instrumentation panel appears here. The wire duct is the grey-colored rectangular plastic channel set vertically and horizontally inside the panel, while the loom is a grey-colored plastic spiral surrounding the bundle of wires near the door hinge:




“Channelocks,” because those particular brands have become popular enough to represent an entire class.

8.3.4 Signal coupling and cable separation

If sets of wires lie too close to one another, electrical signals may “couple” from one wire (or set of wires) to the other(s). This can be especially detrimental to signal integrity when the coupling occurs between AC power conductors and low-level instrument signal wiring such as thermocouple or pH sensor cables.


Two mechanisms of electrical “coupling” exist: *capacitive* and *inductive*. Capacitance is a property intrinsic to any pair of conductors separated by a dielectric (an insulating substance), whereby energy is stored in the electric field formed by voltage between the wires. The natural capacitance existing between mutually insulated wires forms a “bridge” for AC signals to cross between those wires, the strength of that “bridge” inversely proportional to the capacitive reactance ($X_C = \frac{1}{2\pi fC}$). Inductance is a property intrinsic to any conductor, whereby energy is stored in the magnetic field formed by current through the wire. Mutual inductance existing between parallel wires forms another “bridge” whereby an AC current through one wire is able to induce an AC voltage along the length of another wire.

Capacitive coupling between an AC power conductor and a DC sensor signal conductor is shown in the following diagram:



If the voltage-generating sensor happens to be a thermocouple and the receiving instrument a temperature indicator, the result of this capacitive coupling will be a “noisy” temperature signal interpreted by the instrument. This noise will be proportional to both the voltage and the frequency of the AC power.

Inductive coupling between an AC power conductor and a DC sensor signal conductor is shown in the following diagram:




Whereas the amount of noise induced into a low-level signal via capacitive coupling was a function of *voltage* and frequency, the amount of noise induced into a signal via inductive coupling is a function of *current* and frequency¹¹.

A good way to minimize signal coupling is to simply separate conductors carrying incompatible signals. This is why electrical power conductors and instrument signal cables are almost never found in the same conduit or in the same ductwork together. Separation decreases capacitance between the conductors (recall that $C = \frac{A\epsilon}{d}$ where d is the distance between the conductive surfaces). Separation also decreases the coupling coefficient between inductors, which in turn decreases mutual inductance (recall that $M = k\sqrt{L_1L_2}$ where k is the coupling coefficient and M is the mutual inductance between two inductances L_1 and L_2). In control panel wiring, it is customary to route AC power wires in such a way that they do not lay parallel to low-level signal wires, so that both forms of coupling may be reduced.

¹¹The principle at work here is the strength of the *field* generated by the noise-broadcasting conductor: electric field strength (involved with capacitive coupling) is directly proportional to voltage, while magnetic field strength (involved with inductive coupling) is directly proportional to current.


If conductors carrying incompatible signals *must* cross paths, it is advisable to orient the conductors perpendicular to each other rather than parallel, like this:



Perpendicular conductor orientation reduces both inter-conductor capacitance *and* mutual inductance by two mechanisms. Capacitance between conductors is reduced by means of minimizing overlapping area (A) resulting from the perpendicular crossing. Mutual inductance is reduced by decreasing the coupling coefficient (k) to nearly zero since the magnetic field generated perpendicular to the current-carrying wire will be *parallel* and not perpendicular to the “receiving” wire. Since the vector for induced voltage is perpendicular to the magnetic field (i.e. parallel with the current vector in the “primary” wire) there will be no voltage induced along the length of the “receiving” wire.

The problem of power-to-signal line coupling is most severe when the signal in question is *analog* rather than *digital*. In analog signaling, even the smallest amount of coupled “noise” corrupts the signal. A digital signal, by comparison, will become corrupted only if the coupled noise is so great that it pushes the signal level above or below a detection threshold it should not cross. This disparity is best described through illustration.

Two signals are shown here, coupled with equal amounts of noise voltage:




The peak-to-peak amplitude of the noise on the analog signal is almost 20% of the entire signal range (the distance between the lower- and upper-range values), representing a substantial degradation of signal integrity. Analog signals have infinite resolution, which means *any* change in signal amplitude has meaning. Therefore, any noise whatsoever introduced into an analog signal will be interpreted as variations in the quantity that signal is supposed to represent.

That same amount of noise imposed on a digital signal, however, causes no degradation of the signal except for one point in time where the signal attempts to reach a “low” state but fails to cross the threshold due to the noise. Other than that one incident represented in the pulse waveform, the rest of the signal is completely unaffected by the noise, because digital signals only have meaning above the “high” state threshold and below the “low” state threshold. Changes in signal voltage level caused by induced noise will not affect the meaning of digital data unless and until the amplitude of that noise becomes severe enough to prevent the signal’s crossing through a threshold (when it should cross), or causes the signal to cross a threshold (when it should not).


From what we have seen here, digital signals are far more tolerant of induced noise than analog signals, all other factors being equal. If ever you find yourself in a position where you must route a signal wire near AC power conductors, and you happen to have the choice whether it will be an analog signal (e.g. 4-20 mA, 0-10 V) or a digital signal (e.g. EIA/TIA-485, Ethernet), your best option is to choose the digital signal to coexist alongside the AC power wires.

8.3.5 Electric field (capacitive) de-coupling

The fundamental principle invoked in *shielding* signal conductor(s) from external electric fields is that no substantial electric field can exist within a solid conductor. Electric fields exist due to imbalances of electric charge. If such an imbalance of charge ever were to exist within a conductor, charge carriers (typically electrons) in that conductor would quickly move to equalize the imbalance, thus eliminating the electric field. Another way of saying this is to state that electric fields only exist between points of different potential, and therefore cannot exist between equipotential points. Thus, electric flux lines may be found only in the dielectric (insulating media) between conductors, not within a solid conductor:




This also means electric flux lines cannot span the diameter of a hollow conductor:




The electrical conductivity of the hollow sphere's wall ensures that all points on the circumference of the sphere are equipotential to each other. This in turn prohibits the formation of any electric flux lines within the interior air space of the hollow sphere. Thus, all points within the hollow sphere are *shielded* from any electric fields originating outside of the sphere.


The only way to allow an external electric field to penetrate a hollow conductor from the outside is if that conductive shell is left “floating” with respect to another conductor placed within the shell. In this case the lines of electric flux do not exist between different points on the conductive sphere, but rather between the shell of the sphere and the conductor at the center of the sphere because those are the points between which a potential difference (voltage) exists. To illustrate:



However, if we make the hollow shell electrically common to the negative side of the high-voltage source, the flux lines inside the sphere vanish, since there is no potential difference between the internal conductor and the conductive shell:



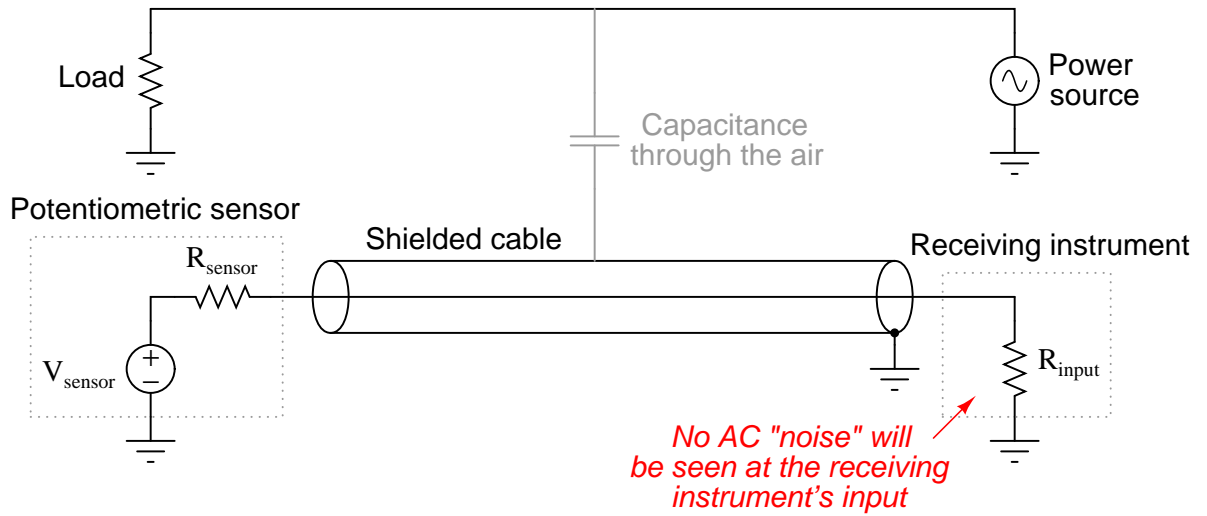
If the conductor within the hollow sphere is elevated to a potential different from that of the high-voltage source's negative terminal, electric flux lines will once again exist inside the sphere, but they will reflect this second potential and not the potential of the original high-voltage source. In other words, an electric field will exist inside the hollow sphere, but it will be completely isolated from the electric field outside the sphere. Once again, the conductor inside is *shielded* from external electrostatic interference:



If conductors located inside the hollow shell are thus shielded from external electric fields, it means there cannot exist any capacitance between external conductors and internal (shielded) conductors. If there is no capacitance between conductors, there will never be capacitive coupling of signals between those conductors, which is what we want for industrial signal cables to protect those signals from external interference¹².

¹²Incidentally, cable shielding likewise guards against strong electric fields *within* the cable from capacitively coupling with conductors outside the cable. This means we may elect to shield “noisy” power cables instead of (or in addition to) shielding low-level signal cables. Either way, good shielding will prevent capacitive coupling between conductors on either side of a shield.


All this discussion of hollow metal spheres is just an introduction to a discussion of *shielded cable*, where electrical cables are constructed with a conductive metal foil wrapping or conductive metal braid surrounding the interior conductors. Thus, the foil or braid creates a conductive *tube* which may be connected to ground potential (the “common” point between external and internal voltage sources) to prevent capacitive coupling between any external voltage sources and the conductors within the cable:



The following photograph shows a set of signal cables with braided shield conductors all connected to a common copper “ground bus.” This particular application happens to be in the control panel of a 500 kV circuit breaker, located at a large electrical power substation where strong electric fields abound:




This next photograph shows a four-conductor USB cable stripped at one end, revealing a metal-foil shield as well as silver-colored wire strands in direct contact with the foil, all wrapped around the four colored power and signal conductors:



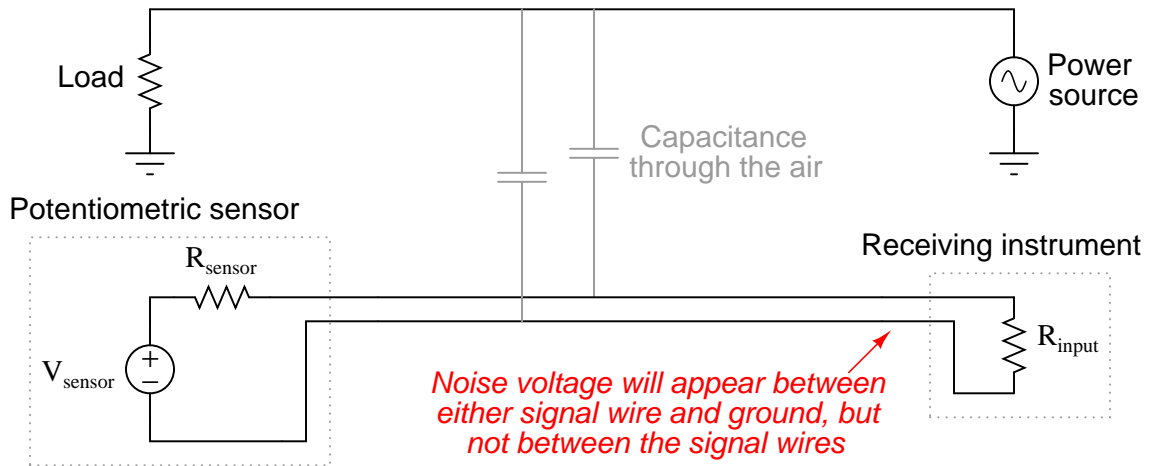
At the terminating end we typically twist the loose shield conductor strands together to form a wire which is then attached to a ground point to fix the cable’s shield at Earth potential.

It is very important to ground *only one end* of a cable's shield, or else you will create the possibility for a *ground loop*: a path for current to flow through the cable's shield resulting from differences in Earth potential at the cable ends. Not only can ground loops induce noise in a cable's conductor(s), but in severe cases it can even overheat the cable and thus present a fire hazard:

A ground loop: something to definitely avoid!



An important characteristic of capacitively-coupled noise voltage is that it is *common-mode* in nature: the noise appears equally on every conductor within a cable because those conductors lie so close to each other (i.e. because the amount of capacitance existing between each conductor and the noise source is the same). One way we may exploit this characteristic in order to help escape the unwanted effects of capacitive coupling is to use *differential signaling*. Instead of referencing our signal voltage to ground, we let the signal voltage “float.” The following schematic diagram illustrates how this works:




The lack of a ground connection in the DC signal circuit prevents capacitive coupling with the AC voltage from corrupting the measurement signal “seen” by the instrument. Noise voltage *will* still appear between either signal wire and ground as a common-mode voltage, but noise voltage *will not* appear *between* the two signal wires where our signal of interest exists. In other words, we

side-step the problem of common-mode noise voltage by making common-mode voltage irrelevant to the sensor and to the signal receiver.

Some industrial data communications standards such as EIA/TIA-485 (RS-485) use this technique to minimize the corrupting effects of electrical noise. To see a practical example of how this works in a data communications circuit, refer to the illustration in section 15.6.2 beginning on page 1055 of this book.

8.3.6 Magnetic field (inductive) de-coupling


Magnetic fields, unlike electric fields, are exceedingly difficult to completely shield. Magnetic flux lines do not terminate, but rather *loop*. Thus, one cannot “stop” a magnetic field, only re-direct its path. A common method for magnetically shielding a sensitive instrument is to encapsulate it in an enclosure made of some material having an extremely high magnetic permeability (μ): a shell offering much easier passage of magnetic flux lines than air. A material often used for this application is *mu-metal*, or μ -*metal*, so named for its excellent magnetic permeability:



This sort of shielding is impractical for protecting signal cables from inductive coupling, as mu-metal is rather expensive and must be layered relatively thick in order to provide a sufficiently low-reluctance path to shunt most of the external magnetic flux lines.

The most practical method of granting magnetic field immunity to a signal cable follows the differential signaling method discussed in the electric field de-coupling section, with a twist (literally). If we *twist* a pair of wires rather than allow them to lie along parallel straight lines, the effects of electromagnetic induction are vastly minimized.

The reason this works is best illustrated by drawing a differential signal circuit with two thick wires, drawn first with no twist at all. Suppose the magnetic field shown here (with three flux lines entering the wire loop) happens to be *increasing* in strength at the moment in time captured by the illustration:




According to Lenz's Law, a current will be induced in the wire loop in such a polarity as to oppose the increase in external field strength. In other words, the induced current tries to "fight" the imposed field to maintain zero net change. According to the right-hand rule of electromagnetism (tracing current in conventional flow notation), the induced current must travel in a counter-clockwise direction as viewed from above the wire loop in order to generate a magnetic field opposing the rise of the external magnetic field. This induced current works against the DC current produced by the sensor, detracting from the signal received at the instrument.

When the external magnetic field strength diminishes, then builds in the opposite direction, the induced current will reverse. Thus, as the AC magnetic field oscillates, the induced current will also oscillate in the circuit, causing AC "noise" voltage to appear at the measuring instrument. This is precisely the effect we wish to mitigate.

Immediately we see a remarkable difference between noise voltage induced by a magnetic field versus noise voltage induced by an electric field: whereas capacitively-coupled noise was always common-mode, here we see inductively-coupled noise as *differential*¹³.

¹³This is not to say magnetic fields cannot induce common-mode noise voltage: on the contrary, magnetic fields are

If we twist the wires so as to create a series of loops instead of one large loop, we will see that the inductive effects of the external magnetic field tend to cancel:



Not all the lines of flux go through the same loop. Each loop represents a reversal of direction for current in the instrument signal circuit, and so the direction of magnetically-induced current in one loop directly opposes the direction of magnetically-induced current in the next. So long as the loops are sufficient in number and spaced close together, the net effect will be complete and total opposition between all induced currents, with the result of no net induced current and therefore no AC “noise” voltage appearing at the instrument.

In order to enjoy the benefits of magnetic *and* electric field rejection, instrument cables are generally manufactured as *twisted, shielded pairs*. The twists guard against magnetic (inductive) interference, while the grounded shield guards against electric (capacitive) interference. If multiple wire pairs are twisted within the same cable, the twist rates of each pair may be made different so as to avoid magnetic coupling from pair to pair¹⁴.

capable of inducing voltage in any electrically-conductive loop. For this reason, both differential and ground-referenced signals are susceptible to interference by magnetic fields.

¹⁴An example of this is the UTP (Unshielded, Twisted Pair) cabling used for Ethernet digital networks, where four pairs of wires having different twist rates are enclosed within the same cable sheath.

8.3.7 High-frequency signal cables

Electronic signals used in traditional instrumentation circuits are either DC or low-frequency AC in nature. Measurement and control values are represented in *analog* form by these signals, usually by the magnitude of the electronic signal (how many volts, how many milliamps, etc.). Modern electronic instruments, however, often communicate process and control data in *digital* rather than analog form. This digital data takes the form of high-frequency voltage and/or current pulses along the instrument conductors. The most capable *fieldbus* instruments do away with analog signaling entirely, communicating all data in digital form at relatively high speeds.

If the time period of a voltage or current pulse is less than the time required for the signal to travel down the length of the cable (at nearly the speed of light!), very interesting effects may occur. When a pulse propagates down a two-wire cable and reaches the end of that cable, the energy contained by that pulse must be absorbed by the receiving circuit or else be *reflected* back down the cable. To be honest, this happens in all circuits no matter how long or brief the pulses may be, but the effects of a “reflected” pulse only become apparent when the pulse time is short compared to the signal propagation time. In such short-pulse applications, it is customary to refer to the cable as a *transmission line*, and to regard it as a circuit component with its own characteristics (namely, a continuous impedance as “seen” by the traveling pulse). For more detail on this subject, refer to section 5.10 beginning on page 475.

This problem has a familiar analogy: an “echo” in a room. If you step into a large room with hard wall, floor, and ceiling surfaces, you will immediately notice echoes resulting from any sound you make. Holding a conversation in such a room can be quite difficult, as the echoed sounds superimpose upon the most recently-spoken sounds, making it difficult to discern what is being said. The larger the room, the longer the echo delay, and the greater the conversational confusion.

Echoes happen in small rooms, too, but they are generally too short to be of any concern. If the reflected sound(s) return quickly enough after being spoken, the time delay between the spoken (incident) sound and the echo (reflected) sound will be too short to notice, and conversation will proceed unhindered.

We may address the “echo” problem in two entirely different ways. One way is to eliminate the echoes entirely by adding sound-deadening coverings (carpet, acoustic ceiling tiles) and/or objects (sofas, chairs, pillows) to the room. Another way to address the problem of echoes interrupting a conversation is to *slow down the rate of speech*. If the words are spoken slowly enough, the time delay of the echoes will be relatively short compared to the period of each spoken sound, and conversation may proceed without interference¹⁵ (albeit at a reduced speed).

Both the problem of and the solutions for reflected signals in electrical cables follow the same patterns as the problem of and solutions for sonic echoes in a hard-surfaced room. If an electronic circuit receiving pulses sent along a cable receives both the incident pulse and an echo (reflected pulse) with a significant time delay separating those two pulses, the digital “conversation” will be impeded in the same manner that a verbal conversation between two or more people is impeded by echoes in a room. We may address this problem either by eliminating the reflected pulses entirely (by ensuring all the pulse energy is absorbed by an appropriate load placed at the cable’s end) or by slowing down the data transfer rate (i.e. longer pulses, lower frequencies) so that the reflected and incident pulse signals virtually overlap one another at the receiver.

¹⁵A student of mine once noted that he has been doing this out of habit whenever he has a conversation with anyone in a racquetball court. All the hard surfaces (floor, walls) in a racquetball court create severe echoes, forcing players to speak slower in order to avoid confusion from the echoes.

High-speed “fieldbus” instrument networks apply the former solution (eliminate reflections) while the legacy HART instrument signal standard apply the latter (slow data rate). Reflections are eliminated in high-speed data networks by ensuring the two furthest cable ends are both “terminated” by a resistance value of the proper size (matching the characteristic impedance of the cable). The designers of the HART analog-digital hybrid standard chose to use slow data rates instead, so their instruments would function adequately on legacy signal cables where the characteristic impedance is not standardized.

The potential for reflected pulses in high-speed fieldbus cabling is a cause for concern among instrument technicians, because it represents a new phenomenon capable of creating faults in an instrument system. No longer is it sufficient to have tight connections, clean wire ends, good insulation, and proper shielding for a signal cable to faithfully convey a 4-20 mA DC instrument signal from one device to another. Now the technician must ensure proper termination and the absence of any discontinuities¹⁶ (sharp bends or crimps) along the cable’s entire length, in addition to all the traditional criteria, in order to faithfully convey a digital fieldbus signal from one device to another.

Signal reflection problems may be investigated using a diagnostic instrument known as a *time-domain reflectometer*, or *TDR*. These devices are a combination of pulse generator and digital-storage oscilloscope, generating brief electrical pulses and analyzing the returned (echoed) signals at one end of a cable. If a TDR is used to record the pulse “signature” of a newly-installed cable, that data may be compared to future TDR measurements on the same cable to detect cable degradation or wiring changes.

8.4 Fiber optics

*Light*¹⁷ has long been used as a long-range signaling medium. While communication by light through open air is still possible using modern technology, it is far more practical in most cases to channel the light signals through a special strand of optically transparent material called an *optical fiber*. When packaged in a protective sheath, it is known as a fiber optic cable.

The transmission of light through a “light pipe” was demonstrated as early as 1842 by Daniel Colladon and Jacques Babinet in Paris, using a running stream of water to guide a beam of light. Many modern houses in the United States are equipped with light-pipes¹⁸ directing natural sunlight into rooms for illumination, without the use of “skylight” ceiling windows. Modern fiber optic cables apply similar optical principles to very small-diameter fibers of transparent material (usually ultra-pure glass), able to convey optical energy and optically-encoded information.

¹⁶The characteristic, or “surge,” impedance of a cable is a function of its conductor geometry (wire diameter and spacing) and dielectric value of the insulation between the conductors. Any time a signal reaches an abrupt change in impedance, some (or all) of its energy is reflected in the reverse direction. This is why reflections happen at the unterminated end of a cable: an “open” is an infinite impedance, which is a huge shift from the finite impedance “seen” by the signal as it travels along the cable. This also means any sudden change in cable geometry such as a crimp, nick, twist, or sharp bend is capable of reflecting part of the signal. Thus, high-speed digital data cables must be installed more carefully than low-frequency or DC analog signal cables.


¹⁷Smoke signals are an ancient form of light-based communication!

¹⁸These are thin plastic or sheet metal tubes with mirrored internal surfaces, extending from a collector dome (made of glass or plastic) outside the dwelling to a diffusion lens inside the dwelling.

8.4.1 Fiber optic data communication

Simply put, an optical fiber is a “pipe” through which light flows. This is, of course, merely an analogy for how an optical fiber works, but it conveys the basic idea. The interface between a piece of electronic equipment and an optical fiber consists of an optical *source* (typically an LED or a semiconductor laser) to generate light signals from electrical signals, and an optical *detector* (typically a photodiode or phototransistor) to generate electrical signals from received light signals.

The predominant use of optical fiber in modern industry is as a data communication medium between digital electronic devices, replacing copper-wire signal and network cabling. An illustration showing two digital electronic devices communicating over a pair of optical fibers appears here, each fiber “conducting” pulses of light (representing serial digital data) from an LED source to a photodiode detector:




The following photograph shows a serial converter (the black rectangular plastic box with a blue label) used to convert optical data pulses entering and exiting through orange-jacketed optical cables (on the left) into EIA/TIA-232 compliant electrical signals through a DB-9 connector (on the right) and vice-versa, allowing the electronic serial data device on the right-hand side of the photograph to communicate via fiber optic cabling:



Note how the two optical fiber ports on the converter body are labeled “R” and “T” for *Receive* and *Transmit*, respectively. Serial devices with built-in electronic/optical converters will similarly label their optical ports.

For this device, connection to each of the optical fibers is made using an “ST” style connector, with a quarter-turn locking ring holding each one in place (much like the quarter-turn barrel body of a “BNC” style electrical connector). The next photograph shows a pair of optical fibers terminated with ST-style connectors. White plastic caps cover the connector tips, keeping the glass fiber ends protected from dust and abrasion:



As a data pathway, optical fiber enjoys certain advantages over electrical cable, including:

- Much greater bandwidth (data-carrying capacity), estimated to be in the terahertz range
- Much less equivalent signal power loss per unit cable length (less than 1 dB per kilometer compared with 25 dB per kilometer for coaxial cable)
- Complete immunity to external “noise” sources
- No radiation of energy or data from the cable, thus will not create interference nor be liable to eavesdropping
- No electrical conductivity, allowing safe routing of cables near high voltage conductors
- Total galvanic isolation (i.e. no electrically conductive connection) between data devices, allowing operation at different electrical potentials
- Safe for use in areas with explosive vapors, dust, and/or fibers

These advantages deserve some elaboration. The superior bandwidth of fiber-optic cable is so dramatic that the present-day (2015) limitation on data transfer rates for most fiber-optic installations is the electronic devices at each end, and not the optical fiber itself! This, combined with the low inherent power loss of optical fiber, makes it an ideal medium for long-range data communication such as telephone and internet. Thousands of miles of optical fiber cable have

been buried in underground trenches, laid down on sea floors, strung as overhead lines, and used as “patch” cables in room-scale applications since the advent of affordable optical cabling in the 1980’s. The “tech boom” of the 1990’s saw an impressive amount of trans-continental and inter-continental optical fiber installation, paving the way for the global expansion of internet services into the 21st century. The limitations of electronics at each end of these long fibers means we have not yet begun to tap their full data-carrying capacity, either. Conveying data in photonic – as opposed to electronic – form means there is absolutely no such thing as capacitive or inductive coupling with external systems as there is with conductive wire cable, which not only means optical fiber communication is immune to external interference but also that the optical signals cannot create interference for any other system. Since optical fibers are customarily manufactured from glass which is electrically non-conductive, it is possible to route optical fibers alongside high-voltage power lines, and also connecting together devices at vastly different electrical potentials from each other, with no risk of bridging those differing potentials. Finally, the low power levels associated with optical fiber signals also makes this technology completely safe in areas where explosive compounds in the atmosphere might otherwise be ignited by faults in electrical communications cable.

Optical fibers also suffer from some unique limitations when compared against electrical cable, including:

- Need to avoid tight bend radii for optical cables
- Connections need to be extremely clean
- Specialized tools and skills necessary for installation and maintenance
- Expensive testing equipment

While electrical “transmission line” signal cables must also avoid sharp bends and other discontinuities caused by cramped installations, this need is especially pronounced for optical fiber (for reasons which will be explained later in this section). Since fiber-to-fiber connections consist of glass pressed against glass, the presence of even microscopic contaminants such as dust particles may damage fiber optic connectors if they aren’t cleaned¹⁹ prior to insertion. Cutting, preparing, and terminating optical fiber cables requires its own set of specialized tools and skills, and is not without unique hazards²⁰. Lastly, the test equipment necessary to check the integrity of an optical pathway is similarly specialized and typically quite expensive.

¹⁹Technicians working with optical fiber typically carry pressurized cans of dust-blowing air or other gas to clean connectors and sockets prior to joining the two.


²⁰Chief of which is the potential to get optical fibers embedded in the body, where such transparent “slivers” are nearly impossible to find and extract.

8.4.2 Fiber optic sensing applications

Optical fibers find applications beyond electronic data cable replacement, though, which means they will be a growing presence in the field of industrial instrumentation above and beyond their use as serial data communication cables. Some industrial process transmitters use optical fibers to send and receive light between the transmitter electronics and an optically-based primary sensing element. This may be as simple as a non-contact proximity switch using light to sense the presence of an object within a gap between the two fibers' ends, or as sophisticated as a chemical analyzer relying on the absorption of specific light wavelengths to detect the presence of a chemical substance in a solution.

Turbine flowmeter sensing

One example of a specialized application for optical fibers is shown in this photograph of a paddlewheel-style liquid flowmeter using a pair of optical fibers to convey light to and from the paddlewheel assembly, where the spinning paddlewheel serves to “chop” the light beam and thereby represent liquid flow rate as a frequency of pulsing light:




An end-view of the two optical fibers is shown in the next photograph. When installed as a working system, these two fibers will plug into a flow transmitter device sending a continuous beam of light through one fiber and sensing the pulsed light signal coming back from the paddlewheel through the other fiber:



Fabry-Perot interferometry temperature measurement

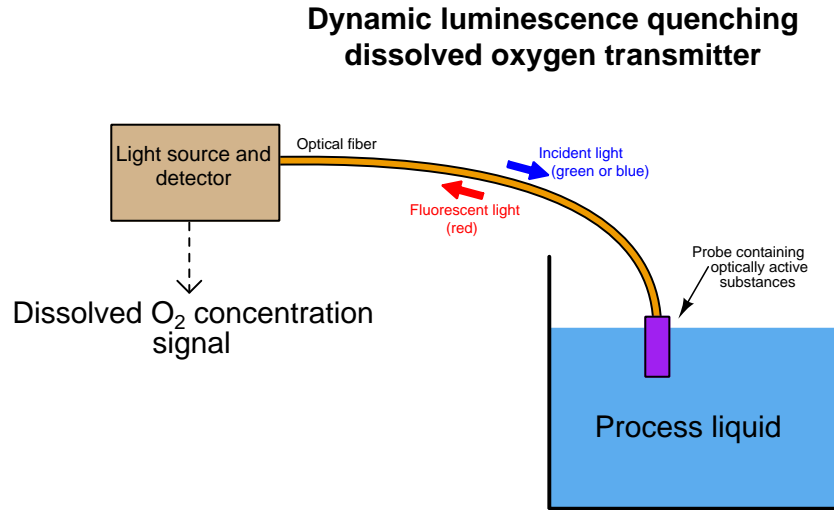
Another example of a specialized application for optical fibers is measurement of high temperatures using the *Fabry-Perot interferometry* method. This technology utilizes a small, thin disk of sapphire as a temperature sensor. The thickness of this disk as well as the speed of light through the sapphire are both temperature-dependent, which means a photon²¹ of light shot at the face of the disk will reflect off the back face of the disk and return to the source at different times depending on the temperature of the disk. In a Fabry-Perot interferometer instrument, the “optical thickness” of the sapphire disk is measured by sending a continuous beam of white light to the disk and receiving the reflected light from the disk through a single optical fiber, the optical interference resulting from the incident and reflected light beams representing the disk’s temperature. This novel method of temperature measurement shows promise for certain challenging industrial process applications such as high-temperature measurement inside slagging coal gasifiers used to efficiently extract energy and chemical feedstocks from coal:



²¹A “photon” is a quantity of light energy represented as a particle, along the same scale as an electron. It isn’t entirely fair to characterize light as either consisting of waves or as consisting of particles, because light tends to manifest properties of both. Actually, this may be said of any sub-atomic particle (such as an electron) as well: under certain conditions these particles act like clumps of matter, and under different conditions they tend to act as waves of electromagnetic energy. This particle-wave duality lies at the heart of quantum physics, and continues to be something of a philosophical mystery simply because the behavior defies the macroscopic constructs we are accustomed to using when modeling the natural world.

Dissolved oxygen measurement

Yet another example of a specialized application for optical fibers is the measurement of dissolved oxygen in aqueous solutions using the *dynamic luminescence quenching* or *fluorescence quenching* method. This technology uses a thin layer of solid material containing molecules known to fluoresce²² with red light when exposed to visible light of a shorter wavelength (typically green or blue). Oxygen molecules present in the liquid solution tend to bond with the fluorescing molecules in the sensor and inhibit that fluorescence, thus providing a means of measuring oxygen concentration near the sensor: the less O₂ dissolved in solution, the stronger the fluorescence (i.e. more red light received, for a longer duration); the more O₂, the less fluorescence. Optical fibers convey both the incident (green or blue) and returned (red) light between the wet sensing element and the interpreting electronics.



In both the Fabry-Perot interferometry and the fluorescence quenching sensors, the function of the fiber optic cable is to physically separate the sensing element from the sophisticated and fragile electronic transmitter needed to interpret the optical signal as a process variable measurement.

²²Fluorescence is the phenomenon of a substance emitting a long-wavelength (low-energy) photon when "excited" by a short-wavelength (high-energy) photon. Perhaps the most familiar example of fluorescence is when certain materials emit visible light when exposed to ultraviolet light which is invisible to the human eye. The example of fluorescence discussed here with dissolved oxygen sensing happens to use two different colors (wavelengths) of visible light, but the basic principle is the same.

Arc flash detection

There is at least one application where the optical fiber itself *is* the sensing element: arc flash detection within high-voltage switchgear cabinets. “Arc flash” is the phenomenon of intense heat and light developed at a high-current electrical arc, especially a phase-to-phase arc between electric power conductors where there is little circuit resistance to limit fault current. High-voltage switchgear is constructed in such a way as to extinguish the arc normally developed at the contacts during each “opening” cycle, but certain faults within a piece of switchgear may inhibit this extinguishing function. In such cases the potential for equipment damage and threat to human health and life is severe.


If a bare (unjacketed) optical fiber is properly arranged within a piece of switchgear, an arc flash event will inject enough light through the fiber that some of it will be detected at the far end where it meets a light-sensitive receiver. Since the fiber itself is not electrically conductive, there is no risk of conducting a high-voltage arc back to this receiver. The receiver, meanwhile, serves the purpose of commanding any “upstream” switchgear to trip open in the event of a detected arc fault. The early detection of arc flash by optical means rather than the time-delayed detection of the same fault by overcurrent or current-imbalance detection results in much faster clearing of the faulted switchgear from the power grid, both limiting equipment damage and limiting the potential for injury or death.

8.4.3 Fiber optic cable construction

Communication-grade optical fibers are manufactured from fused silica (SiO_2) glass of exceptional purity²³. A single strand of optical fiber made from this glass called the “core” serves as a waveguide for the light. The core is surrounded by another layer of glass called the “cladding” which has a different index of refraction²⁴ necessary to “channel” the majority of the optical energy through the core and inhibit “leakage” of optical power from the cable. Additional layers of plastic and other materials around the core/cladding center provide coloring (for fiber identification in multi-fiber cables), protection against abrasion, and tensile strength so the cable will not suffer damage when pulled through conduit.

The purpose of building a fiber optic cable with a core and a cladding having different refractive indices (i.e. different speeds of light) is to exploit a phenomenon called *total internal reflection*, whereby rays of light reflect off the interface between core and cladding to prevent its unintentional escape from the core at any point along the length of the fiber.

When light crosses an interface between two materials having different speeds, the light beam will become *refracted* as a function of those two speeds as described by *Snell's Law*:




Snell's Law relates the sine of the incident angle to the sine of the refracted angle as a ratio to each material's speed of light, the material possessing the greatest speed of light (i.e. the lowest refractive index value) exhibiting the greatest angle as measured from perpendicular to the interface:

$$\frac{\sin \theta_1}{v_1} = \frac{\sin \theta_2}{v_2}$$

²³Impurities such as metals and water are held to values less than 1 part per *billion* (ppb) in modern optical fiber-grade glass.

²⁴The “index of refraction” (n) for any substance is the ratio of the speed of light through a vacuum (c) compared to the speed of light through that substance (v): $n = \frac{c}{v}$. For all substances this value will be greater than one (i.e. the speed of light will always be greatest through a vacuum, at 299792458 meters per second or 186282.4 miles per second). Thus, the refractive index for an optically transparent substance is analogous to the reciprocal of the *velocity factor* of an electrical transmission line, where the permittivity and permeability of the cable materials act to slow down the propagation of electric and magnetic fields through the cable.

According to Snell's Law, there will be a critical angle at which the incident light ray will refract to being parallel to the interface. Beyond this critical angle, the light ray ideally reflects off the interface and never enters the second material at all. This is the condition of *total internal reflection*, and it is what we desire in an optical fiber where the core is the first material and the cladding is the second material:



Both the core and cladding of an optical fiber are manufactured from the same base material of ultra-pure fused silica, but “doped” with specific impurities designed to alter the refractive index of each one (raising the refractive index of the core to decrease its optical velocity and lowering the refractive index of the cladding to increase its optical velocity).

The diameter of core and cladding vary with the type of optical fiber, but several standard sizes have emerged in the industry, each one specified by the diameter of the core followed by the diameter of the cladding expressed in *microns* (millionths of a meter). A common optical fiber standard in the United States is 62.5/125 (62.5 micron core diameter, 125 micron cladding diameter), and 50/125 in Europe. Some less common standard core/cladding diameters²⁵ include 85/125 and 100/140.

²⁵All of these sizes refer to *glass* fibers. Plastic-based optical fibers are also manufactured, with much larger core diameters to offset the much greater optical losses through plastic compared to through ultra-pure glass. A typical plastic optical fiber (POF) standard is specified at a core diameter of 980 microns and a cladding diameter of 1000

To give some perspective on the physical size of an optical fiber core, the following photograph shows the end-view of an “ST” style fiber optic connector for a 50/125 micron cable, held by my hand. A green LED light source is shining into the other end of this cable, the tiny green dot visible at the center of the ST connector revealing the diameter of the 50 micron core:



Several other layers of material must be placed over the core and cladding to form a rugged optical fiber. A plastic *jacket* with a typical diameter of 250 microns (0.25 mm) covers the cladding, and provides a base for color-coding the fiber. This three-layer construction of core, cladding, and jacket is known in the industry as *Primary Coated Optical Fiber*, or *PCOF*.

PCOF is still too fragile for end-user applications, and so another layer of plastic is typically added (900 microns in diameter) to make the fiber *Secondary Coated Optical Fiber*, or *SCOF*. When wrapped with fiberglass or Kevlar fibers around the secondary jacket for tensile strength, and a protective PVC plastic outer layer to protect against abrasion, the cable becomes suitable for indoor use. Cables suitable for outdoor, direct burial, and undersea applications usually take the form of groups of PCOF fibers packaged within an extremely rugged encasement with metal strands for tensile strength. Sometimes a gel material helps cushion the fibers from each other within the confines of the cable sheath.


microns (1 millimeter)!

8.4.4 Multi-mode and single-mode optical fibers

In any sort of waveguide – optical, electrical, or even acoustical (sound) – the signal energy may be able to propagate down the waveguide in different orientations. This is true for optical fibers where the core diameter is relatively large²⁶ compared to the wavelength of the light: there will be many alternative pathways for light to travel along the length of a fiber’s core. Optical fibers with core diameters of 50 microns or more are referred to as *multi-mode* fibers, because multiple independent pathways, or “modes”, of light are possible within the core’s width.

If an optical fiber’s core is manufactured to be small enough, relative to the wavelength of the light used, the fiber will only support one “mode” or pathway down its core. Such fiber is called *single-mode*. Single-mode fiber cores typically range from 4 to 10 microns in diameter, with 8 micron being typical.


The purpose of single-mode optical fiber is to avoid a problem called *modal dispersion*. When multiple “modes” of light propagate down the length of an optical fiber, they don’t all have the same length. That is to say, some modes actually take a straighter (and more direct) path down the fiber’s core than others. The reason this is a problem is that this phenomenon corrupts the integrity of high-speed (i.e. short-period) pulses. An exaggerated illustration of this problem appears here, showing the relative path lengths of three different light rays, each one entering the fiber core at a slightly different angle. The light ray closest to parallel with the core’s centerline finds the shortest “mode” to the fiber’s end, and arrives in the least amount of time:



With different “modes” of light arriving at different times from the same incident pulse, the received light pulse at the exiting end of the fiber will no longer possess a crisp “square-wave” shape. Instead, the pulse will be “smeared” over time, occupying a larger time span. This poses a bandwidth limit on the fiber, as there will be some maximum pulse frequency at which adjacent pulses will begin to merge together and become indistinguishable. The longer the length of optical fiber, the more pronounced this dispersion will be. This problem is most evident in applications where the fiber length is very long (hundreds of miles) and the data rate is very high (hundreds of megahertz). Thus, it is a significant problem for long-distance data trunk cables such as those used for transcontinental and intercontinental internet traffic.

²⁶A common core size for “multi-mode” optical fiber is 50 microns, or 50 micro-meters. If a wavelength of 1310 nanometers (1.31 microns) is used, the core’s diameter will be $\frac{50}{1.31}$ or over 38 times the wavelength.

Single-mode optical fiber completely averts this problem by eliminating²⁷ multiple modes within the fiber core. When there is only one mode (pathway) for light to travel, there will be exactly one distance for light to travel from one end of the fiber to the other. Therefore, all portions of the incident light pulse experience the same travel time, and the light pulse arrives at the far end of the cable suffering no modal dispersion:




As you can imagine, single-mode fiber is more challenging to splice than multi-mode fiber, as the smaller core diameter provides less room for alignment error.

A compromise solution to the problem of modal dispersion in multi-mode fibers is to manufacture the core glass with a *graded* index of refraction rather than a homogeneous index of refraction. This means the concentration of doping material in the glass varies from the center of the core to the outer diameter of the core where it interfaces with the cladding. The result of this graded dispersion is that modes traveling closest to the core's centerline will experience a slower speed of light (i.e. greater index of refraction) than modes near the edge of the core, which means the difference in travel time from one mode to the next will be less pronounced than within normal "step-index" fibers. Of course, this also means graded-index optical fiber is more costly to manufacture than step-index optical fiber.

²⁷The most straight-forward way to make an optical fiber single-mode is to manufacture it with a skinnier core. However, this is also possible to achieve by increasing the wavelength of the light used! Remember that what makes a single-mode optical fiber only have one mode is the diameter of its core *relative to the wavelength of the light*. For any optical fiber there is a *cutoff wavelength* above which it will operate as single-mode and below which it will operate as multi-mode. However, there are practical limits to how long of a wavelength we can make the light before we run into other problems, and so single-mode optical fiber is made for standard light wavelengths by manufacturing the cable with an exceptionally small core diameter.

8.4.5 Fiber optic cable connectors, routing, and safety

One of the most popular styles of single-fiber connector is the so-called “ST” style, which uses a quarter-turn locking barrel to secure the connector into its matching socket:



Communication patch cables such as the one shown above come in *pairs* of fibers, one for receiving and one for transmitting. Note how the plastic strain-relief grips between the metal barrel of each connector and each orange-jacketed cable are color-coded (one white, one black) for easy identification at each end of the cable.

An older style of connector based on the type used to connect small coaxial cables together is the “SMA” style, which used a threaded barrel to lock each fiber in place. The SMA-style connectors were very secure, but laborious to engage and disengage due to the fine pitch of the barrel’s threads and the subsequent need to turn the barrel multiple rotations (versus one-quarter turn of the barrel for an ST connector).

Given that communication patch cables typically have two fibers (one for each direction of data flow), connector styles have emerged to accommodate fiber pairs. One such style is the so-called “SC” connector, with a pair of side-by-side plugs accommodating twin optical fibers.


Terminating a bare cable of fibers with individual connectors is a time-consuming process, requiring the technician to unbundle the individual fibers, strip the jacketing off of each one to reveal the core and cladding, cleave each glass fiber to give it a flat end, and finally insert and secure each fiber into its respective connector. Typical fiber connectors use either a “hot-melt” or a chemical epoxy system of attachment, where the glue adheres to the strain-relief fibers of the cable for tensile strength, while the central glass fiber protrudes through a small hole in the center of the connector tip. This protruding glass fiber must be carefully cut and polished to produce a flat end suitable for engagement with another optical fiber aligned to its center.

Optical fibers may be spliced mid-way in a cable run, although this practice should be avoided whenever possible. If the fibers are multi-mode, the splicing may be done using “butt” connectors

but the power losses may be unacceptable. Alternatively, stripped fibers may be inserted into both ends of a small-diameter tube filled with gel having the same index of refraction as the core glass, to “conduct” light with as little loss as possible from one fiber core to the other.

A very good technique often applied to single-mode fiber is that of *fusion splicing*, where two single-mode fiber ends are literally melted together using an electric arc so that they form one seamless glass fiber. The alignment of fibers prior to fusion is done under the view of a microscope, and often with the aid of a light source on one end and an optical power meter on the other end to give a quantitative measurement of alignment accuracy. When the two fibers are aligned as close as possible, the electric arc is fired to melt the two fibers together, creating a single fiber. Fusion splicing is the method of choice for long-distance runs of single-mode fiber, where low power loss and high integrity of the splice are paramount factors.

When laying optical fiber in wire trays, pulling through rigid conduit, or arranging it in connection panels, an important physical consideration is to maintain a minimum bend radius²⁸ at all points along the fiber’s length. This is important because sharp bends will cause light to “leak” out of the fiber core and into the cladding where it may then escape the cable altogether. A sharp bend in an optical fiber will cause the angle between the light ray and the core/cladding interface to reach the critical point where total internal reflection no longer occurs:



The light leakage from an optical fiber may be dramatic if the bend is sharp enough. On an indoor cable, using visible laser light, you can actually see the light “leak” through to the PVC outer coating on the outside of the cable!

Junction boxes and connection panels where excess lengths of fiber optic cabling may be coiled will typically provide plastic forms over which those loops of cable may be bent, the radius of that plastic form exceeding the manufacturer’s specification for minimum bend radius.

A common way in which the minimum bend radius requirement is violated is when a cable tie is used to anchor a fiber optic cable to some sturdy surface such as a wiretray or a cabinet post. The sharp bend created by the tension of a tightened cable tie on the fiber optic cable will easily exceed the minimum bend radius for that cable, creating light leakage and subsequent performance problems. Therefore, a good installation practice for fiber optic cables is to always leave cable ties loose enough that they do not tightly grip the cable.

²⁸Typically a few inches for multi-mode fiber.

There are multiple safety concerns when working with optical fibers, both when installing them and when doing maintenance-type work. Installation hazards center around dangers of the glass fiber itself, while maintenance hazards center around the light sources used to “power” the optical fibers.

Installation of fiber optic cable requires that individual glass fibers be separated from each other in a multi-fiber cable and each one terminated with a connector, and this requires at some point that the technician strip each fiber down to its glass core and cladding. Both the core and the cladding are extremely small in diameter, and are made of ultra-pure glass. If a piece of core/cladding breaks off the fiber and penetrates the skin, the resulting “sliver” will be nearly invisible due to its exceptional transparency. Its outer surface is also very smooth, making extraction difficult. Unextracted pieces of an optical fiber, if left in the body, can actually migrate through the victim’s flesh and become buried even deeper to the point where they can cause serious health problems.

Technicians working with optical fiber typically lay a length of adhesive tape, sticky-side up, on whatever workbench or table they are using to prepare the cable, as a tool to catch any loose fiber ends they cut off. At the conclusion of the job, this length of tape is carefully rolled up and then disposed of in the same manner that “sharps” may be disposed of in a medical environment.

Maintenance technicians working with functioning fiber optic systems need to be careful when disconnecting “hot” fibers, due to the intensity of the light used in some systems. This is especially true of long-distance telecommunication fibers using laser sources rather than regular LEDs, which may have power levels reaching a half watt or so. One-half of a watt doesn’t sound like very much power, but when you consider this power level is concentrated over a circular area with a diameter less than 10 microns (for single-mode fiber), the watt-per-square-meter value is actually large enough to cause significant temperature increases wherever the light beam happens to fall. In fact, you can actually damage a fiber-optic connector on such a system by disconnecting the fiber with the fiber “powered”, the laser light being intense enough to burn and pit the aluminum ferrule of the connector!

Even standard LED light sources may pose a hazard if a technician directly views the end of the cable with his or her eyes, due to the focused nature of the light beam. The retina of your eye is extremely sensitive to light, and may easily be damaged by viewing such an intensely focused beam coming out of an optical fiber, where the entire LED’s light output is channeled into a core just a fraction of a millimeter in diameter. The optical hazard is even greater when infra-red light sources are used, because there is no visible indication of the light’s presence. A technician won’t even be able to see the light coming out, yet it could still be intense enough to damage their retina(s).

Laser-sourced fibers should never be unplugged from the equipment. One should treat a laser-sourced fiber with the same respect as a “live” electrical circuit, and use the same lockout/tagout procedures to ensure personnel safety. In systems using visible light wavelengths, a safe way to view the light coming out the end of an optical fiber is to point the fiber end at a piece of paper and look for the colored dot falling on the paper. The paper’s rough surface scatters the light so that it is no longer a focused beam.

The only time it is truly safe to view the end of an optical fiber to check for light is when the light source is something diffuse such as natural sunlight or a flashlight. It is common for technicians to use a flashlight to identify fibers from one end of a multi-fiber cable to the other, one technician shining the flashlight at the end of one fiber while another technician views all the fibers at the other end of the cable to see which one is lit.

Some optical communications equipment come equipped with a feature called an *Open Fiber Control* (OFC) safety system, which turns off all light sources on a channel whenever an interruption of light is detected at the receiver port. Since most duplex (two-way) optical fiber channels consist of two fibers (one for each direction of light), a break in any one fiber will darken one receiver, which then commands the transmitter port on that equipment to darken as well to prevent anyone getting injured from the light. It also completely disrupts communication in that channel, requiring a re-initialization of the channel after the fiber is plugged back in.

8.4.6 Fiber optic cable testing


Optical fibers, like electrical communications cable, may need to be tested to measure certain performance characteristics. Such testing is commonplace for new installations of fiber optic cabling to ensure all installed cable lengths and connectors are functioning properly. Repeated tests over time, compared with the initial installation test, quantifies any degradation of cables or connectors. Another common testing procedure, called *acceptance testing*, tests the optical cable while it is still on the spool prior to installation.

Two basic types of optical fiber tests are presented here: one where the power level of light is measured at the far end of the fiber from a source of known optical power, and another where a pulse of light is sent down a fiber and the light received at the same end of the fiber is analyzed. The former test is simply a measurement of optical power, while the latter test is a sophisticated analysis of light over very brief periods of time (time domain reflectometry).

Optical power loss testing


Perhaps the simplest quantitative test of an optical fiber consists of shining a light source of known optical power at one end of a fiber and monitoring the amount of optical power received at the other end of the fiber. This type of test is typically performed with two pieces of equipment: the source and the power meter.

First, the optical power meter and light source are short-coupled together using a pair of patch cables and a single “butt” connector:



Once light is received by the optical power meter, the technician presses the “zero” button to set the baseline or reference point for all future power measurements. Although some light will be lost in the two patch cables and connector, this amount of loss will also be present in the final test and so it must be ignored.

After “zeroing” the optical power meter, the actual fiber to be tested is connected between the light source and the power meter. Any *additional* light lost within the tested fiber²⁹ will register at the power meter as a negative decibel figure:



Recall that the definition of a “decibel” is 10 times the common logarithm of the power ratio between output and input for any system:

$$\text{dB} = 10 \log \left(\frac{P_{out}}{P_{in}} \right)$$

Thus, the power loss of -0.6 dB shown in the illustration represents 87.1% of the optical source power received by the optical power meter. Decibels are very commonly used as an expression of power gain and loss in communication system testing, because dB figures directly add when components are connected in series with each other. For example, if we knew that a certain type of “butt” connector for optical fiber exhibited a typical power loss of -1.2 dB and that three of these connectors would be used to join a single run of fiber, we would know to expect a total connector loss of -3.6 dB (i.e. 3×-1.2 dB).

²⁹Not just light lost along the length of the fiber, but also at each connector on the fiber, since placing the test fiber within the optical path between the light source and optical power meter necessarily introduces another pair of connectors where light may be lost.


Excessive optical power losses may be caused by a number of different factors, including:

- Poor alignment between fibers in a connector
 - Connector flaw causing fibers to be mis-aligned (e.g. angular misalignment)
 - Fiber flaw causing mis-alignment in a good connector (e.g. cores not concentric)
- Mismatched fiber sizes (e.g. 62.5 micron core sending light into a 50 micron core)
- Oil or debris on the end of a connector
- Rough (improperly polished) end on one or more fibers
- Minimum bend radius violated at any point along the fiber's length
- Cracked fiber core

Unfortunately, a power meter test will not indicate what kind of flaw is causing excessive power loss, nor where that flaw might be located. If the cable in question has removable connectors mid-way in its length, the power meter and/or source may be relocated to test portions of the cable to determine which section contributes more to the power loss, but an end-to-end power test cannot pinpoint the location or the type of fault.


OTDR testing

An *Optical Time-Domain Reflectometer* or *OTDR* is a sophisticated test instrument used to probe the characteristics of long optical fibers. They work by injecting a very brief pulse of light into one end of a long optical fiber, then monitoring any light received at that same end of the fiber. As the light pulse travels down the length of the fiber, it continuously loses some of its magnitude due to scattering in the glass. Some of this scattered light returns back to the source-end of the fiber, presenting a sort of “continuous echo” of the moving pulse. This continuous echo is analogous to the noise heard from an object moving away from the listener. As the light pulse encounters flaws and other discontinuities in the fiber and/or connectors along its length, the echoed signal changes in magnitude. This received signal is displayed as a time-domain plot on the OTDR viewing screen, and will look something like this:



The “trace” shown on the display screen of an OTDR is a plot of the received optical signal strength over time. A large “spike” at the left-hand side of this trace marks the incident pulse of light injected into the optical fiber by the OTDR from the traveling pulse as it propagates down the length of the fiber. All signals after that (to the right of that initial “spike”) represent light received from that same end of the optical fiber. In a completely uniform fiber the resulting “echo” would trace a downward-sloping straight line as the traveling light pulse gradually weakens. In an imperfect fiber, any discontinuities such as splices, connector joints, sharp bends, cracks, etc. will cause the traveling light pulse to lose more photons than usual at the location of the discontinuity: sometimes returning a strong echo back toward the OTDR and other times not. A discontinuity such as a mis-aligned fiber connector will tend to return a strong echo as part of the traveling light pulse reflects off the mis-aligned connector end and returns to the OTDR. A discontinuity such as a mal-formed fusion splice merely scatters a greater-than-normal amount of light out through the fiber’s cladding, in which case there is no echo “pulse” received by the OTDR but rather just a further weakening of the echo signal.

The OTDR trace shown in the previous illustration demands further explanation. Shown here is a magnified view of it, complete with numbers to identify each noteworthy event:



Legend:

1. Incident pulse output by the OTDR, and injected into the launch fiber
2. Reflection off the face of the near-end connection between the launch fiber and the fiber under test
3. Loss of light due to a non-reflective discontinuity (e.g. sharp bend, splice)
4. Loss of light due to a reflective discontinuity (e.g. mis-aligned connector)
5. Reflection off the face of the far-end connection at the end of the fiber under test
6. The “noise floor”

As you can see, an OTDR trace provides much more information about the performance of an optical fiber than a simple power test. Each flaw in the cable or its associated connectors appears as a deviation from the normal downward-sloped line of the trace, the location in time revealing the distance between the OTDR and the flaw. Thus, an OTDR not only indicates the nature of each flaw, and the amount of optical power lost at each flaw, but also the *location*³⁰ of each flaw along the fiber’s length. One important caveat exists for this distance calculation, and that is the fact that the length of a fiber in a multi-fiber cable will always be somewhat longer than the length of the cable itself, since individual fibers inside a cable are often “wound” in a spiral configuration or otherwise deviating from the straight centerline of the cable. “Loose tube” cables, for example, often exhibit fiber lengths 5% to 10% greater than the physical length of the cable itself.

³⁰Since distance along any path is simply the product of speed and time ($x = vt$), and the speed of light through an optical fiber is a well-defined quantity ($v = \frac{c}{n}$ where n is the core’s index of refraction), the distance between the OTDR and the flaw is trivial to calculate.

8.5 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Sanitary applications:** In processes where bacterial growth cannot be tolerated, there must never be stagnant pockets in piping systems for cultures to reside.
- **Fluid seals:** accomplished by maintaining tight contact between solid surfaces. The shapes of these surfaces are generally conical (e.g. tapered pipe threads, tubing ferrules).
- **Electrical connection integrity:** the integrity of electrical connections is absolutely essential to system reliability. Maintaining firm, clean contact between mating conductor surfaces is necessary to ensure this.
- **Analog vs. digital signals:** analog signals have infinite resolution but are susceptible to corruption by noise. Digital signals have limited resolution but are tolerant of any noise measuring less than the difference in thresholds between the high and low states.
- **Electromagnetic induction:** occurs only when magnetic fields are perpendicular to the conductor. Relevant to signal coupling in cables, mitigated by twisting cable conductors to form opposing loops from the perspective of an external magnetic field.
- **Electrostatic coupling:** occurs when electric fields bridge between conductors, and cannot occur "behind" a grounded conductor. Relevant to signal coupling in cables, mitigated by building cables with shield conductors and grounding the shield at one cable end.
- **Lenz's Law:** any magnetic field arising from electromagnetic induction opposes the inducing field. Relevant to determining the directions of induced current in a twisted-pair cable exposed to an AC magnetic field.
- **Capacitance:** $C = \frac{\epsilon A}{d}$, capacitance being proportional to the area of two overlapping conductors (A), the permittivity of the insulating (dielectric) substance between them (ϵ), and the distance (d) separating the conductors. Relevant to electrostatic coupling, where the degree of coupling between two electrical conductors is directly proportional to the overlapping area and inversely proportional to the distance between the conductors.
- **Transmission lines:** short-duration (pulsed) electrical signals travel along a cable at nearly the speed of light, reflecting off the end of that cable if not properly terminated. Relevant to signal cables carrying high-frequency signals.
- **Decibels:** used to express the ratio of one power to another in logarithmic form, such that the sum of component dB values is equivalent to the product of those same component gain/loss ratios. Decibels may also be used to express a power quantity relative to some reference power value such as 1 milliwatt (dBm) or 1 watt (dBW). Decibels are an example of a mathematical *transform function*, whereby one type of mathematical problem (multiplication/division) is transformed into an easier type of problem (addition/subtraction).

- **Time-Domain Reflectometry:** the analytical technique of sending a pulse signal down a transmission line or waveguide and analyzing the characteristics of that transmission line or waveguide by the reflected signal. Relevant to the testing of electrical cables as well as optical fibers.
- **Snell's Law:** relates the angle of a refracted light ray at the interface of two transparent substances to the relative speeds of light through those two substances. Relevant to the core and cladding materials of an optical fiber, the speeds of light for those two substances altered in such a way as to produce total internal reflection of the light rays within the core. This keeps the light in the fiber, inhibiting leakage and subsequent signal loss.

References

“An Overview of Fiber Optic Technology” whitepaper BB-WP13-r0-0305, B&B Electronics Manufacturing Company.

Austin, George T., *Shreve’s Chemical Process Industries*, McGraw-Hill Book Company, New York, NY, 1984.

“Connections” report FL-01 2-94, EBAA Iron Sales, Inc., 1994.

“CPITM Tube Fittings”, catalog 4230, Parker Hannifin Corporation, Cleveland, OH, 2000.

Croft, Terrell and Summers, Wilford I., *American Electrician’s Handbook*, Eleventh Edition, McGraw-Hill Book Company, New York, NY, 1987.

Dutton, Harry J.R., *Understanding Optical Communications*, First Edition, document SG24-5230-00, International Business Machines Corporation, September 1998.

“Fitting Installation Manual”, Hoke Incorporated, Spartanburg, SC, 1999.

“Gaugeable Tube Fittings and Adapter Fittings”, document MS-01-140, revision 7, Swagelok Company, MI, 2004.

Graves, W.V., *The Pipe Fitters Blue Book*, W.V. Graves Publisher, Webster, TX, 1973.

“Industrial Pipe Fittings and Adapters”, catalog 4300, Parker Hannifin Corporation, Columbus, OH, 2000.

Ivanov, Georgi, *Fabry-Perot Sapphire Temperature Sensor for use in Coal Gasification*, Master of Science thesis paper in Electrical Engineering, Blacksburg, VA, 3 May 2011.

Morrison, Ralph, *Grounding and Shielding Techniques in Instrumentation*, John Wiley and Sons, Inc., NY, 1967.

“Oxymax COS61D dissolved oxygen sensor with Memosens protocol operating instructions”, BA00460C/07/EN13.12, Endress+Hauser.

“Pipe Fittings”, document MS-01-147, revision 3, Swagelok Company, MI, 2002.

“Piping Joints Handbook”, document D/UTG/054/00, BP Amoco, 2000.

“Thread and End Connection Identification Guide”, document MS-13-77, revision 3, Swagelok Company, 2005.

Chapter 9

Discrete process measurement

The word “discrete” means *individual* or *distinct*. In engineering, a “discrete” variable or measurement refers to a true-or-false condition. Thus, a discrete sensor is one that is only able to indicate whether the measured variable is above or below a specified setpoint.


Discrete sensors typically take the form of *switches*, built to “trip” when the measured quantity either exceeds or falls below a specified value. These devices are less sophisticated than so-called *continuous* sensors capable of reporting an analog value, but they are quite useful in industry.

Many different types of discrete sensors exist, detecting variables such as position, fluid pressure, material level, temperature, and fluid flow rate. The output of a discrete sensor is typically electrical in nature, whether it be an active voltage signal or just resistive continuity between two terminals on the device.

9.1 “Normal” status of a switch


Perhaps the most confusing aspect of discrete sensors is the definition of a sensor’s *normal* status. Electrical switch contacts are typically classified as either *normally-open* or *normally-closed*, referring to the open or closed status of the contacts under “normal” conditions. But what exactly defines “normal” for a switch? The answer is not complex, but is often misunderstood due to the ambiguous nature of the word *normal*.

The “normal” status for a switch is the status its electrical contacts are in *during a condition of no physical stimulation*. Another way to think of the “normal” status is to think in terms of the switch being *at rest*. For a momentary-contact pushbutton switch, this would be the status of the switch contact when it is *not* being pressed. Electrical switches are always drawn in schematic diagrams in their “normal” statuses, regardless of their application. For instance, the following diagram shows a normally-open pushbutton switch controlling a lamp on a 120 volt AC circuit (the “hot” and “neutral” poles of the AC power source labeled L1 and L2, respectively):




We can tell this switch is a normally-open (NO) switch because it is drawn in an open position. The lamp will energize only if someone presses the switch, holding its normally-open contacts in the “closed” position. Normally-open switch contacts are sometimes referred to in the electrical industry as *form-A* contacts.

If we had used a normally-closed pushbutton switch instead, the behavior would be exactly opposite. The lamp would energize if the switch was left alone, but it would turn off if anyone pressed the switch. Normally-closed switch contacts are sometimes referred to in the electrical industry as *form-B* contacts:



This seems rather simple, don’t you think? What could possibly be confusing about the “normal” status of a switch? The confusion becomes evident, though, when you begin to consider *process* switches (i.e. switches actuated by process measurements such as pressure, flow, level, etc.). In order to better understand this concept, we will consider a simple application of a *flow* switch: a switch built to actuate when a sufficient rate of fluid flows through a pipe.

A flow switch is built to detect fluid flow through a pipe. In a schematic diagram, the switch symbol appears to be a toggle switch with a “flag” hanging below. The schematic diagram, of course, only shows the circuitry and not the pipe where the switch is physically mounted:



This particular flow switch is used to trigger an alarm light if coolant flow through the pipe ever falls to a dangerously low level, and the contacts are *normally-closed* as evidenced by the closed status in the diagram. Here is where things get confusing: even though this switch is designated as “normally-closed,” it will spend most of its lifetime being held in the open status by the presence of adequate coolant flow through the pipe. Only when the flow through the pipe slows down enough will this switch return to its “normal” status and conduct electrical power to the lamp. In other words, the “normal” status for this switch (closed) is actually an *abnormal* status for the process it operates within (low flow), for the simple reason that the switch should be stimulated and not at rest while the process is operating as it should.


Students often wonder why process switch contacts are labeled according to this convention of “no stimulation” instead of according to the typical status of the process in which the switch is used. The answer to this question is that the manufacturer of the switch has no idea whatsoever as to your intended use. A flow switch manufacturer does not know or care whether their product gets used as a low-flow detector or as a high-flow detector. In other words, the manufacturer cannot predict what the typical status of *your* process will be, and so the definition of “normal” status for the switch must be founded on some common criterion unrelated to your particular application. That common criterion is the resting status: when the sensor is exposed to the *least* (or no) amount of stimulation from the process it senses.

Here is a listing of “normal” definitions for various process switch types:

- **Limit switch:** target not contacting the switch
- **Proximity switch:** target far away
- **Pressure switch:** low pressure (or even a vacuum)
- **Level switch:** low level (empty)
- **Temperature switch:** low temperature (cold)
- **Flow switch:** low flow rate (fluid stopped)

These are the conditions represented by the switch statuses shown in a schematic diagram. These may very well *not* be the statuses of the switches when they are exposed to *typical* operating conditions in the process.


A helpful tip to remember about process switches and their respective schematic diagram symbols is that the symbols are conventionally drawn in such a way that an *upward* motion of the movable switch element represents *increasing stimulus*. Here are some examples of this, showing various process switch types and NO/NC contact configurations, comparing their states with no stimulus versus when the stimulus exceeds the each switch’s threshold or “trip” setting. The *normal* status of each switch as defined by the manufacturer is labeled in green text:



It is imperative¹ to remember that the way a switch is drawn in a schematic diagram merely represents its “normal” status as defined by the manufacturer. This may or may not be the switch’s status during “typical” operation of the process, and it may or may not be the status of that switch at the time of concern when you are examining the schematic! The “normal” status of a switch means just one thing: what that switch will do when subjected to minimum stimulus – that is to say, what it will do when its stimulus is less than the actuation threshold of the switch.

¹Mistaken interpretations of switch status remains one of the most resilient misconceptions for students first learning this topic. It seems that a great many students prefer to think of a switch’s drawn status as its status at the present moment (e.g. when the process is running as expected). I believe the heart of this misconception is the meaning of the word “normal,” which to most peoples’ minds refers to “the way things typically are.”

For example, the LED in this circuit will turn on if the liquid level rises above 14 inches *and* the pressure falls below 22 PSI *and* either the flow is less than 3 gallons per minute *or* the temperature is greater than 125 degrees Fahrenheit:




Since we know we need a switch to be *closed* in order to conduct electricity and provide a path for current in this circuit, we are looking for the necessary conditions to close each switch. For any normally-closed (NC) switch, this means a stimulus value less than the actuation threshold. For any normally-open (NO) switch, this means a stimulus value great enough to exceed the threshold and “hold” the switch in its actuated state. Since the flow and pressure switches in this circuit are both NC, we are looking for flow and pressure values less than the respective settings. Since the level and temperature switches are both NO, we are looking for level and temperature values in excess of their respective settings.

The present status of a switch may be determined by comparing its stimulating quantity against its trip (threshold) setting. A switch will be in its “normal” (resting) state when the stimulus value is less than the threshold value. Conversely, a switch will be in its “actuated” state when the stimulus value exceeds the threshold value. Determination of a switch’s status, therefore, is a matter of comparing the stimulus quantity to the threshold “trip” setting. One cannot simply look at the schematic diagram to tell what the switch is doing – one must *compare* the switch’s setting versus against a known stimulus value in order to tell whether it will be in its resting state or not.

Likewise, if we happen to know the switch’s present status in a system, we may qualitatively determine the stimulating quantity by comparing the present status against the “normal” (resting) status. If a switch is in its resting state, then the stimulating quantity must be less than the trip threshold. If a switch is in its actuated (non-normal) state, then the stimulating quantity must be greater than the trip threshold². The next example showcases these determinations.

²In this discussion I am deliberately omitting the detail of *deadband* for process switches, for the sake of simplicity.

In this next example, we see a pictorial representation of multiple switches wired to the discrete input channels of a programmable logic controller (PLC), with red LED indicators denoting the real-time status of each input on the PLC:




We may determine a switch's degree of stimulation by comparing its present status with its "normal" status. If a switch happens to be in the same state as its normal state (i.e. resting), we know its stimulus must be less than the threshold (trip) value. If a switch happens to be in the opposite state as its normal state (i.e. actuated), we know its stimulus has exceeded the threshold (trip) value. The states of these four switches may be shown in a table:

Switch	Normal status	Present status	Trip value	Stimulus
Pressure	Normally-closed (NC)	Closed	12 PSI	$P < 12 \text{ PSI}$
Temperature	Normally-open (NO)	Closed	235 °F	$T > 235 \text{ °F}$
Level	Normally-closed (NC)	Closed	4 inches	$L < 4 \text{ inches}$
Limit	Normally-open (NO)	Open	n/a	<i>no contact</i>


9.2 Hand switches

A *hand switch* is exactly what the name implies: an electrical switch actuated by a person's hand motion. These may take the form of toggle, pushbutton, rotary, pull-chain, etc. A common form of industrial pushbutton switch looks something like this:




The threaded neck inserts through a hole cut into a metal or plastic panel, with a matching nut to hold it in place. Thus, the button faces the human operator(s) while the switch contacts reside on the other side of the panel.


When pressed, the downward motion of the actuator breaks the electrical bridge between the two NC contacts, forming a new bridge between the NO contacts:



The schematic diagram symbol for this type of switch looks much like the real thing, with the normally-closed contact set on top and the normally-open contact set below:




Connecting two of the terminals together makes this form of switch electrically identical to a *Form C*:




This switch contact arrangement is sometimes referred to as a *form-C* contact set, since it incorporates both a form-A contact (normally-open) as well as a form-B contact (normally-closed).

Most industrial hand switches are available in modular form, where sets of switch contact blocks may be “stacked” together to be actuated by the same pushbutton or rotary knob. This allows an almost unlimited number of switch contacts to be simultaneously actuated by a single actuating mechanism. Different actuator types such as pushbuttons, rotary selectors, knobs, and keyswitches may also be interchanged with contact modules for maximum flexibility:



9.3 Limit switches

Limit switch symbols




A *limit switch* detects the physical motion of an object by direct contact with that object. An example of a limit switch is the switch detecting the open position of an automobile door, automatically energizing the cabin light when the door opens.

Recall from section 9.1 that the “normal” status of a switch is the resting condition of *no stimulation*. A limit switch will be in its “normal” status when it is not in contact with anything (i.e. nothing touching the switch actuator mechanism).


Limit switches find many uses in industry, particular in robotic control and CNC (Computer Numerical Control) machine tool systems. In many motion-control systems, the moving elements have “home” positions where the computer assigns a position value of zero. For example, the axis controls on a CNC machine tool such as a lathe or mill all return to their “home” positions upon start-up, so the computer can know with confidence the starting locations of each piece. These home positions are detected by means of limit switches. The computer commands each servo motor to travel fully in one direction until a limit switch on each axis trips. The position counter for each axis resets to zero as soon as the respective limit switch detects that the home position has been reached.

A typical limit switch design uses a roller-tipped lever to make contact with the moving part. Screw terminals on the switch body provide connection points with the NC and NO contacts inside the switch. Most limit switches of this design share a “common” terminal between the NC and NO contacts like this:



This switch contact arrangement is sometimes referred to as a *form-C* contact set, since it incorporates both a form-A contact (normally-open) as well as a form-B contact (normally-closed).

A close-up view of several limit switches (used on a drum sequencer) shows the arrangement of connection terminals for form-C contacts. Each limit switch has its own “NO” (normally-open), “NC” (normally-closed), and “C” (common) screw terminal for wires to attach:



A limit switch assembly attached to the stem of a rotary valve – used to detect the fully-closed and fully-open positions of the valve – is shown in the following photograph:



9.4 Proximity switches


A *proximity switch* is one detecting the proximity (closeness) of some object. By definition, these switches are *non-contact sensors*, using magnetic, electric, or optical means to sense the proximity of objects.

Recall from section 9.1 that the “normal” status of a switch is the resting condition of *no stimulation*. A proximity switch will be in its “normal” status when it is distant from any detectable object.

Being non-contact in nature, proximity switches are often used instead of direct-contact limit switches for the same purpose of detecting the position of a machine part, with the advantage of never wearing out over time due to repeated physical contact.

Most proximity switches are *active* in design. That is, they incorporate a powered electronic circuit to sense the proximity of an object. *Inductive* proximity switches sense the presence of metallic objects through the use of a high-frequency magnetic field. *Capacitive* proximity switches sense the presence of non-metallic objects through the use of a high-frequency electric field. *Optical* proximity switches detect the interruption of a light beam by an object. *Ultrasonic* proximity switches sense the presence of dense matter by the reflection of sound waves.


The schematic diagram symbol for a proximity switch with mechanical contacts is the same as for a mechanical limit switch, except the switch symbol is enclosed by a diamond shape, indicating a powered (active) device:



Many proximity switches, though, do not provide “dry contact³” outputs. Instead, their output elements are transistors configured either to *source* current or *sink* current. The terms “sourcing” and “sinking” are best understood by visualizing electric current in the direction of *conventional flow* rather than *electron flow*.

³This curious label is used to describe switch contacts lacking their own built-in power source, as opposed to a switch contained in a device that also provides power to drive the switch circuit. Dry contacts may be mechanical in nature, or they may be electronic (e.g. transistor). By contrast, a “wet” contact is one already connected to an internal power source, ready to drive a load with no other external components needed.

The following schematic diagrams contrast the two modes of switch operation, using red arrows to show the direction of current (conventional flow notation). In both examples, the load being driven by each proximity switch is a light-emitting diode (LED):





A common coloring convention for electronic proximity switches is brown for +V power supply, blue for ground (– pole of power supply), and black for the switched output signal. This convention is common to sinking and sourcing proximity switches alike.

An electronic switch designed to sink current through its signal wire is alternatively referred to as an *NPN* switch due to the type of transistor used in its output. Conversely, an electronic switch designed to source current through its signal wire may be referred to as a *PNP* switch. The key to understanding these labels is to recognize the emitter terminal of the output transistor is always the one connected to the power supply rail. For a sinking switch, this means the emitter must connect to the negative rail, necessitating⁴ an NPN transistor to do the switching. For a sourcing switch,

⁴To be honest, one *could* use an NPN transistor to source current or a PNP to sink, but it would require the

this means the emitter must connect to the positive rail, in which case only a PNP transistor will suffice.

Yet another convention for describing sourcing and sinking transistor switches is to refer to them as *high-side* switches and *low-side* switches, respectively. A sourcing transistor (PNP) has its emitter terminal attached directly to the “high” rail (+) of the DC power supply. A sinking transistor (NPN), by contrast, has its emitter terminal attached directly to the “low” rail (–) of the DC power supply.



transistor be used in the *common-collector* configuration which does not allow for saturation. The engineers designing these proximity switches strive for complete saturation of the transistor, in order to achieve minimum “on” resistance, and that requires a common-emitter configuration.

These photographs show two different styles of electronic proximity switch:



Many industrial proximity switches have built-in LED indicator lamps to help technicians diagnose circuit problems by directly indicating switch status. With just a glance, one may tell whether or not the proximity switch is detecting the presence of an object.

The next photograph shows a proximity switch detecting the passing of teeth on a chain sprocket, generating a slow square-wave electrical signal as the sprocket rotates. Such a switch may be used as a rotational speed sensor (sprocket speed proportional to signal frequency) or as a broken chain sensor (when sensing the rotation of the driven sprocket instead of the drive sprocket):



Like other types of process switches, proximity switches come in both “normally open” (NO) and “normally closed” (NC) varieties. This distinction has nothing whatsoever to do with sourcing versus sinking (PNP vs. NPN), but rather what the status of the proximity switch will be when no objects are near. Thus, it is possible to find normally-open proximity switches that are sinking (NPN) as well as normally-open proximity switches that are sourcing (PNP), and normally-closed proximity switches in either sinking or sourcing designs as well.


These switch characteristics are commonly fixed, and must be specified when ordering the device. Likewise, the detection range of a proximity switch is usually a fixed parameter rather than being adjustable.

9.5 Pressure switches

A *pressure switch* is one detecting the presence of fluid pressure. Pressure switches often use diaphragms or bellows as the pressure-sensing element, the motion of which actuates one or more switch contacts.

Recall from section 9.1 that the “normal” status of a switch is the resting condition of *no stimulation*. A pressure switch will be in its “normal” status when it senses minimum pressure (e.g. an applied pressure, or in some cases a vacuum condition)⁵. For a pressure switch, “normal” status is any fluid pressure *below* the trip threshold of the switch.

Pressure switch symbols



The following photograph shows two pressure switches sensing the same fluid pressure as an electronic pressure transmitter (the device on the far left) because they are all plumbed to a common tube:




⁵If the trip setting of a pressure switch is below atmospheric pressure, then it will be “actuated” at atmospheric pressure and in its “normal” status only when the pressure falls below that trip point (i.e. a vacuum).

A legacy design of pressure switch uses a bourdon tube as the pressure-sensing element, and a glass bulb partially filled with mercury as the electrical switching element. When applied pressure causes the bourdon tube to flex sufficiently, the glass bulb tilts far enough to cause the mercury to fall against a pair of electrodes, thus completing an electrical circuit. A great many pressure switches of this design were sold under the brand name of “Mercoid,” with a few appearing in this photograph of a steam boiler (the round-shaped units with glass covers allowing inspection of the bourdon tube and mercury tilt switch):



A close-up photograph of one of these pressure switches appears here. The bourdon tube is grey in color, and almost as wide in diameter as the circular switch housing. The mercury tilt switch bottles have yellow-colored plastic caps covering up their external electrical contacts:



The next set of photographs show a mercury tilt switch removed from the pressure switch mechanism, so you may see the switch in two different states (contact open on the left, and closed on the right):



Advantages of mercury tilt switches include immunity to switch contact degradation from harmful atmospheres (oil mist, dirt, dust, corrosion) as well as safety in explosive atmospheres (since a spark

contained within a hermetically sealed glass bulb cannot touch off an explosion in the surrounding atmosphere). Disadvantages include the possibility of intermittent electrical contact resulting from mechanical vibration, as well as sensitivity to mounting angle (i.e. you would *not* want to use this kind of switch aboard a moving vehicle!).

A pressure switch manufactured by the Danfoss corporation appears in the next photograph. This particular model of pressure switch has windows on the front cover allowing a technician to see the pressure limit setting inside:



This switch balances the force generated by a pressure-sensing element against a mechanical spring. Tension on the spring may be adjusted by a technician, which means the trip point of this switch is adjustable.

One of the settings on this switch is the *deadband* or *differential* pressure setting, seen in the lower window. This setting determines the amount of pressure change required to re-set the switch to its normal state after it has tripped. For example, a high-pressure switch with a trip point of 67 PSI (changes state at 67 PSI, increasing) that re-sets back to its normal state at a pressure of 63 PSI decreasing has a “deadband” or “differential” pressure setting of 4 PSI (67 PSI – 63 PSI = 4 PSI).

The “differential” pressure setting of a gauge pressure switch is not to be confused with a true *differential pressure* switch. In the next photograph, we see a pressure switch truly actuated by *differential* pressure (the difference in fluid pressure sensed between two ports):



The electrical switch element is located underneath the blue cover, while the diaphragm pressure element is located within the grey metal housing. The net force exerted on the diaphragm by the two fluid pressures varies in magnitude and direction with the magnitude of those pressures. If the two fluid pressures are precisely equal, the diaphragm experiences no net force (zero differential pressure).



Like the Danfoss gauge pressure switch seen previously, this differential pressure switch has a “trip” or “limit” setting as well as a “dead-band” or “differential” setting. It is important to recognize and clearly distinguish the two meanings of *differential pressure* in the context of this device. It senses differences in pressure between two input ports (“differential pressure” – the difference between two different fluid pressure connections), but being a switch, it also exhibits some dead band in its action (“differential pressure” – a change in pressure required to re-set the switch’s state).

9.6 Level switches

A *level switch* is one detecting the level of liquid or solid (granules or powder) in a vessel. Level switches often use floats as the level-sensing element, the motion of which actuates one or more switch contacts.

Recall from section 9.1 that the “normal” status of a switch is the resting condition of *no stimulation*. A level switch will be in its “normal” status when it senses minimum level (e.g. an empty vessel). For a level switch, “normal” status is any fluid level *below* the trip threshold of the switch.



Level switch symbols



9.6.1 Float-type level switches

Some level switches use a *float* to sense the level of a liquid surface, actuating an electrical switch by the motion of the float. The electrical schematic symbol for a level switch is actually based on this type of mechanism, with a round “ball” float drawn as the actuating element:


Level switch symbols



An example of this technology is a level switch manufactured by Magnetrol, with two such switches shown in the following photograph of a steam boiler. These switches sense water level in the steam drum of the boiler:




The Magnetrol float switch mechanism uses a mercury tilt bulb, tilted by a magnet's attraction to a steel rod lifted into position by a float. The float directly senses liquid level, positioning the steel rod either closer to or farther away from the magnet. If the rod comes close enough to the magnet, the mercury bottle tilts to change the switch's electrical status:



A feature of this design is complete isolation between the electrical components and the “wet” components sensing the liquid level. The steel rod moves inside a non-magnetic metal tube, with the tube sealing process fluid pressure from escape while still allowing the magnetic tilt switch to sense float position.

Simpler float switch designs also exist for direct installation in open (vented) process vessels, resembling the float valve assembly on a water-flush toilet. Any “limit” style switching element will work here, including inductive proximity switches, to sense the float's position in an environment where no isolation need exist between the switch and the process fluid(s):



9.6.2 Tuning fork level switches

This level switch uses a metal *tuning fork* structure to detect the presence of a liquid or solid (powder or granules) in a vessel:



An electronic circuit continuously excites the tuning fork, causing it to mechanically vibrate. When the prongs of the fork contact anything with substantial mass, the resonant frequency of the fork decreases. The circuit detects this frequency change and indicates the presence of mass contacting the fork. The forks' vibrating motion tends to shake off any accumulated material, such that this style of level switch tends to be resistant to fouling.

It should be noted that the previous photograph of the tuning-fork style level switch is complete: the fork "paddles" are only a couple of inches long and require no physical extensions in order to properly detect liquid or solid material at that point.

9.6.3 Paddle-wheel level switches

A more primitive variation on the theme of a “tuning fork” level switch is the *rotating paddle* switch, used to detect the level of powder or granular solid material. This level switch uses an electric motor to slowly rotate a metal paddle inside the process vessel. If solid material rises to the level of the paddle, the material’s bulk will place a mechanical load on the paddle. A torque-sensitive switch mechanically linked to the motor actuates when enough torsional effort is detected on the part of the motor. A great many level switches of this design sold in the United States under the trade-name *Bindicator* (so-called because they detected the level of solid material in storage *bins*).

A “Bindicator” style of level switch appears in this photograph (painted black, against a white-painted hopper), used to detect the presence of soda ash powder in a hopper at a water treatment plant:



9.6.4 Ultrasonic level switches

Yet another style of electronic level switch uses ultrasonic sound waves to detect the presence of process material (either solid or liquid) at one point:



Sound waves pass back and forth within the gap of the probe, sent and received by piezoelectric transducers. The presence of any substance other than gas within that gap affects the received audio power, thus signaling to the electronic circuit within the bulkier portion of the device that process level has reached the detection point. The lack of moving parts makes this probe quite reliable, although it may become “fooled” by heavy fouling.

9.6.5 Capacitive level switches


Another electronic liquid level switch technology is *capacitive*: sensing level by changes in electrical capacitance between the switch and the liquid. The following photograph shows a couple of capacitive switches sensing the presence of water in a plastic storage vessel:



9.6.6 Conductive level switches

Perhaps the simplest (and oldest) form of electrical level detection is where a pair of metal electrodes contacts the process material to form a complete electrical circuit, actuating a relay. This type of switch, of course, only works with granular solids and liquids that are electrically conductive (e.g. potable or dirty water, acids, caustics, food liquids, coal, metal powders) and not with non-conducting materials (e.g. ultra-pure water, oils, ceramic powders).

A legacy design for conductive level switches is the model 1500 “induction relay” originally manufactured by B/W Controls, using a special transformer/relay to generate an isolated AC probe voltage and sense the presence of liquid:




Line voltage (120 VAC) energizes the primary coil, sending a magnetic field through the laminated ferrous⁶ core of the relay. This magnetic field easily passes through the center of the secondary coil when the secondary circuit is open (no liquid closing the probe circuit), thus completing the magnetic “circuit” in the core. With the magnetic circuit thus completed, the armature will not be attracted to the core. However, when a circuit is completed by liquid level rising to contact both probes, the secondary coil’s resulting current “bucks” the magnetic flux⁷ through its center, causing more

⁶ “Ferrous” simply means any iron-containing substance.

⁷ The reason for this opposition is rooted in the roles of primary and secondary coils as power *load* and *source*, respectively. The voltage across each coil is a function of Faraday’s Law of Electromagnetic Induction: $V = N \frac{d\phi}{dt}$. However, since the primary coil acts as a load (drawing power from the 120 VAC source) and the secondary coil acts as a source (sending power to the probes), the directions of current through the two coils will be opposite despite their common voltage polarities. The secondary coil’s opposite current direction causes an opposing magnetic force in that section of the core, reducing the magnetic flux there. In a normal power transformer, this reduction in magnetic

magnetic flux to bypass to the end poles where it attracts the ferrous armature toward the core frame. This physical attraction actuates switch contacts which then signal the presence of liquid level at the probes.

The following pair of illustrations shows the two conditions of this level switch, with the magnetic lines of flux highlighted as dashed lines through the core:



The “transformer” design of this particular conductive level switch not only provides electrical isolation between the probes and the energizing (120 VAC) circuit, but it also enables a wide range of detection voltages to be generated for the probes just by altering the number of wire “turns” in the secondary coil. The B/W Controls model 1500 inductive relay is available in a variety of AC voltage output levels, ranging from 12 VAC (for detecting metallic substances) to 800 VAC for use with demineralized water such as that found in steam boiler systems.

More modern variations on the same design theme use much lower AC voltages⁸ to energize the probes, employing sensitive semiconductor amplifier circuits to detect probe current and signal liquid level.

flux caused by secondary current is also felt by the primary coil (since there is only one magnetic “path” in a power transformer’s core), which then causes the primary coil to draw more current and re-establish the core flux at its original magnitude. With the inductive relay, however, the opposing magnetic force created by the secondary coil simply forces more of the primary coil’s magnetic flux to bypass to the alternate route: through the armature.

⁸The B/W Controls model 5200 solid-state relay, for example, uses only 8 volts AC at the probe tips.

9.7 Temperature switches


A *temperature switch* is one detecting the temperature of some substance. Temperature switches often use bimetallic strips as the temperature-sensing element, the motion of which actuates one or more switch contacts. An alternative design uses a metal bulb filled with a fluid that expands with temperature, causing the switch mechanism to actuate based on the pressure this fluid exerts against a diaphragm or bellows. This latter temperature switch design is really a pressure switch, whose pressure is a direct function of process temperature by virtue of the physics of the entrapped fluid inside the sensing bulb.

Recall from section 9.1 that the “normal” status of a switch is the resting condition of *no stimulation*. A temperature switch will be in its “normal” status when it senses minimum temperature (i.e. cold, in some cases a condition colder than ambient)⁹. For a temperature switch, “normal” status is any sensed temperature *below* the trip threshold of the switch.

Temperature switch symbols



Normally-open
(NO)



Normally-closed
(NC)

⁹If the trip setting of a temperature switch is below ambient temperature, then it will be “actuated” at ambient temperature and in its “normal” status only when the temperature falls below that trip point (i.e. colder than ambient).

The following photograph shows a temperature-actuated switch manufactured by the Ashcroft corporation:



Like all other process switches, temperature switches exhibit a certain amount of *deadband* in their switching action. A temperature switch that trips at 300 °F *rising*, for example, will not re-set at 300 °F *falling*. That switch would more likely reset at some lower temperature such as 295 °F. With mechanical switch designs, some amount of deadband is inevitable due to friction inside the mechanism. However, process switch deadband is actually a useful characteristic as it helps avoid repeated “nuisance” alarms from happening.

To understand this concept, it is helpful to imagine a scenario where the process variable is at or very near the trip point. For our hypothetical temperature switch with a trip point of 300 °F (rising), imagine a situation where the process temperature is precisely 300.0 °F. Any further rise in temperature will of course trip the switch (sounding an alarm). With no deadband, however, the switch will immediately re-set when the temperature falls back down to 300.0 °F. This means the switch may possibly “cycle” back and forth between its trip and reset states with just a minute change in process temperature (300.0 °F to 300.1 °F and back again). If the temperature switch is activating an alarm every time it trips, it will create a series of alarm events prompting operators to repeatedly acknowledge the alarm. This is a nuisance to operations personnel, as it distracts them

from addressing what they already realize is a process problem. It is better for the switch to trip at 300.0 °F rising *and remain in that tripped state* until the temperature falls down to some degree substantially below the trip point. This way, the operators only receive *one* alarm event rather than multiple alarm events for each process temperature excursion.

Some mechanical temperature switches come equipped with a separate adjustment for deadband (also called *differential*). Setting this deadband adjustment in a mechanical temperature switch requires the technician to repeatedly subject the sensing element to a rising and falling temperature, to check that the switch trips at the proper setting and resets at the proper setting. This is analogous to cycling the process variable back and forth when adjusting the “zero” and “span” settings of an analog transmitter: checking to see that the transmitter repeatedly outputs a 0% signal at the lower range value (LRV) and a 100% signal at the upper range value (URV).

For discrete temperature-sensing applications demanding high accuracy and repeatability, *electronic* temperature switch circuits using thermocouples, RTDs, or thermistors may be used instead of a mechanical (bi-metallic or filled bulb) sensing element. The operation and configuration of discrete electronic temperature switches is very similar to that of continuous electronic temperature transmitters.

An example of an electronic temperature switch module is the Moore Industries model SPA (“Site Programmable Alarm”), shown here:



Not only is this particular model capable of directly interpreting both RTD and thermocouple signals, but it will also input 4-20 mA loop current signals as well. This means you may build a temperature switch system out of a 4-20 mA loop-powered temperature transmitter (located in the field) and an SPA switch module (located in a protected environment such as a control room or electrical enclosure). Many other manufacturers and models of electronic switching units exist, with the Moore Industries SPA being just one example.

With electronic temperature switches, the adjustment of deadband (differential) is both precise and flexible. Unlike mechanical switches where deadband is primarily a function of friction, and therefore liable to change over time as the device wears, electronic switching circuits may be precisely set for any trip and reset points along its measurement range, remaining very stable over time.

9.8 Flow switches

A *flow switch* is one detecting the flow of some fluid through a pipe. Flow switches often use “paddles” as the flow-sensing element, the motion of which actuates one or more switch contacts.

Recall from section 9.1 that the “normal” status of a switch is the resting condition of *no stimulation*. A flow switch will be in its “normal” status when it senses minimum flow (i.e. no fluid moving through the pipe). For a flow switch, “normal” status is any fluid flow rate *below* the trip threshold of the switch.

Flow switch symbols



Normally-open
(NO)



Normally-closed
(NC)

A simple paddle placed in the midst of a fluid stream generates a mechanical force which may be used to actuate a switch mechanism, as shown in the following photograph:



Like all other process switches, flow switches exhibit *deadband* (also called *differential*) in their switching action. A flow switch that trips at 15 GPM *rising*, for example, will not re-set at 15 GPM *falling*. That switch would more likely reset at some lower flow rate such as 14 GPM. With the “trip” and “reset” points being different, the switch avoids unnecessary cycling if the process flow rate hovers near one critical value.

9.9 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **“Normal” switch status:** the “normal” status of a switch contact as defined by the manufacturer is its *resting* condition (minimum stimulus).
- **Sourcing versus sinking:** whether the electronic device in question is outputting (conventional flow) current or inputting current. Relevant to the proper connection of electronic switches (especially proximity switches).
- **Deadband and hysteresis:** the difference in response with the independent variable increasing versus decreasing. Usually caused by friction in a mechanism. Relevant to the “trip” settings of process switches: the value at which a switch changes state when its stimulus increases is not the same value it changes state when its stimulus decreases.

References

“Installation and Service Manual – 1500 Induction Control Relays”, document 511 1500.M4R 02/05.Z145 10M, AMETEK Automation and Process Technologies, Clawson, MI, 2005.

“Installation and Service Manual – 5200 Solid State Relays”, document 432 5200.M4R 05/07.Z152, AMETEK Automation and Process Technologies, Clawson, MI, 2007.

Chapter 10


Discrete control elements

The word “discrete” means *individual* or *distinct*. In engineering, a “discrete” variable or measurement refers to a true-or-false condition. Thus, a discrete control element is one that has but a limited number of states (usually two: on and off). In the case of valves, this means a valve designed to operate either in “open” mode or “closed” mode, not in-between.

10.1 On/off valves

An on/off valve is the fluid equivalent of an electrical switch: a device that either allows unimpeded flow or acts to prevent flow altogether. These valves are often used for routing process fluid to different locations, starting and stopping batch processes, and engaging automated safety (shutdown) functions.

Valve styles commonly used for on/off service include ball, plug, butterfly (or disk), gate, and globe. Large on/off valves are generally of such a design that the full-open position provides a nearly unimpeded path for fluid to travel through. Ball, plug¹, and gate valves provide just this characteristic:



¹A *plug* valve is very much like a ball valve, the difference being the shape of the rotating element. Rather than a spherical ball, the plug valve uses a truncated cone as the rotary element, a slot cut through the cone serving as the passageway for fluid. The conical shape of a plug valve's rotating element allows it to wedge tightly into the "closed" (shut) position for exceptional sealing.

A series of photographs showing a cut-away ball valve (hand-actuated) in three different positions reveals the inner workings common to all ball valve mechanisms:




The left-hand image shows the valve in the shut position, with the bore axis facing the viewer (preventing fluid flow). The right-hand image shows the valve in the open position, with the bore axis perpendicular to view and allowing flow. The middle image shows the valve in a partially-open condition.

10.2 Fluid power systems

Given the ability of pressurized fluids to transmit force over long distances, it is not surprising that many practical “fluid power systems” have been built using fluid as a mechanical power-conducting media. Fluid systems may be broadly grouped into *pneumatic* (gas, usually air) and *hydraulic* (liquid, usually oil²).


Although there is no particular reason why a fluid power system must be discrete and not continuous, the majority of fluid power systems operate in an on/off control mode rather than throttling, which is why this subject is covered in the “Discrete Control Elements” chapter.

As usual for technical specialties, fluid power has its own unique symbology for describing various components and their interconnections. The following diagram shows some common symbols used in fluid power system diagrams. Lines connecting components together in a fluid power diagram indicate pipes, hoses, or tubes, much like lines connecting components together in an electronic schematic diagram represent wires:




²While it would be technically possible to use water instead of oil in a hydraulic power system, oil enjoys some distinct advantages. First, oil is a lubricating substance, and non-corrosive, unlike water. Second, oil enjoys a wider operating temperature range than water, which tends to both freeze and boil more readily.

Many of these symbols are self-explanatory, especially the pumps, motors, and cylinders. What seems to cause the most confusion for people new to this symbology are the spool valve symbols. A “spool” valve is a special type of flow-directing valve used in pneumatic and hydraulic systems to direct the pressurized fluid to different locations. The symbology for a spool valve is a set of boxes, each box containing arrows or other symbols showing the intended direction(s) for the fluid’s travel. Take for instance this pneumatic reversing cylinder control system:




The proper way to interpret a spool valve symbol is to see only one “box” active at any given time. As the actuator (in this case, a hand-actuated lever) is moved one way or the other, the boxes “shift” laterally to redirect the flow of fluid from source to load.


For example, when the spool valve in this reversing control system is in its center position, the outer boxes in the symbol are inactive. This is emphasized in the following diagram by coloring the outer boxes grey. In this position, the spool valve neither admits compressed air to the cylinder nor vents any air from the cylinder. As a result, the piston within the cylinder holds its position:



If the spool valve is actuated in one direction, the spool piece inside the valve assembly shifts, directing compressed air to one side of the cylinder while venting air from the other side. This is shown in the following diagram by shifting the boxes to one side, lining up the “active” box with the cylinder and air supply/vent connections:




If the spool valve is actuated in the other direction, the spool piece inside the valve assembly shifts again, switching the directions of air flow to and from the cylinder. Compressed air still flows from the supply to the vent, but the direction within the cylinder is reversed. This causes the piston to reverse its mechanical travel:



Note that the boxes in a spool valve symbol are never shifted or grayed-out in color like this to represent the valve’s state in a real fluid power diagram. The previous illustrations were drawn this way only as an aid to your understanding, teaching you how to interpret the meaning of the symbols when you see them in real fluid power diagrams. Like electrical switches represented in schematic diagrams, spool valve symbols are always drawn with the boxes aligned in their “resting” states, and with all portions identically colored.

Hydraulic systems require more components, including filters and pressure regulators, to ensure proper operation. Shown here is a simple uni-directional hydraulic motor control system:




Note the placement of the pressure relief valve: it is a *shunt* regulator, bleeding excess pressure from the discharge of the hydraulic pump back to the reservoir³. A “shunt” regulator is necessary because hydraulic pumps are *positive displacement*, meaning they discharge a fixed volume of fluid with every revolution of the shaft. If the discharge of a positive-displacement pump is blocked (as it would be if the spool valve were placed in its default “off” position, with no shunt regulator to bleed pressure back to the reservoir), it will mechanically “lock” and refuse to turn. This would overload the electric motor coupled to the pump, if not for the pressure regulating valve providing an alternative route for oil to flow back to the reservoir. This shunt regulator allows the pump to discharge a fixed rate of oil flow (for a constant electric motor speed) under all hydraulic operating conditions.

³Note also how identical reservoir symbols may be placed at different locations of the diagram although they represent the exact same reservoir. This is analogous to “ground” symbols in electronic schematic diagrams, every ground symbol representing a common connection to the same zero-potential point.

An alternative to using a shunt regulating valve in a hydraulic system is to use a *variable-displacement pump*. Variable-displacement pumps still output a certain volume of hydraulic oil per shaft revolution, but that volumetric quantity may be varied by moving a component within the pump. In other words, the pump's per-revolution displacement of oil may be externally adjusted.


If we connect the variable-displacement mechanism of such a hydraulic pump to a pressure-sensing element such as a bellows, in a way where the pump senses its own discharge pressure and adjusts its volumetric output accordingly, we will have a pressure-regulating hydraulic system that not only prevents the pump from “locking” when the spool valve turns off, but also saves energy by not circulating pressurized oil all the time:



Note the placement of a filter at the inlet of the pump in all hydraulic systems. Filtration is an absolute essential for any hydraulic system, given the extremely tight dimensional tolerances of components inside pumps, motors, spool valves, and cylinders. Even very small concentrations of particulate impurities in the hydraulic fluid may drastically shorten the life of these precision-machined components.

Hydraulic fluid also acts as a heat-transfer medium, and as such must be kept cool enough to prevent thermal damage to components. Large hydraulic systems are equipped with coolers, which are just heat exchangers designed to extract heat energy from the fluid and transfer it to either cooling water or ambient air. Small hydraulic systems dissipate heat at a fast enough rate through their components that coolers are often unnecessary.

An interior view of a simple “2-way” spool valve such as that used in the hydraulic motor system previously examined reveals why cleanliness and temperature stability is important. The spool valve is shown here in both positions, with its accompanying schematic symbol:



Both the spool and the valve body it moves in are circular in cross-section. The spool has wide areas called “lands” that act to cover and uncover ports in the valve body for fluid to flow through. The precise fit between the outside diameter of the lands and the inside diameter of the valve body’s bore is the only factor limiting leakage through this spool valve in the closed state. Dirty hydraulic fluid will wear at this precise fit over time until the valve is no longer capable of sealing fluid in its “closed” position. Extreme cycles in temperature will also compromise the precise fit between the spool and the valve body.

Pneumatic fluid power systems require cleanliness as well, since any particulate contamination in the air will likewise cause undue wear in the close-tolerance compressors, motors, valves, and cylinders. Unlike hydraulic oil, compressed air is not a natural lubricant, which means many pneumatic power devices benefit from a small concentration of oil vapor in the air. Pneumatic “oilers” designed to introduce lubricating oil into a flowing air stream are generally located very near the point of use (e.g. the motor or the cylinder) to ensure the oil does not condense and “settle” in the air piping.

Fluid power systems in general tend to be inefficient, requiring much more energy input to the fluid than what is extracted at the points of use⁴. When large amounts of energy need to be transmitted over long distances, electricity is the a more practical medium for the task. However, fluid power systems enjoy certain advantages over electric power, a few of which are listed here:

- Fluid power motors and cylinders do not overload at low speeds or under locked conditions
- Fluid power systems present little hazard of accidentally igniting flammable atmospheres (i.e. no sparks produced)
- Fluid power systems present little or no fire hazard⁵
- Fluid power systems present no hazard of electric shock or arc flash
- Fluid power systems are often easier to understand and troubleshoot than electric systems
- Fluid power systems may be safely used in submerged (underwater) environments
- Pneumatic systems are relatively easy to equip with back-up energy reserve (e.g. liquefied nitrogen serving as a back-up gas supply in the event of compressor shut-down)
- Pneumatic systems are self-purging (i.e. enclosures housing pneumatic devices will be naturally purged of dusts and vapors by leaking air)

Another important consideration for fluid power systems is the ongoing maintenance work they require for reliable operation. Hydraulic power systems will suffer rapid wear if the hydraulic oil is not clean and chemically stable. The fluid in a hydraulic system not only transmits mechanical power, but it also lubricates and stabilizes the temperature of components as they transfer that power between different forms. Regular filter changes and oil changes (especially if the fluid is subject to contamination from the process) is necessary for long service life of any hydraulic system.

Pneumatic (instrument air) systems must be free of water vapor and particulate contamination for much the same reason. Water is perhaps the most common contaminant in instrument air systems, causing corrosion of metal components and subsequent clogging of orifices. Special devices called *air dryers* installed in instrument air systems use solid materials called *desiccants* to absorb water entrained in the compressed air. The desiccant material is “regenerated” by the dryer mechanism on a regular cycle, but must be periodically replaced when its water-absorbing ability wanes.

⁴Close-coupled hydraulic systems with variable-displacement pumps and/or motors may achieve high efficiency, but they are the exception rather than the rule. One such system I have seen was used to couple a diesel engine to the drive axle of a large commercial truck, using a variable-displacement pump as a continuously-variable transmission to keep the diesel engine in its optimum speed range. The system was so efficient, it did not require a cooler for the hydraulic oil!

⁵Many kinds of hydraulic oils are flammable, so this is not a perfectly true statement. However, fire-resistant fluids such as *Skydrol* (introduced to the aviation industry for safety) are commercially available.

This next photograph shows a high-capacity industrial air dryer, with two large chambers holding desiccant:



A valving system directs the main flow of compressed air through one of these desiccant chambers at a time, allowing the desiccant to absorb water vapor in the air. Meanwhile, the unused chamber is purged of its collected water by venting low-pressure air through it to the atmosphere. An electronic timer unit (or PLC) controls the cycling of this valve system to ensure adequate drying and maximized desiccant service life.

Moisture content in instrument air is often expressed by the term *dew point*. This is the temperature at which water vapor suspended in the instrument air will condense into water droplets, at atmospheric pressure. The “drier” the air, the lower the dew point temperature; the “wetter” the air, the higher the dew point temperature. Sometimes the “dryness” of instrument air is expressed in terms of *pressure dew point* (PDP), which is the temperature of water condensation at system pressure rather than at atmospheric pressure. Pressure dew point is always a higher temperature value than atmospheric dew point, since greater air pressures force condensation to occur more readily. Pressure dew point is a more practical value than atmospheric dew point for an instrument air system, as PDP directly indicates the ambient temperature at which water will condense in an *operating* pneumatic system. A low dew point value means that the air dryer is working as it should. A high dew point value means condensation is more likely to form in the compressed air system piping.


A simple way to help extract water from an instrument air system is an accessory called a *water trap*, usually found on air pressure regulators. The following photograph shows a Fisher pneumatic regulator equipped with such a trap on the bottom:



A shiny metal “wingnut” drain appears at the very bottom of the regulator, acting as a manual valve for purging collected water from the basin where compressed air enters the regulator mechanism. Periodic opening of this drain valve by maintenance or operations personnel allows collected water to be blown out of the regulator.

Another way to help minimize the amount of water reaching pneumatic devices is to properly orient all connections to the main air pipe (called a *header*). Ideally, each instrument air tap coming off a header should do so on the *top* of the header, not the bottom. This way, collected condensation inside the header will not go directly to the points of use, but rather will drain downhill to the lowest point in the header where a drain valve may be placed.

This next photograph shows an *incorrect* installation, where air is drawn off the bottom of the main header line:

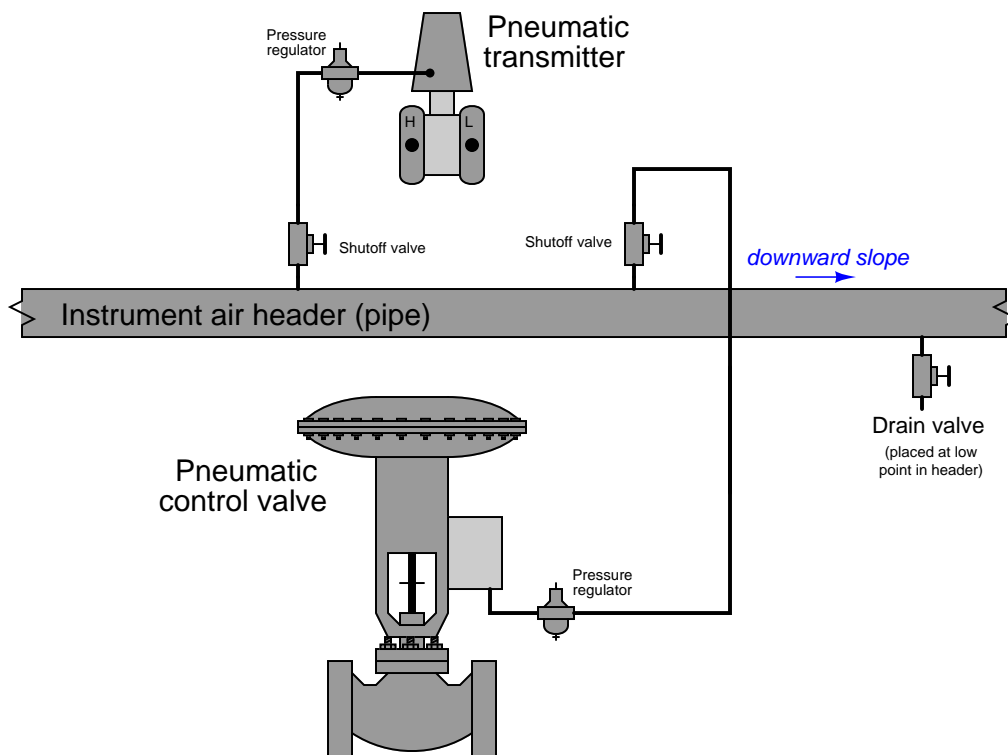


Such an installation invites trouble, as every bit of water condensed inside the header pipe is *guaranteed* to find its way by gravity to the instruments connected to the underside of that header.

One good feature of this installation is the use of stainless steel as the piping material. Copper, brass, plastic⁶, and stainless steel are the preferred materials for instrument air piping, tubing, valves, and fittings, as standard (iron) pipe will inevitably rust in the presence of condensation. Particles of rust created inside an instrument air system plays havoc with the tiny ports, nozzles, and orifices of pneumatic instruments.

⁶Certain types of plastic pipe such as PVC should never be used in compressed air systems because it becomes brittle and liable to fracture over time. If you are considering the use of plastic for a high-pressure compressed air system, be sure the type of plastic is engineered for air pressure service!

The proper way to make instrument air connections to the air header is as such:



In order to facilitate draining of the header, the header should be slightly inclined, with the drain valve installed at the lowest point. This drain valve should then be periodically opened on a regular maintenance schedule in order to prevent the header from slowly filling up with condensed water over time.

10.3 Solenoid valves

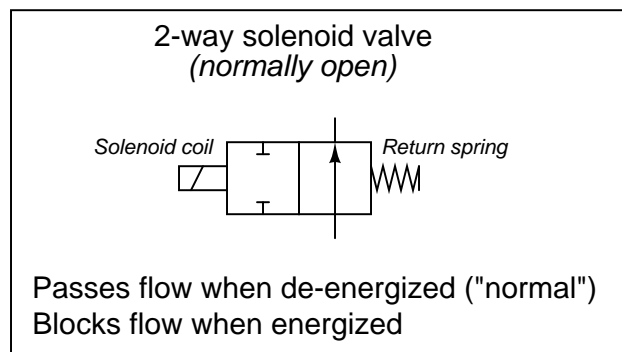
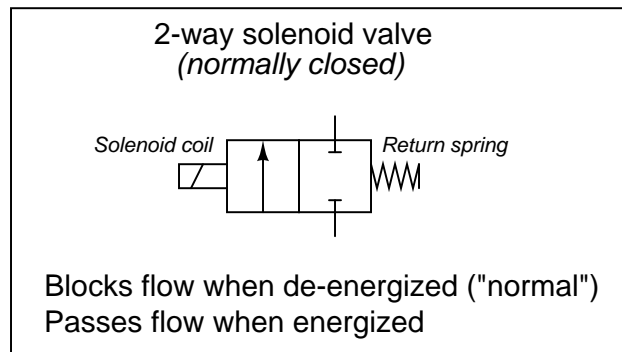
A very common form of on/off valve used for pneumatic and hydraulic systems alike is the *solenoid valve*. A "solenoid" is nothing more than a coil of wire designed to produce a magnetic field when energized. Solenoid actuators work by attracting a movable ferrous *armature* into the center of the solenoid coil when energized, the force of this attraction working to actuate a small valve mechanism.

Solenoid-actuated valves are usually classified according to the number of ports ("ways"). A simple on/off solenoid valve controlling flow into one port and out of another port is called a *2-way* valve. Another style of solenoid valve, where flow is directed in one path or to another path – much like a single-pole double-throw (SPDT) electrical switch – is called a *3-way* valve because it has three fluid ports.


10.3.1 2-way solenoid valves

2-way solenoid valves operate in a manner analogous to single-pole single-throw (SPST) electrical switches: with only one path for flow.

Solenoid valve symbols often appear identical to fluid power valve symbols, with “boxes” representing flow paths and directions between ports in each of the valve’s states. Like electrical switches, these valve symbols are always drawn in their “normal” (resting) state, where the return spring’s action determines the valve position:



A good way to make sense of these “box” valve symbols is to imagine the boxes sliding back and forth as the actuating elements work. For example, the two boxes in a normally-closed solenoid valve symbol may be thought of in terms of being *pushed* to the left by the spring when de-energized and *pushed* to the right by the solenoid’s force when energized. Here, the color grey de-emphasizes the unselected box in each of the valve’s two states:




As with electrical switches in schematic diagrams, fluid control valve symbols are always drawn in their “normal” (resting) states. For example, a normally-closed valve will always be drawn so that the box with the blocked ports aligns with the tubes leading to and from the valve. What you see in the above illustration are “dramatized” symbols, highlighting the valve’s action by color and by re-positioning the boxes, strictly for the purpose of making it easier for you to grasp the concept. This sort of coloring and re-positioning is *never* shown in a real schematic diagram. In a fluid control schematic, it is left to the reader to visualize the valve symbol boxes moving to and fro, determining the flow path of fluid through the valve.

Unlike electrical switches, of course, the terms *open* and *closed* have opposite meanings for valves. An “open” electrical switch constitutes a break in the circuit, ensuring no current; an “open” valve, by contrast, freely allows fluid flow through it. A “closed” electrical switch has continuity, allowing current through it; a “closed” valve, on the other hand, shuts off fluid flow.

The arrow inside a solenoid valve symbol actually denotes a preferred direction of flow. Most solenoid valves use a “globe” or “poppet” style of valve element, where a metal plug covers up a hole (called the “seat”). Process fluid pressure should be applied to the valve in such a way that the pressure difference tends to hold the solenoid valve in its “normal” position (the same position as driven by the return spring). Otherwise⁷, enough fluid pressure might override the return spring’s action, preventing the valve from achieving its “normal” state when de-energized. Thus, we see that the label “2-way” does not refer to two directions of flow as one might assume, but rather two *ports* on the valve for fluid to travel through.

Some solenoid valves are designed in such a way that the direction of fluid flow through them is irrelevant. In such valves, the arrow symbols will be double-headed (one head at each end, pointing in opposite directions) to show the possibility of flow in either direction.

Bidirectional solenoid valve

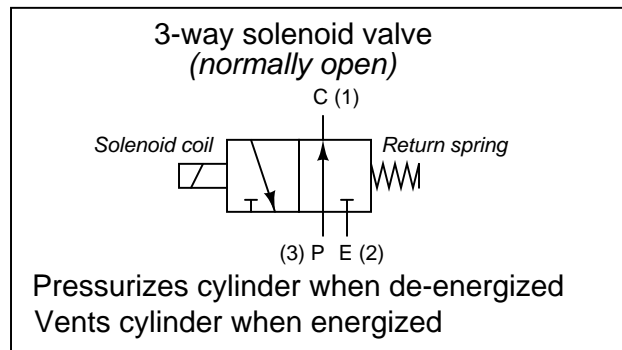



⁷One could argue that enough fluid pressure could override the solenoid’s energized state as well, so why choose to have the fluid pressure act in the direction of helping the return spring? The answer to this (very good) question is that the solenoid’s energized force greatly exceeds that of the return spring. This is immediately obvious on first inspection, as the solenoid *must* be stronger than the return spring or else the solenoid valve would never actuate! Furthermore, the solenoid’s force must be *significantly* stronger than the spring, or else the valve would open rather slowly. Fast valve action demands a solenoid force that greatly exceeds spring force. Realizing this, now, we see that the spring is the weaker of the two forces, and thus it makes perfect sense why we should use the valve in such a way that the process pressure helps the spring: the solenoid’s force has the best chance of overcoming the force on the plug produced by process pressure, so those two forces should be placed in opposition, while the return spring’s force should work *with* (not against) the process pressure.

10.3.2 3-way solenoid valves

3-way solenoid valves operate in a manner analogous to single-pole double-throw (SPDT) electrical switches: with two paths for flow sharing one common terminal.

3-way solenoid valves have three ports for fluid, and like 2-way valves may be referred to either as *normally-open* and *normally-closed*. Ports on a pneumatic 3-way valve are commonly labeled with the letters “P,” “E,” and “C,” representing *Pressure* (compressed air supply), *Exhaust* (vent to atmosphere), and *Cylinder* (the actuating mechanism), respectively. Alternatively, you may see the cylinder port labeled “A” (for *actuator*) instead of “C”. If the solenoid valve is intended for use in a hydraulic (liquid) system, the letter “T” is customarily used to identify the return port rather than “E” (i.e. *Tank* rather than *Exhaust*):



The letters used to label ports on a valve such as this not only denote those ports’ destinations, but also serve to mark which “box” of the valve’s symbol is the normal (resting) state. In all fluid power diagrams you will see that only one of the boxes on each spool valve will have lines connecting to it and/or labels at the fluid ports, and that box is the one which will be aligned when the valve is not being actuated.

Alternatively, the numbers 1, 2, and 3 may be used to label the same ports. However, the numbers do not consistently refer to pressure source (P) and exhaust (E) ports, but rather to the 3-way valve's "normal" versus "actuated" statuses. A 3-way valve will pass fluid between ports 1 and 3 in its "normal" (resting) state, and pass fluid between ports 1 and 2 in its energized state. The following table shows the correspondence between port numbers and port letters for both styles of 3-way solenoid valve:


Valve type	Pressure port (P)	Exhaust port (E)	Cylinder port (C)
Normally-closed	2	3	1
Normally-open	3	2	1

Another way to think of this labeling is to consider port 1 the *common*, port 2 the *normally-closed*, and port 3 the *normally-open*, in a manner similar to SPDT (form-C) electrical switches. Again, bear in mind that the words "open" and "closed" do not mean the same for fluid valves as they do for electrical switches. A "normally-open" port on the valve permits fluid flow in its "normal" state, whereas a "normally-open" switch contact *prevents* electric current flow in its "normal" state.


As with 2-way solenoid valves, the arrows denote preferred direction of fluid flow. Bidirectional 3-way valves will be drawn with double-headed arrows (pointing both directions).

A different symbology is used in loop diagrams and P&IDs than that found in fluid power diagrams – one more resembling general instrumentation (ISA) valve symbols:



2-way solenoid valve



3-way solenoid valve




Regrettably, these symbols are not nearly as descriptive as those used in fluid power diagrams. In order to show directions of flow (especially for 3-way valves), one must add arrows showing "normal" (resting, *DE*) flow directions:


3-way solenoid valve
(normally-closed)3-way solenoid valve
(normally-open)

Alternatively, a *pair* of arrows shows the directions of flow in both energized (*E*) and de-energized (*D*) states:

3-way solenoid valve
(normally-closed)



3-way solenoid valve
(normally-open)



Photographs of an actual 3-way solenoid valve (this one manufactured by ASCO) appear here:




A view of the nameplate for this particular solenoid valve reveals some of its ratings and characteristics:




10.3.3 4-way solenoid valves

When a pneumatic actuator requires air pressure applied to two different ports in order to move two different directions (such as the case for cylinders lacking a return spring), the solenoid valve supplying air to that actuator must have four ports: one for air supply (P), one for exhaust (E), and two for the cylinder ports (typically labeled A and B). The following diagram shows a 4-way solenoid valve connected to the piston actuator of a larger (process) ball valve:



The same diagram could be drawn using the “triangle” solenoid valve symbols rather than the “block” symbols more common to fluid power diagrams:









Here, the letters “D” and “E” specify which directions air is allowed to flow when the solenoid is de-energized and energized, respectively.

In both of the examples shown above, the solenoid valve forces the piston-actuated valve stem to move down (shut off) when the solenoid is de-energized. When the solenoid is energized, air is directed to the bottom of the piston (with the top of the piston becoming vented to atmosphere), causing the piston-actuated valve stem to move up (open wide).

An interior view of a standard spool-type 4-way valve of the kind commonly used for directional hydraulic⁸ controls is shown here, along with its accompanying schematic symbol:

Cut-away illustration

Schematic symbol









Note that the actuator (e.g. hand lever, solenoid armature, etc.) has been omitted from this illustration for simplicity. Only the spool and valve body are shown.

⁸In hydraulics, it is common to use the letter “T” to represent the *tank* or *reservoir* return connection rather than the letter “E” for *exhaust*, which is why the supply and vent lines on this valve are labeled “P” and “T”, respectively.

A variation on this theme uses a shorter spool allowing the two control ports to freely pass fluid in the “normal” position:

Cut-away illustration

Schematic symbol




Such a 4-way valve is useful for applications where the final control element (motor, cylinder) must be free to move rather than be locked in place with the valve in the middle position.

If no center “off” position is needed, the lands may be shortened in such a way that they cannot fully cover the “P,” “1,” and “2” ports simultaneously, making the valve useful only in its two extreme positions:


Cut-away illustration

Schematic symbol



Not all 4-way valves use the spool-type design. However, the spool valve enjoys the advantage of having *pressure balance* on its one moving part. If you examine these cut-away illustrations closely, you will see that the two lands present equal surface areas to the two pressures (pump and tank, “P” and “T”) in perfect vertical symmetry, such that any forces acting on the two lands from fluid pressure do so in opposite directions. This means there will be no net hydraulic force acting on the spool to interfere with its positioning, thus making it very easy to position by hand lever, solenoid, piston, etc.

A photograph of a Parker brand 4-way pneumatic solenoid valve appears here:




This particular solenoid valve is spring-centered, with one solenoid coil at each end to provide actuation in two different directions. The middle position is one where all ports are blocked, providing a “locked” control position for the pneumatic actuating element fed air by this solenoid valve.

10.3.4 Normal energization states

Solenoid valves may be used in such a way that they spend most of their time de-energized, energizing only for brief periods of time when some special function is required. Alternatively, solenoids may be maintained in an energized state, and de-energized to perform their design function. The choice to use a solenoid's energized or de-energized state to perform a specific function is left to the system designer, but nevertheless it is important for all maintenance personnel to know in order to perform work on a solenoid-controlled system.

Take the following segment of an actual P&ID for a steam turbine-driven pump control system for example, where a pair of 3-way solenoid valves control instrument air pressure⁹ to a piston-actuated steam valve to start the turbine in the event that an electric motor-driven pump happens to fail:



If *either*¹⁰ of the two solenoid valves de-energizes, instrument air pressure will vent from the top of the piston actuator to atmosphere, causing the steam valve to “fail” to the full-open position and send steam to the turbine. This much is evident from the curved arrows showing air flowing to the “Vent” ports in a de-energized (DE) condition. An additional valve (PY-590) guarantees the piston actuator’s upward motion by simultaneously applying air pressure to the bottom¹¹ of the actuator

⁹The letters “IAS” refer to *instrument air supply*.

¹⁰This solenoid valve arrangement would be designated *1oo2* from the perspective of starting the turbine, since only one out of the two solenoids needs to trip in order to initiate the turbine start-up.

¹¹If you examine this diagram closely, you will notice an error in it: it shows the top and bottom of the piston actuator *connected together* by air tubing, which if implemented in real life would prevent air pressure from imparting any force to the valve stem at all! Connecting the top and bottom of the actuator together would ensure the piston always sees zero differential pressure, and thus would never develop a resultant force. The output tube of PY-590

if ever air is vented from the top. As an additional feature, the left-hand solenoid valve (SOV-590A) has a manual “Reset” lever on it, symbolized by the letter “R” inside a diamond outline.


The only indication of the solenoids’ typical status (energized or de-energized) comes from the letters “NE” next to each solenoid coil. In this case, “NE” stands for *normally energized*. Therefore, this steam turbine control system, which serves as a back-up to an electric motor-driven pump, relies on either (or both) of the solenoid valves *de-energizing* to make the turbine start up. Under “normal” conditions, where the turbine is not needed, the solenoids remain energized and the steam valve remains shut.

Unfortunately, this use of the word “normal” is altogether different from the use of the word “normal” when describing a solenoid valve’s open/close characteristics. Recall that a *normally open* solenoid valve allows fluid to pass through when it is de-energized. A *normally closed* solenoid valve, by contrast, shuts off fluid flow when de-energized. In this context, the word “normally” refers to the *unpowered* state of the solenoid valve. This is quite similar to how the word “normally” is used to describe switch contact status: a normally-open (NO) electrical switch is open when unactuated (at rest); a normally-closed (NC) electrical switch is closed when unactuated (at rest). In both cases, with solenoid valves and with electrical switches, the word “normally” refers to the resting condition of *no stimulation*.

However, when an engineer designs a solenoid control system and declares a solenoid to be “normally energized,” that engineer is describing the *typical* status of the solenoid valve *as it is intended to function in the process*. This may or may not correspond to the *manufacturer’s* definition of “normally,” since the solenoid manufacturer cannot possibly know which state any of their customers intends to have their solenoid valve typically operate in. To illustrate using the previous steam turbine control system P&ID, those two solenoid valves would be considered *normally closed* by the manufacturer, since their de-energized states block air flow from the “P” port to the “C” port and vent air pressure from the “C” port to the “E” (vent) port. However, the engineer who designed this system wanted both solenoids to be energized whenever the turbine was not needed to run (the “normal” state of the process), and so the engineer labeled both solenoid coils as *normally energized*, which means both solenoids will be actuated to pass air pressure from their “P” ports to their “C” ports (and close off the vent ports) under typical conditions. Here, we see the manufacturer’s definition of “normal” for the solenoid valve body (i.e. the *resting* condition) is precisely opposite that of the process engineer’s definition of “normal” for the solenoid coil (i.e. the typical operating status of the process) in this particular application.

should only connect to the bottom of the piston actuator, not to the bottom *and* the top. A more minor error in this diagram snippet is the labeling of SOV-590A: it actually reads “SOV-59DA” if you look closely enough! My first inclination when sampling this real P&ID for inclusion in the book was to correct the errors, but I think an important lesson may be taught by leaving them in: documentation errors are a realistic challenge you will contend with on the job as an instrumentation professional!

The manufacturer's and process engineer's definitions of "normally" are not always in conflict. Take for example this P&ID segment, showing the solenoid control of an air vent door on a large furnace, designed to open up if the forced-draft fan (blowing combustion air into the furnace) happens to stop for any reason:



Here we have a *normally open* solenoid valve, designed by the manufacturer to pass instrument air pressure from the pressure ("P") port to the cylinder ("C") port when de-energized. The straight arrow with the "DE" label next to it reveals this to be the case. Instrument air pressure sent to the air door actuator holds the door shut, meaning the air door will swing open if ever instrument air pressure is vented by the solenoid. For this particular solenoid, this would require an *energized* condition.

The process engineer designing this Emergency Air Door control system chose to let the solenoid be in its de-energized state under typical operating conditions (when the furnace air door should be shut), a fact revealed by the letters "NDE" (normally de-energized) next to the solenoid coil symbol. Therefore, the "normal" process operating condition for this solenoid happens to be de-energized, which makes the manufacturer's definition of "normal" match the engineer's definition of "normal." The solenoid valve should be open (passing air to the door's actuating cylinder) under "normal" operating conditions.


10.4 On/off electric motor control circuits

An electric motor is often used as a discrete control element in a control system if driving a pump, conveyor belt, or other machine for the transportation of a process substance. As such, it is important to understand the functioning of motor control circuits.

Of all the available electric motor types, the most common found in industrial applications (by far) is the three-phase AC induction motor. For this reason, this section of the book will focus exclusively on this type of motor as a final control element.

10.4.1 AC induction motors

The basic principle of an AC induction motor is that one or more out-of-phase AC (sinusoidal) currents energize sets of electromagnet coils (called *stator* coils or windings) arranged around the circumference of a circle. As these currents alternately energize the coils, a magnetic field is produced which “appears” to rotate around the circle. This *rotating magnetic field* is not unlike the appearance of motion produced by an array of *chaser lights* blinking on and off in sequence: although the bulbs themselves are stationary, the out-of-phase sequence of their on-and-off blinking makes it appear as though a pattern of light “moves” or “chases” along the length of the array¹². Likewise, the superposition of magnetic fields created by the out-of-phase coils resembles a magnetic field of constant intensity revolving around the circle. The following images show how the magnetic field vector (the red arrow) is generated by a superposition of magnetic poles through one complete cycle (1 revolution), viewing the images from left to right, top to bottom (the same order as you would read words in an English sentence)¹³:



It should come as no surprise that the combined effect of these three-phase currents will be to create a resultant magnetic field vector rotating in a particular direction. After all, this is precisely how three-phase electric power is generated: by spinning a single magnet at the center of three sets of coils offset by 120 degrees. The rotating magnetic field generated by the stator windings of a three-phase motor is merely a reproduction of the rotor’s magnetic field inside the generator supplying the three-phase power!

¹²To view a flip-book animation of this sequence, turn to Appendix A.1 beginning on page 2813.

¹³To view a flip-book animation of this same sequence, turn to Appendix A.2 beginning on page 2840.

If a permanent magnet were placed within the center of this machine on a shaft such that it was free to rotate, the magnet would spin at the exact same speed as the rotating magnetic field. If the magnetic field completes one full revolution in $\frac{1}{60}$ of a second, the rotating speed of the magnet will be 60 revolutions per second, or 3600 revolutions per minute (3600 RPM). Since the magnet follows in lock-step with the rotating magnetic field, its rotational speed is said to be *synchronous*. We would thus identify this motor as a *synchronous AC motor*.

If an electrically conductive object were placed within the center of this same machine on a shaft such that it was free to rotate, the relative motion between the rotating magnetic field and the conductive object (rotor) would induce electric currents in the conductive object, producing magnetic fields of their own. Lenz's Law tells us that the effect of these induced magnetic fields would be to try to oppose change: in other words, the induced fields react against the rotating magnetic field of the stator coils in such a way as to minimize the relative motion. This means the conductive object would begin to rotate in the same direction as stator's rotating magnetic field, always trying to "catch up" to the rotating magnetic field. However, the conductive rotor could never exactly match the speed of the rotating magnetic field as in the case of a synchronous motor. If the rotor ever did achieve synchronous speed, there would no longer be any relative motion between the rotor and the rotating magnetic field, which means the induction would cease. No induction would mean no electric currents induced in the rotor, which would mean no reactive magnetic field, which would mean no torque to motivate the rotor. Thus, the electrically conductive rotor's speed must always slightly lag ("slip") behind the rotating magnetic field's synchronous speed in order to experience induction and thereby be able to create a torque¹⁴. We call this type of motor an *induction AC motor*.



It may come as a surprise for some to learn that *any* conductive object – ferromagnetic or not – will experience a torque when placed inside the rotating magnetic field generated by the stator coils. So long as the object is electrically conductive¹⁵, electromagnetic induction will ensure the creation of electric currents in the rotor, and these currents will produce their own magnetic fields which react against the stator's rotating magnetic field to produce a torque on the rotor.

¹⁴A helpful analogy for this effect is to imagine a sailboat traveling directly downwind, its motive force provided by a sail oriented perpendicular to the direction of travel. It should be obvious that in this configuration the sailboat cannot travel faster than the wind. What is less obvious is the fact that the sailboat can't even travel as fast as the wind, its top speed in this configuration being slightly less than the wind speed. If the sailboat somehow did manage to travel exactly at the wind's speed, the sail would go slack because there would be no relative motion between the sail and the wind, and therefore the sail would cease to provide any motive force. Thus, the sailboat must "slip" or "lag" behind the wind speed just enough to fill the sails with enough force to overcome water friction and maintain speed.

¹⁵As a vivid illustration of this concept, I once worked at an aluminum foundry where an AC induction motor stator assembly was used to electromagnetically spin molten aluminum inside the mold as it cooled from molten to solid state. Even though aluminum is a non-magnetic material, it was still spun by the stator's rotating magnetic field due to electromagnetic induction and Lenz's Law.

The effect of Lenz's Law between a magnet and a conductive object may be demonstrated by using a powerful permanent magnet and a strip of light-weight aluminum foil. Aluminum, of course, is electrically conductive but non-magnetic. However, despite the lack of magnetic attraction between the magnet and the foil, the foil will nevertheless experience a motive force if the magnet is swept past its surface rapidly, due to Lenz's Law:

Demonstration of Lenz's Law between a moving magnet and a strip of conductive metal



This very same principle is what makes an induction AC motor function: a rotating magnetic field induces electric currents in an electrically-conductive rotor, which then spins in the same direction as the magnetic field. An induction motor's rotor can never achieve synchronous speed on its own, for if it ever did the induction would cease due to a lack of relative motion between the rotating magnetic field and the rotor. The same is true of the aluminum foil strip experiments: the foil strip can never fully "catch up" to the moving magnet, for if it ever did the induction would cease and the motive force would disappear. Thus, induction machines always spin a bit slower than synchronous speed.

A typical "two-pole"¹⁶ induction motor operating at a power line frequency of 60 Hz has a synchronous speed of 3600 RPM (i.e. the rotating magnetic field is spinning 60 revolutions per second), but the rotor may only achieve a full-load speed of approximately 3540 RPM. Similarly, a typical "four-pole" induction motor with a synchronous speed of 1800 RPM¹⁷ may only attain a rotor speed of approximately 1760 RPM.

Induction motors are by far the most popular design in industry. The most common variant of the induction motor is the so-called *squirrel-cage* design, where the rotor is made up of aluminum bars joining two aluminum "shorting rings," one at either end of the rotor. Ferrous metal (iron alloy) fills

¹⁶Two magnetic poles in the stator *per phase*, which is the lowest number possible because each phase naturally produces both a "north" and a "south" pole when energized. In the case of a three-phase induction or synchronous motor, this means a total of *six* magnetic stator poles.

¹⁷Doubling the number of magnetic poles increases the number of AC power cycles required for the rotating magnetic field to complete one full revolution. This effect is not unlike doubling the number of light bulbs in a chaser light array of fixed length, making it seem as though the light sequence is moving slower because there are more bulbs to blink along the same distance.

the spaces between the rotor bars to provide a lower-reluctance magnetic “circuit” between stator poles than would otherwise be a large air gap if the rotor were simply made of aluminum. If the ferrous metal were removed from the rotor, the remaining aluminum bars and shorting rings would resemble the cage-wheel exercise machine used by hamsters and other pet rodents, hence the name.

A photograph of a small, disassembled three-phase AC induction “squirrel-cage” motor is shown here, revealing the construction of the stator coils and the rotor:



Given the extremely simple construction of AC induction motors, they tend to be very reliable. So long as the stator coil insulation is not damaged by excessive moisture, heat, or chemical exposure, these motors will continue to operate indefinitely. The only “wearing” components are the bearings supporting the rotor shaft, and those are easily replaced.

Starting a three-phase induction motor is as simple as applying full power to the stator windings. The stator coils will instantly produce a magnetic field rotating at a speed determined by the frequency of the applied AC power, and the rotor will experience a large torque as this high-speed (relative to the rotor’s stand-still speed of zero) magnetic field induces large electric currents in it. As the rotor comes up to speed, the relative speed between the rotating magnetic field and the rotating rotor diminishes, weakening the induced currents and also the rotor’s torque.

One way to “model” an AC induction motor is to think of it as an AC transformer with a short-circuited, movable secondary winding. When full power is first applied, the initial current drawn by the stator (primary) windings will be very large, because it “sees” a short-circuit in the rotor (secondary) winding. As the rotor begins to turn, however, this short-circuit draws less and less current until the motor reaches full speed¹⁸ and the line current approaches normal. As with

¹⁸As mentioned previously, the rotor can never fully achieve synchronous speed, because if it did there would be zero relative motion between the rotating magnetic field and the rotating rotor, and thus no induction of currents in the rotor bars to create the induced magnetic fields necessary to produce a reaction torque. Thus, the rotor must “slip” behind the speed of the rotating magnetic field in order to produce a torque, which is why the full-load speed of an induction motor is always just a bit slower than the synchronous speed of the rotating magnetic field (e.g. a 4-pole motor with a synchronous speed of 1800 RPM will rotate at approximately 1750 RPM).

a transformer, where a reduction in secondary current (from a load change) results in a reduction in primary current, the reduction in induced rotor current (from reduced slip speed) results in a reduction in stator winding current.

The huge surge of current at start-up time (as much as ten times the normal running current!) is called *inrush* current, causing the rotor to produce a large mechanical torque. As the rotor gains speed, the current reduces to a normal level, with the speed approaching the “synchronous” speed of the rotating magnetic field. If somehow the rotor achieves synchronous speed (i.e. the slip speed becomes zero), stator current will fall to an absolute minimum. If a mechanical power source “overdrives” a powered induction motor, forcing it to spin faster than synchronous speed, it will actually begin to function as a generator¹⁹ and source electrical power.

Any mechanical load causing the motor to spin slower likewise causes the stator windings to draw more current from the power supply. This is due to the greater slip speed causing stronger currents²⁰ to be induced in the rotor. Stronger rotor currents equate to stronger stator currents, just like a transformer where a heavier load on the secondary winding causes greater currents in both secondary and primary windings.

Reversing the rotational direction of a three-phase motor is as simple as swapping any two out of three power conductor connections. This has the effect of reversing the *phase sequence* of the power “seen” by the motor²¹. The flip-book animation beginning in Appendix A.1 beginning on page 2813 shows how reversing two of the three lines has the effect of reversing phase sequence.

An interesting problem to consider is whether it is possible to make an AC induction motor function on *single-phase* power rather than polyphase power. After all, it is the three-step phase sequence of three-phase AC power that gives the stator windings’ magnetic field its definite rotational direction. If we have only one sine wave supplied by the AC power source, is it possible to generate a truly *rotating* magnetic field? At best, it seems all we could ever produce with single-phase AC power is a *pulsing* or “blinking” magnetic field. If you imagine a string of light bulbs blinking on and off 180° out of phase (i.e. ABABABAB), one could argue the sequence is marching from A to B, or alternatively from B to A – there is no definite direction to the lights’ “motion.”


¹⁹In this mode, the machine is called an *induction alternator* rather than an *induction motor*.

²⁰Faraday’s Law of Electromagnetic Induction describes the voltage induced in a wire coil of N turns as proportional to the *rate of change* of the magnetic flux: $V = N \frac{d\phi}{dt}$. The greater the difference in speed between the rotor and the rotating magnetic field, the greater $\frac{d\phi}{dt}$, inducing greater voltages in the rotor and thus greater currents in the rotor.

²¹This principle is not difficult to visualize if you consider the phase sequence as a repeating pattern of letters, such as ABCABCABC. Obviously, the reverse of this sequence would be CBACBACBA, which is nothing more than the original sequence with letters A and C transposed. However, you will find that transposing *any* two letters of the original sequence transforms it into the opposite order: for example, transposing letters A and B turns the sequence ABCABCABC into BACBACBAC, which is the same *order* as the sequence CBACBACBA.

Since single-phase AC induction motors obviously exist, there must be a solution to this problem. In order to give the magnetic field within a single-phase stator assembly a definite rotation, we must *artificially create a second phase* within the motor itself. One common way to do this is to add a second set of stator windings offset from the first and energize those windings through a high-voltage capacitor, which creates a leading phase shift in winding current. This phase shift creates an out-of-step magnetic field in the second winding set, providing a definite direction of rotation. Once the motor comes up to speed, this auxiliary winding may be disconnected by a speed-sending switch, since a spinning motor will happily run²² on single-phase AC. This is called a *capacitor-start* induction motor, and it is the design used for most single-phase AC induction motors requiring a high starting torque (e.g. pumps, shop grinders, drill presses, etc.):

Single-phase, capacitor-start induction motor



One of the major principles to grasp about AC induction motors is that they *must start as polyphase machines, although they may continue to run as single-phase machines.*

²²I once encountered a washing machine induction motor with an “open” fault in the start winding. When energized, this motor remained still and hummed because it had no second phase to give its magnetic field a rotation. However, if you used your hand to give the motor a spin in either direction, *the motor would accelerate to full speed in that direction!*

A capacitor-start, single-phase electric motor is shown in the following photograph. My hand is touching the capacitor enclosure for the motor's starting winding. The speed switch is internal to the motor and cannot be seen in this photograph:



Capacitor-start motors are often designed in such a way that the starting winding draws much more current than the “run” winding, in order to provide a strong starting torque. This is important when the mechanical load being turned by the motor requires a great deal of torque to get moving, such as in the case of a reciprocating gas compressor or a fully-loaded conveyor belt. Due to this high current draw, starting windings are not rated for continuous duty, but rather must be de-energized shortly after starting the motor in order to avoid overheating.

Smaller AC motors, such as those used inside bench-top and rack-mount electronic equipment, use a completely different method for generating a rotating magnetic field from single-phase AC power. The following photograph shows one such motor, employing copper *shading coils* at the corners of the magnetic stator poles. The rotor has been removed, held by my fingers for inspection:



Instead of a capacitor creating a leading phase-shift for current through a special stator winding, this “shaded-pole” induction motor uses a pair of copper loops wrapped around the corners of the magnetic poles to produce a lagging phase shift in the magnetic field at those corners. The copper shading coils act as inductors, delaying the magnetic field through them by -90° , creating a secondary magnetic field that is out-of-step with the main magnetic field generated by the rest of the pole face. The out-of-step magnetic field together with the main magnetic field adjacent to it creates a definite direction of rotation²³.

An interesting experiment you can try yourself is to obtain²⁴ one of these small shaded-pole AC motors and make it rotate by applying pulsed DC power to it, from a battery. Every time you connect the stator winding to the battery, the increasing magnetic flux will lead at the non-shaded

²³In this example, the direction of rotation is counter-clockwise. The shaded poles are oriented counter-clockwise of center, which means their delayed magnetic fields create an “appearance” of rotation in that direction: the magnetic field achieves its peak strength first at the pole centers, and then later (delayed) at the shaded poles, as though there were an actual magnet rotating in that direction.

²⁴A convenient source of small shaded-pole motors is your nearest home improvement or hardware store, where they likely sell replacement electric motors for bathroom fans. Of course, you may also find such motors inside of a variety of discarded electric appliances as well. Being rather rugged devices, it is quite common to find the shaded-pole motor inside of an electrical appliance in perfect condition even though other parts of that appliance may have failed with age. In fact, the shaded-pole motor shown in the preceding photograph was salvaged from a “water-pic” electric toothbrush, the motor used to drive a small water pump (which in this case had mechanically failed) delivering water to the head of the toothbrush.


pole faces and lag at the shaded pole faces. Every time you disconnect the stator winding from the battery, the decreasing magnetic flux will lead at the non-shaded pole faces and lag at the shaded pole faces. In either case, the shaded poles' magnetic flux will lag behind that of the non-shaded poles, causing the rotor to rotate slightly in one definite direction.

The fact that all AC induction motors must start as polyphase machines even though they can run as single-phase machines means that an AC motor designed to run on three-phase power may actually continue to run if one or more of its phases are “lost” due to an open wire connection or blown fuse. The motor cannot deliver full-rated mechanical power in this condition, but if the mechanical load is light enough the motor will continue to spin even though it no longer has multiple phases powering it! A three-phase motor, however, *cannot start from a stand-still* on just one phase of AC power. The loss of phases to an AC induction motor is called *single-phasing*, and it may cause a great deal of trouble in an industrial facility. Three-phase electric motors that become “single-phased” from a fault in one of the three-phase power lines will refuse to start. Those that were already running under heavy (high-torque) mechanical load will stall. In either case, the stopped motors will simply “hum” and draw large amounts of current.

10.4.2 Motor contactors


To start up and shut down a three-phase AC induction motor, any three-pole switch with a suitable current rating will suffice. Simply closing such a switch to send three-phase power to the motor will cause it to start up, while opening the three-pole switch will cut power to the motor to make it turn off. If we desire to have *remote* start and stop control over a three-phase motor, we need a special relay with switch contacts big enough to safely conduct the motor's inrush current over many start and stop cycles. Large, high-current-rated electromechanical relays built for this very purpose are commonly referred to as *contactors* in industry.

A schematic diagram of a three-phase contactor connected to a three-phase motor (with fuses for overcurrent protection) is shown here:




Energizing terminals A1 and A2 magnetizes the electromagnet coil, causing all three switch contacts to simultaneously close, sending three-phase AC power to the motor. De-energizing the coil causes it to de-magnetize, releasing the armature and enabling a return spring inside the contactor to snap all three contacts to the open (off) position.

A contactor rated at 75 horsepower (at 480 volt AC 3-phase power) is shown here, both assembled and with the top cover removed to reveal the three sets of high-current electrical switch contacts:



Each phase switch contact is actually a series pair of contacts that make and break simultaneously with the actuation of a ferrous armature attracted by an electromagnet coil in the base of the contactor assembly. The operation of the three contact sets may be seen in this pair of photographs, the left-hand image showing the contacts in their normal (open) state, and the right-hand image showing the contacts closed (the armature “pulled in”) by the force of my finger:



Of course, it would be very dangerous to touch or manually actuate the contacts of a motor starting relay with the cover removed as shown. Not only would there be an electric shock hazard from touching any one of the bare copper contacts with your finger, but the arcing produced by closing and opening such contacts would pose *arc flash* and *arc blast* hazards. This is why all modern motor contactors are equipped with arc shield covers. The actual switch contact pads are not made of pure copper, but rather silver (or a silver alloy) designed to survive the repeated arcing and blasting action of large AC currents being initiated and interrupted.

Below the main power connection terminals (L1-L3, T1-T3) on this contactor hide two small screw terminals (commonly denoted A1 and A2) providing connection points to the electromagnet coil actuating the contactor:



Like most three-phase contactors, this one’s coil is rated for 120 volts AC. Although the electric motor may operate on three-phase, 480 volt AC power, the contactor coil and the rest of the control circuitry operates on a lower voltage for reasons of safety. Like all electromechanical relays, motor contactors use a low-power signal to control higher-power electric current to the load. This


“amplifying” action enables relatively small control switches, PLCs, and relay circuits to start and stop relatively large (high-current) electric motors.

10.4.3 Motor protection

An essential component of any high-power motor control circuit is some device to detect a condition of excessive *overload* and interrupt power to the motor before thermal damage occurs. A very simple and common overload protective device is known as an *overload heater*, consisting of resistive elements connected in series with the three lines of a 3-phase AC motor, designed to heat and to cool at rates modeling the thermal characteristics of the motor itself.


Fuses and circuit breakers also protect against overcurrent, but for different reasons and for different parts of the motor circuit. Both fuses and circuit breakers tend to be fast-acting devices, intended to interrupt overcurrent resulting from an electrical fault such as a phase-to-ground short circuit. They are sized to protect the wiring delivering power to a load, not (necessarily) the load itself. Thermal overload heaters, by contrast, are specifically designed to protect an electric motor from damage resulting from mild overcurrent conditions, such as what might be experienced if the motor becomes mechanically overloaded. The sizing of overload heaters is unrelated to wire ampacity, and therefore unrelated to the ratings of the fuses or circuit breakers delivering line power to the motor.

A schematic diagram of a three-phase overload connected to a three-phase contactor and three-phase motor is shown here:




Both contacts inside the overload assembly will remain in their resting (“normal”) states so long as the heater elements (the back-to-back “hook” symbols seen in the above diagram) remain cool. If one or more of the resistive heaters becomes too warm, however, the contacts will actuate and change state. The normally-closed overload contact (terminals 95 and 96) is typically wired in series with the contactor coil (terminals A1 and A2), so that a detected overload condition forces the contactor to de-energize and interrupt power to the motor.


The following photograph shows a three-phase contactor relay joined together with a set of three “overload heaters” through which all of the motor’s current flows. The overload heaters appear as three brass-colored metal strips near a red push-bar labeled “Reset”. The entire assembly – contactor plus overload heaters – is referred to as a *starter*:



Removing one of the heater elements reveals its mechanical nature: a small toothed wheel on one side engages with a lever when it is bolted into place in the overload assembly. That lever connects to a spring-loaded mechanism charged by the manual actuation of the red “Reset” push-bar, which in turn actuates a small set of electrical switch contacts:




The purpose of the overload heater is to heat up as the motor draws excessive current. The small toothed wheel is held in place by a rod immersed in a solidified mass of solder, encased in a brass cylinder underneath the heater strip. The next photograph shows the underside of the heater element, with the toothed wheel and brass cylinder plainly visible:



If the heater element becomes too hot (due to excessive motor current), the solder inside the brass cylinder will melt, allowing the toothed wheel to spin. This will release spring tension in the overload mechanism, allowing the small electrical switch to spring to an open state. This “overload contact” then interrupts current to the contactor’s electromagnet coil, causing the contactor to de-energize and the motor to stop.

Manually pressing the “Reset” push-bar will re-set the spring mechanism and re-close the overload contact, allowing the contactor to energize once more, but only once the overload heater element has cooled down enough for the solder inside the brass cylinder to re-solidify. Thus, this simple mechanism prevents the overloaded motor from being immediately re-started after a thermal overload “trip” event, giving it time to cool down as well.

A typical “trip curve” for a thermal overload unit is shown here, with time plotted against the severity of the overcurrent level:

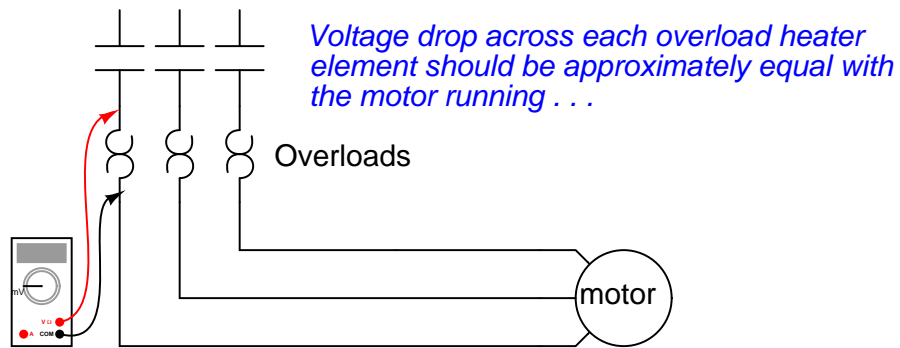


In contrast to a circuit breaker or fuse – which is sized to protect the power wiring from overcurrent heating – the overload heater elements are sized specifically to protect the *motor*. As such, they act as thermal models of the motor itself, heating to the “trip” point just as fast as the motor itself will heat to the point of maximum rated temperature, and taking just as long to cool to a safe temperature as the motor will. Another difference between overload heaters and breakers/fuses is that the heaters are not designed to directly interrupt current by opening²⁵, as fuses or breakers do. Rather, each overload heater serves the simple purpose of *warming* proportionately to the magnitude and time duration of motor overcurrent, causing a different electrical contact to open, which in turn triggers the contactor to open and interrupt motor current.

Of course, overload heaters only work to protect the motor from thermal overload if they experience similar ambient temperature conditions. If the motor is situated in a very hot area of the industrial process unit, whereas the overload elements are located in a climate-controlled “motor control center” (MCC) room, they may fail to protect the motor as designed. Conversely, if the overload heaters are located in a hot room while the motor is located in a freezing-cold environment (e.g. the MCC room lacks air conditioning while the motor is located in a freezer), they may “trip” the motor prematurely.

²⁵This is not to say overload heaters cannot fail open, because they can and will under extraordinary circumstances. However, opening like a fuse is not the design function of an overload heater.

An interesting “trick” to keep in mind for motor control circuit diagnosis is that overload heaters are nothing more than low-value resistors. As such, they will drop small amounts of voltage (usually quite a bit less than 1 volt AC) under full load current. This voltage drop may be used as a simple, qualitative measure of motor phase current. By measuring the voltage dropped across each overload heater (with the motor running), one may ascertain whether or not all phases are carrying equal currents. Of course, overload heaters are not precise enough in their resistance to serve as true current-measuring “shunts,” but they are more than adequate as qualitative indicators of relative phase current, to aid you in determining (for instance) if the motor suffers from an open or high-resistance phase winding:



As useful as thermal overload “heaters” are for motor protection, there are more effective technologies available. An alternative way to detect overloading conditions is to monitor the temperature of the stator windings directly, using thermocouples or (more commonly) RTDs, which report winding temperatures to an electronic “trip” unit with the same control responsibilities as an overload heater assembly. This sophisticated approach is used on large (thousands of horsepower) electric motors, and/or in critical process applications where motor reliability is paramount. Machine vibration equipment used to monitor and protect against excessive vibration in rotary machines is often equipped with such temperature-sensing “trip” modules just for this purpose. Not only can motor winding temperatures be monitored, but also bearing temperatures and other temperature-sensitive machine components so that the protective function extends beyond the health of the electric motor.


Devices specifically constructed to monitor the condition of electrical power components such as motors, generators, transformers, or distribution lines, and take action to protect those components in the event their parameters fall outside safe limits, are generally known as *protective relays*²⁶. A protective relay is designed to monitor physical variables such as line currents and winding temperatures relevant to a large electrical component, then automatically initiate a “trip” action to shut off power to that component by sending a signal to the nearest circuit breaker or other automatic disconnect device²⁷.

Originally, protective relays were electromechanical in nature, using coils, magnets, springs, rotating disks, and other components to detect and act upon out-of-spec electrical measurements. Modern protective relays – for electric motors or for other electric power components such as generators, power lines, and transformers – use microprocessors instead of electromagnetic mechanisms to perform the same basic functions. With microprocessor technology comes vast increases in responsiveness and precision of timing, as well as digital networking capability to share system data among other components and to human operators.

²⁶For a more complete coverage of protective relays, refer to section 25.7 beginning on page 1985.

²⁷One way to help clarify the function of a protective relay is to envision circuit protection without one. Household and low-current industrial circuit breakers are constructed to have their own internal current-sensing elements (either thermal or magnetic) to force the circuit breaker open automatically when current exceeds a pre-set limit. With protective relays, the circuit breaker instead has a “trip coil” which will cause the breaker to trip when energized. The breaker then relies entirely on the (external) protective relay to tell it when to trip. By relegating the function of event detection to a sophisticated, external relay, the circuit breaker may act much “smarter” in protecting against a wider variety of faults and abnormal conditions than if it relied entirely on its own internal overcurrent-sensing mechanism.

A diagram showing how a modern (digital) protective relay would monitor various parameters on a medium-voltage (4160 volts AC, three-phase) industrial electric motor is shown here:



In this example, line voltage (4160 volts AC) and line current are both too great to be directly connected to the protective relay, and so the relay senses line voltage and line current via *potential transformers* (PTs) and *current transformers* (CTs), respectively. A potential transformer²⁸ is a precision device providing a known accurate step-down ratio, usually down to 120 volts or 240 volts AC full-scale, for the protective relay to directly sense. Likewise, a current transformer is a precision device providing a known and accurate step-down ratio for current (actually a step-up from the perspective of voltage), usually down to 1 amp or 5 amps AC full-scale, for the protective relay to directly sense. Both transformers provide *galvanic isolation* (a complete lack of electrical conductivity) between the medium-voltage motor power conductors and the protective relay electronics while still permitting accurate sensing of line voltage and line current.

The *zero sequence* CT is a special current transformer encircling all three motor phase conductors, providing indication of a ground fault within the motor. The fact that this CT measures the instantaneous algebraic sum of currents in and out of the motor means that in ordinary operation it will output absolutely zero signal, since Kirchhoff's Current Law states that the algebraic sum of currents into and out of a node (the motor here is considered a node) must be zero. If, however,

²⁸Potential transformers are also known as *voltage transformers*, abbreviated *VT*.

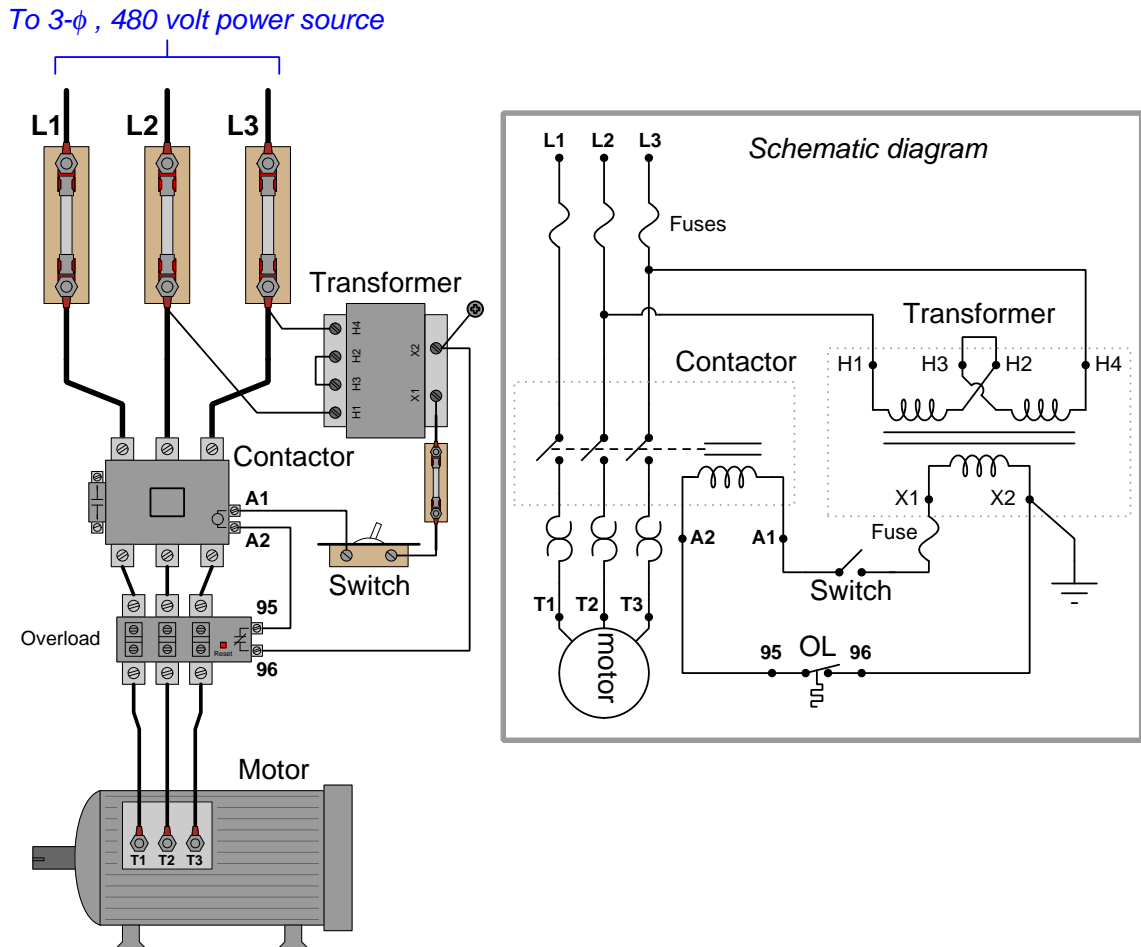
a ground fault develops within the motor where some AC current “leaks” from a stator winding to earth ground to return to the 4160 VAC power source’s neutral connection, that imbalance of phase currents will be sensed by the zero sequence CT, since that ground fault current represents a fourth path for current not accounted for by the three power conductors passing through to the motor.

This next photograph shows the front panel display of a General Electric (“Multilin”) model 369 protective relay for an electric motor:



10.4.4 Motor control circuit wiring


A simple three-phase, 480 volt AC motor-control circuit is shown here, both in pictorial and schematic form. This entire assembly consisting of contactor, overload block, control power transformer, power fuses (or alternatively, a circuit breaker) and associated components is informally referred to as a *bucket*:



Note how a *control power transformer* steps down the 480 volt AC to provide 120 volt AC power for the contactor coil to operate on. Furthermore, note how the overload (“OL”) contact is wired in series with the contactor coil so that a thermal overload event forces the contactor to de-energize and thus interrupt power to the motor even if the control switch is still in the “on” position. The overload heaters appear in the schematic diagram as pairs of back-to-back “hook” shapes, connected in series with the three “T” lines of the motor. Remember that these “OL” heater elements do not directly interrupt power to the motor in the event of an overload, but rather signal the “OL” contact to open up and de-energize the contactor.

In an automatic control system, the toggle switch would be replaced by another relay contact (that relay controlled by the status of a process), a process switch, or perhaps the discrete output channel of a programmable logic controller (PLC).

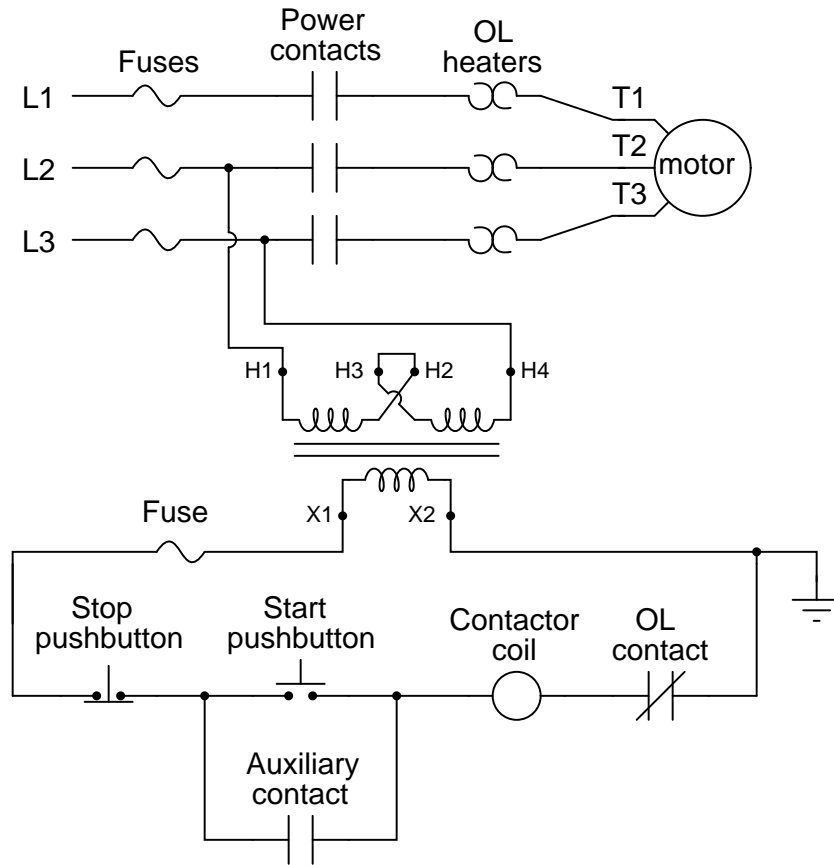
It should be noted that a toggling-style of switch is necessary in order for the motor to continue to run after a human operator actuates the switch. The motor runs when the switch is in the closed state, and stops when the switch opens. An alternative to this design is to build a *latching* circuit allowing the use of momentary contact switches (one to start, and one to stop). A simple latching motor control circuit is shown here:



In this circuit, an *auxiliary contact* actuated by the motor contactor is wired in parallel with the “Start” pushbutton, so that the motor contactor continues to receive power after the operator releases the switch. This parallel contact – sometimes called a *seal-in contact* – latches the motor in an “on” state after a momentary closure of the “Start” pushbutton switch.


A normally-closed “Stop” switch provides a means to “un-latch” the motor circuit. Pressing this pushbutton switch opens the control circuit, forcing current to halt through the coil of the contactor, which then opens the three motor power contacts as well as the auxiliary contact used to maintain the contactor’s energized state.

A simple *ladder diagram* showing the interconnections of all components in this motor control circuit makes this system easier to understand:



Most on/off motor control circuits in the United States are some variation on this wiring theme, if not identical to it. Once again, this system could be automated by replacing the “Start” and “Stop” pushbutton switches with process switches (e.g. pressure switches for an air compressor control system) to make a system that starts and stops automatically. A programmable logic controller (PLC) may also be used to provide the latching function rather than an auxiliary contact on the contactor. Once a PLC is included in the motor control circuit, a great many automatic control features may be added to enhance the system’s capabilities. Examples include timing functions, motor cycle count functions, and even remote start/stop capability via a digital network connecting to operator interface displays or other computers.

In applications where reversing motor control is desired, a pair of contactors may be wired together as shown here:



Note how motor reversal is accomplished by swapping phases L1 and L3: in the forward direction, power line conductor L1 connects to motor terminal T1, L2 connects to T2, and L3 connects to T3. In the reverse direction L2 still connects to T2, but L1 now connects to T3 and L3 now connects to T1. Recall the principle that swapping *any two phases* in a three-phase power system reverses the phase rotation, which in this case make the electric motor spin the other direction.

With two contactors, the control circuit now contains two coils to actuate those contactors: one marked “forward” and the other marked “reverse”. Separate “forward” and “reverse” pushbutton switches send power to those coils, and separate seal-in auxiliary contacts connected in parallel with their respective pushbuttons latch each one.

An important feature of this reversing starter circuit is the inclusion of *interlocking* contacts in each rung of the circuit. In the forward-control circuit, a normally-closed auxiliary contact actuated by the “reverse” contactor is wired in series, and vice-versa in the reverse-control circuit. The purpose of an “interlock” is to prevent incompatible events from happening, in this case preventing the actuation of the “reverse” contactor when the “forward” contactor is already actuated, and vice-versa. If both contactors were to be simultaneously actuated, it would result in a direct phase-to-phase fault (short-circuit) between L1 and L3!

Some reversing motor starters provide a feature called *mechanical interlocking*, where the motion of the armature in each contactor is restrained in such a way that both cannot actuate simultaneously. This usually takes the form of a “rocking beam” lever preventing one contactor armature from being pulled in while the other contactor’s armature is pulled in, similar to a “see-saw” playground toy where only one end can be down at any given time. It is not uncommon for both electrical and mechanical interlocking to be used in the same reversing starter, as a measure of extra protection.


A modern trend in motor control is the use of digital networks to both command the contactor as well as monitor the motor's operating status remotely. This next photograph²⁹ shows a digitally monitored and controlled "bucket," using DeviceNet as the control network:



Using a digital network standard such as Ethernet, DeviceNet, Modbus, Profibus, or any number of others to monitor and control a motor brings a host of benefits for maintenance and operations. Control wiring is vastly simplified with digital networks, as a single network cable is able to address multiple motor buckets. The "smart" network interface module installed in the bucket may be designed to monitor such parameters as line voltage, line current, phase imbalance, and power factor to report these values to the host control system via the network.

²⁹This bucket was still under construction at the time the photograph was taken. As such, none of the motor leads have been connected, which is why there are no power conductors exiting the bottom of the bucket. Instead, all you see are three terminals ready to accept heavy-gauge motor leads.

It is common for the network interface module inside the bucket to have its own digital display for local indication of these parameters as well. A close-up photograph of a Square-D “Motor Logic Plus” unit shows some of its locally-accessible features:



The PLC connected to the network is able to access all these values as well, reporting them to operations and/or maintenance personnel as desired. Instead of individual wires running between the PLC and the motor starter to command each motor to run and stop, the PLC simply transmits “start” and “stop” commands over the network to individually addressed digital starter modules. The network wiring may simply be paralleled (“daisy-chained”) between units, such that several buckets reside on the same physical network, each one programmed with a unique address. A PLC connected to this same network is thus able to access and control all parameters for all motors on that network.

10.5 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **“Normal” switch status:** the “normal” status of a switch contact as defined by the manufacturer is its *resting* condition (minimum stimulus).
- **“Seal-in” circuit:** when an electrical relay uses one of its own switch contacts to continue its own coil energization after the initial triggering event has passed. Relevant to motor starter circuits.
- **Lenz’s Law:** any magnetic field arising from electromagnetic induction opposes the inducing field. Relevant to the operation of induction AC motors.
- **Rotating magnetic field:** this is necessary to cause an AC induction motor to spin in a particular direction, and is generated by polyphase field poles (i.e. multiple magnetic fields that are out-of-phase with each other). All AC induction motors require such a polyphase magnetic field to start up in a particular direction, although a single-phase magnetic field is sufficient to maintain rotation once started.

References

“369 Motor Management Relay Instruction Manual”, part number 1601-0077-BU, GE publication GEK-106288R, GE Multilin, Markham, Ontario, Canada, 2010.

“857 Protection System User Manual”, publication 857-UM001A-EN-P, Rockwell Automation, Inc, Milwaukee, WI, 2009.

“ASCO Nuclear Catalog – Nuclear Products Qualified to IEEE Specifications”, ASCO Valve Inc.

“Compressed Air Manual”, 6th edition, Atlas Copco Compressor AB, Sweden, 1998.

Croft, Terrell and Summers, Wilford I., *American Electrician’s Handbook*, Eleventh Edition, McGraw-Hill Book Company, New York, NY, 1987.

Fitzgerald, A.E. and Higginbotham, David E., *Basic Electrical Engineering*, Second Edition, McGraw-Hill Book Company, New York, NY, 1957.

“Fluid Power Design Engineer’s Handbook”, bulletin no. 0105-B1, Parker-Hannifin Corporation, Cleveland, OH, 1973.

“General Service Solenoid Valves – 3/2 Series 8300/8315”, document 8300R1, ASCO Valve Inc.

“Hydraulic Cartridge Systems – Product Selection Guide”, catalog HY15-3501/US, Parker Hannifin Corporation, Lincolnshire, IL, 2008.

“Hydraulic Valves Industrial Standard”, catalog HY11-3500/UK, Parker Hannifin GmbH & Co. KG, Kaarts, Germany, 2008.


Robichaud, S.N., “Plumb Your Plant Air System with Plastic Pipe”, *Hydraulics & Pneumatics*, June 1, 1999.

Chapter 11

Relay control systems


The word “discrete” means *individual* or *distinct*. In engineering, a “discrete” variable or measurement refers to a true-or-false condition. Thus, a discrete control system is one designed to operate on Boolean (“on” or “off”) signals supplied by discrete sensors such as process switches.

A form of discrete control taught in every introductory course on digital electronics involves the use of circuits called *logic gates*. These circuits input one or more Boolean signals, and output a Boolean signal according to a simple rule such as “AND” or “OR”:




AND

A	B	Output
0	0	0
0	1	0
1	0	0
1	1	1




OR

A	B	Output
0	0	0
0	1	1
1	0	1
1	1	1




XOR

A	B	Output
0	0	0
0	1	1
1	0	1
1	1	0




NAND

A	B	Output
0	0	1
0	1	1
1	0	1
1	1	0



NOR

A	B	Output
0	0	1
0	1	0
1	0	0
1	1	0




XNOR

A	B	Output
0	0	1
0	1	0
1	0	0
1	1	1


Industrial control systems rarely utilize logic gates in a direct fashion for discrete control systems, although the fundamental *concepts* of “AND,” “OR,” and other gate types are universally applied. Instead, control functions are either implemented using electromechanical relays and/or

with programmable digital devices such as PLCs (Programmable Logic Controllers). This chapter focuses on the practical use of both technologies for industrial discrete control.


An “AND” function is equivalent to *series-connected* normally-open contacts in a relay control circuit, because the lamp will energize only if switch A *and* switch B are actuated:



An “OR” function is equivalent to *parallel-connected* normally-open contacts in a relay control circuit, because the lamp will energize if switch A *or* switch B is actuated:



The “NOT” function is equivalent to a single normally-closed contact in a relay control circuit, because the lamp will energize only if the switch is *not* actuated:




11.1 Control relays

An *electromechanical relay* is an electrical switch actuated by an electromagnet coil. As switching devices, they exhibit simple “on” and “off” behavior with no intermediate states. Relays are very useful devices, as they allow a single discrete (on/off) electrical signal to *control* much greater levels of electrical power, and/or multiple power or control signals that are otherwise isolated from each other. For example, a relay may be controlled by a low-voltage, low-current signal that passes through a delicate switch of some sort (e.g. limit switch, proximity switch, optical sensor), and then the switching contacts of that relay may be used to control a much higher-voltage, higher-current circuit, and even multiple circuits given multiple sets of switching contacts.


The electronic schematic symbol for a simple single-pole, single-throw (SPST) relay is shown here:

SPST relay
(NO contact)




A coil of wire wrapped around a laminated ferrous core provides the magnetic field necessary to actuate the switch mechanism. This electromagnet coil’s actuating influence on the relay’s contact(s) is represented by the dashed line. This particular relay is equipped with *normally open* (NO) switch contacts, which means the switch will be in the open (off) state when the relay coil is de-energized. Recall from section 9.1 that the “normal” status of a switch is the resting condition of *no stimulation*. A relay switch contact will be in its “normal” status when its coil is not energized. A single-pole, single-throw relay with a normally-closed (NC) switch contact would be represented in an electronic schematic like this:


SPST relay
(NC contact)




In the electrical control world, the labels “Form-A” and “Form-B” are synonymous with “normally open” and “normally closed” contacts, respectively. Thus, we could have labeled the SPST relay contacts as “Form-A” and “Form-B,” respectively:



An extension of this theme is the single-pole, double-throw (SPDT) relay contact, otherwise known as a “Form-C” contact. This design of switch provides both a normally-open and normally-closed contact set in one unit, actuated by the electromagnet coil:




A further extension on this theme is the double-pole, double-throw (DPDT) relay contact. This design of switch provides two sets of Form-C contacts in one unit, simultaneously actuated by the electromagnet coil:



DPDT relays are some of the most common found in industry, due to their versatility. Each Form-C contact set offers a choice of either normally-open or normally-closed contacts, and the two


sets (two “poles”) are electrically isolated from each other so they may be used in different circuits.

A common package for industrial relays is the so-called *ice cube relay*, named for its clear plastic case allowing inspection of the working elements. These relays plug into multi-pin base sockets for easy removal and replacement in case of failure. A DPDT “ice cube” relay is shown in the following photographs, ready to be plugged into its base (left) and with the plastic cover removed to expose both sets of Form-C contacts (right):





These relays connect to the socket with eight pins: three for each of the two Form-C contact set, plus two more pins for the coil connections. Due to the pin count (8), this style of relay base is often referred to as an *octal* base.

A closer view of one Form-C contact shows how the moving metal “leaf” contacts one of two stationary points, the actual point of contact being made by a silver-coated “button” at the end of the leaf. The following photographs show one Form-C contact in both positions:



Industrial control relays usually have connection diagrams drawn somewhere on the outer shell to indicate which pins connect to which elements inside the relay. The style of these diagrams may vary somewhat, even between relays of identical function. Take for instance the diagrams shown here, photographed on three different brands of DPDT relay:




Bear in mind that these three relays are *identical* in their essential function (DPDT switching), despite differences in physical size and contact ratings (voltage and current capacities). Only two of the three diagrams shown use the same symbols to represent contacts, and all three use unique symbols to represent the coil.

11.2 Relay circuits

Electromechanical relays may be connected together to perform logic and control functions, acting as logic elements much like digital gates (AND, OR, etc.). A very common form of schematic diagram showing the interconnection of relays to perform these functions is called a *ladder diagram*. In a “ladder” diagram, the two poles of the power source are drawn as vertical rails of a ladder, with horizontal “rungs” showing the switch contacts, relay contacts, relay coils, and final control elements (lamps, solenoid coils, motors) drawn in between the power rails.

Ladder diagrams differ from regular schematic diagrams of the sort common to electronics technicians primarily in the strict orientation of the wiring: vertical power “rails” and horizontal control “rungs.” Symbols also differ a bit from common electronics notation: relay coils are drawn as circles, with relay contacts drawn in a way resembling capacitors:


Ladder diagram symbols



Unlike schematic diagrams where the association between relay coils and relay contacts is represented by dashed lines, ladder diagrams associate coils and contacts *by label*. Sometimes you will find relay contacts labeled identically to the coil (e.g. coil labeled CR5 and all contacts for that relay also labeled CR5) while other times you will find suffix numbers used to distinguish individual contacts within each relay from each other (e.g. coil labeled CR5 and its three contacts labeled CR5-1, CR5-2, and CR5-3).


Another notable convention in relay circuits and their ladder diagrams is that each and every wire in the circuit is labeled with a number corresponding to common connection points. That is, wires connected together always bear the same number: the common number designates a condition of electrical commonality (all points bearing the same number are *equipotential* to each other). Wire numbers only change when the connection passes through a switch or other device capable of dropping voltage.

An actual ladder diagram of a relay-based motor control system is shown here, complete with red-line edits showing modifications to the circuit made by an industrial electrician:



Perhaps the most confusing aspect of relay control circuits for students to grasp is the meaning of *normal* as it applies to the status of relay contacts. As discussed previously, the word “normal” in this context – whether it be the status of hand switches, process switches, or the switch contacts inside control relays – means “in a condition of rest” or no stimulation. In other words, a “normally-open” relay contact is open *when the relay coil is unpowered* and closed when the relay coil is powered. Likewise, a “normally-closed” relay contact is closed *when the relay coil is unpowered* and open when the relay coil is powered.


To illustrate this concept, let us examine a relay control circuit where a pressure switch activates an alarm light:



Here, both the pressure switch and the relay contact (CR1-1) are drawn as normally-closed switch contacts. This means the pressure switch contact will be closed when the applied pressure is less than its trip point (50 PSI), and the relay switch contact will be closed when the relay coil is de-energized.

When analyzing the operation of a relay control system, it is helpful to have some way to temporarily denote the conductive status of switch contacts and the energization status of relay coils (i.e. a notation we might sketch using pencil on a diagram to help us follow the operation of the circuit). A symbology I recommend is the use of arrow and “X” symbols to represent power flow and no power flow (respectively). These symbols clearly denote component status while avoiding confusion with the symbols used to denote *normal* status of switch contacts¹.


In this next diagram, we assume the applied pressure is less than 50 PSI, leaving the pressure switch in its “normal” (closed) state:



Since the pressure is insufficient to actuate the pressure switch, its contact remains in the “normal” state (closed). This sends power to relay coil CR1, thus actuating contact CR1-1 and holding it in the *open* state. With CR1-1 contact open, the alarm lamp receives no power. In this example we see the pressure switch in its “normal” state but the relay in the *actuated* state.

¹An unfortunately common tendency among novices is to sketch slash marks through relay contact symbols in order to show when they happen to be closed. This is a very bad habit, and should be discouraged at all times! Diagonal lines drawn through a contact symbol are supposed to denote the contact to be *normally*-closed, not *closed*: it shows that a switch contact will be in the closed (conducting) state *when it is at rest*. What we actually need is a different kind of symbol to show when a contact is closed during any arbitrary condition we may imagine. When someone uses this same symbology to denote a contact that happens to be closed during some condition, it needlessly confuses the concepts of *closed* versus *normally-closed*.

Using arrow and “X” symbols again to represent the presence or absence of power in this circuit, we will now analyze its status with an applied switch pressure greater than 50 PSI:




Now that there is sufficient fluid pressure applied to the switch to actuate it, its contact is forced into the actuated state which for this “normally-closed” switch is open. This open condition de-energizes relay coil CR1, allowing relay contact CR1-1 to spring-return to its normal status (closed), thus sending power to the alarm lamp. From this analysis we see that the lamp fulfills the function of a *high pressure alarm*, energizing when the applied pressure exceeds the trip point.

Where students typically find themselves confused is assuming the switch contact will be in the same state it is drawn in. This is not necessarily true. The way switch contacts are drawn merely reflects their *normal* status as defined by the switch manufacturer, which means the status of the switch when there is no (or insufficient) actuating stimulus present. Whether or not the switch will actually be in its normal state at any given time is a question of whether or not a sufficient stimulus is present to actuate that switch. Just because a switch is drawn normally-closed does not necessarily mean it *will* be closed when you go to analyze it. All it means is that the switch will be closed *when nothing actuates it*.


This exact same principle applies to relay ladder-logic programming in electronic control systems called PLCs (Programmable Logic Controllers). In a PLC, a digital microprocessor performs the logic functions traditionally provided by electromechanical relays, with the programming for this microprocessor taking the form of a relay diagram (also called a “ladder-logic” diagram).

Here, we will emulate the exact same high-pressure alarm circuit using an Allen-Bradley MicroLogix 1000 PLC instead of a relay coil:

Wiring diagram:



Ladder-logic program:




Suppose a fluid pressure of 36 PSI is applied to the pressure switch. This is less than the switch's trip setting of 50 PSI, leaving the switch in its "normal" (closed) state. This sends power to input I:0/2 of the PLC. The contact labeled I:0/2 drawn in the ladder-logic program of the PLC acts like a relay contact driven by a coil energized by input terminal I:0/2. Thus, the closed pressure switch

contact energizes input terminal I:0/2, which in turn “closes” the normally-open contact symbol I:0/2 drawn in the ladder-logic program. This “virtual” contact sends virtual power to a virtual coil labeled B3:0/0, which is nothing more than a single bit of data in the PLC’s microprocessor memory. “Energizing” this virtual coil has the effect of “actuating” any contact drawn in the program bearing the same label. This means the normally-closed contact B3:0/0 will now be “actuated” and thus in the open state, not sending virtual power to the output coil O:0/1. With virtual coil O:0/1 “unpowered,” the real-life output O:0/1 on the PLC will be electrically open, and the alarm lamp will be unpowered (off).

If we apply a fluid pressure of 61 PSI to the pressure switch, the normally-closed pressure switch contact will be actuated (forced) into the open state. This will have the effect of de-energizing PLC input I:0/2, thus “opening” the normally-open virtual contact in the PLC program bearing the same label. This “open” virtual contact interrupts virtual power to the virtual coil B3:0/0, causing the normally-closed virtual contact B3:0/0 to “close,” sending virtual power to virtual coil O:0/1. When this virtual output coil “energizes,” the real-life output channel of the PLC activates, sending real power to the alarm light to turn it on, signaling a high-pressure alarm condition.

We may simplify this PLC program further by eliminating the virtual control relay B3:0/0 and simply having input I:0/2 activate output O:0/1 through a “normally-closed” virtual contact:




The effect is the same: the PLC output O:0/1 will activate whenever input I:0/2 de-energizes (whenever the pressure switch is opened by a high pressure), turning on the alarm lamp in a high-pressure condition. In a low-pressure condition, the energized input I:0/2 forces the virtual normally-closed contact I:0/2 to open, thus de-energizing the PLC’s output O:0/1 and turning the alarm lamp off.

Programmable Logic Controllers have not only greatly simplified the wiring of industrial logic controls by replacing multitudes of electromechanical relays with a microprocessor, but they have also added advanced capabilities such as counters, timers, sequencers, mathematical functions, communications, and of course the ability to easily modify the control logic through programming rather than by re-wiring relays. The beauty of ladder-logic programming is that it translates the technician’s understanding of traditional relay control circuits into a virtual form where contacts and coils interact to perform practical control functions. A key concept to master, however, is the association of real-life conditions to switch status based on the “normal” representation of those switch contacts, whether the switches be real (relay) or virtual (PLC). Once this vital concept is

mastered, both hard-wired relay control circuits and PLC programs become possible to understand. Without mastering this vital concept, neither relay control circuits nor PLC programs may be understood.

11.3 Interposing relays


In addition to directly performing logic functions, electromechanical relays may also be used as *interposing* devices between mismatched sensors, controllers, and/or control devices. A very simple example of a relay used to interpose between mismatched devices is shown in the following circuit diagram, where a delicate toggle switch is used to control a bank of high-power lights for an off-road vehicle:



In this circuit the relay performs no logic function whatsoever. Rather, it merely “amplifies” the signal sent by the dashboard toggle switch to send or halt power to the bank of high-power lights. Without the relay, a much heavier-duty toggle switch would have to be installed in the dashboard of this vehicle to safely and reliably make and break the light circuit.

Another example of an interposing relay found in automotive applications is the use of a “solenoid” in the electric starting motor circuit for an internal combustion engine. The “start” control switch is typically actuated by the driver turning a key, that switch mounted on the steering column or dashboard of the vehicle. The starting motor, meanwhile, typically draws *hundreds of amps* of current as it labors to start up the engine. A keyswitch capable of making and breaking hundreds of amps of current would be enormous, and in fact dangerous to locate in the cab of the vehicle. The “solenoid” relay connected between the keyswitch and the starting motor relocates that danger, and allows a relatively delicate keyswitch to safely activate the high-power motor.

An industrial example of an interposing relay between mismatched devices is shown here, where a DC output proximity switch must trigger an input channel to a Programmable Logic Controller (PLC) rated for 120 volts AC:



Again, the relay in this system performs no logic function, but merely allows the proximity switch to drive one of the PLC input channels. Directly connecting the proximity switch to one of the input channels of the PLC is not a practical option, because this particular PLC input requires 120 volts AC to activate, and our proximity switch operates on 24 volts DC. The mismatch between switch voltage and PLC input voltage requires us to use the relay to “interpose” between the switch and PLC. When the proximity switch senses an object nearby, its output activates, which in turn energizes the relay coil. When the relay contact magnetically closes, it completes a circuit for 120 volts AC to reach input channel 0 on the PLC, thereby energizing it.

An important detail in this relay circuit is the inclusion of a *commutating diode* in parallel with the relay coil, the purpose being to dissipate the coil’s stored energy upon de-energization when the proximity switch turns off. Without this diode in place, the coil’s “kickback” voltage (which may reach hundreds of volts in potential) will destroy the proximity switch’s output transistor.


Note how this commutating diode appears to be connected “backwards” with regard to the polarity of the 24 volt DC power source: cathode toward the source’s positive pole and anode toward the source’s negative pole. This is intentional, as we do *not* wish to have the diode conduct when power is applied to the relay coil through the proximity switch². The diode only turns on when the polarity reverses, which is what happens when the proximity switch turns off and the relay coil’s magnetic field collapses (now acting as a source rather than as a load). As the relay coil

²If the diode were connected the other way, it would pass current whenever the proximity switch turned on, shorting past the relay coil and most likely damaging the proximity switch in doing so!

temporarily outputs a “reverse” voltage, the diode gives that coil a continuous path for its current while dropping a low voltage (about 0.7 volts DC), dissipating the coil’s stored energy in the form of heat at the diode.

Interposing relays are also used to connected mismatched PLC outputs and control devices. In this application, the mismatch may be in terms of voltage ratings and/or current ratings. As with the input interposing circuit shown previously, the task of the relay in an output interposing circuit is to be controlled by the PLC’s output channel, and in turn direct power to a field device that is itself incompatible with the PLC’s output.

The following diagram shows an example of an interposing relay connected to a PLC output channel:



In this circuit the PLC’s transistor outputs can only handle 24 volts DC, and at fairly low current. The three-phase contactor³ coil requires 120 volts AC at modest current levels to function, and so the relay interposes between the PLC’s low-voltage and low-current output channel and the relatively high-voltage and high-current demands of the contactor’s coil. Once again we see the use of a commutating diode to dissipate the relay coil’s stored energy whenever the PLC de-energizes it, so that the resulting “kickback” voltage does not damage the fragile transistor output circuitry within the PLC.

³A “contactor” is nothing more than a very large electromechanical relay, and itself is a form of interposing device. Its purpose is to make and break three-phase AC power to a heavy load (e.g. an electric motor) at the command of a much smaller electrical signal, in this case a 120 volt AC signal sent to the coil of the contactor.

11.4 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Amplification:** the control of a relatively large signal by a relatively small signal. Relevant to the role of relays as interposing devices.
- **Interposing:** the use of a relay as an intermediary between electrically incompatible devices.
- **“Normal” switch status:** the “normal” status of a switch contact as defined by the manufacturer is its *resting* condition (minimum stimulus).
- **“Seal-in” circuit:** when an electrical relay uses one of its own switch contacts to continue its own coil energization after the initial triggering event has passed. Relevant to all manner of relay control circuits.

References

Summers, Wilford I. and Croft, Terrell, *American Electrician's Handbook*, Eleventh Edition, McGraw-Hill Book Company, New York, NY, 1987.

Chapter 12

Programmable Logic Controllers

Every control system may be divided into three general sections: *input* devices (sensors), *controllers*, and *output* devices (actuators). The input devices sense what is happening in the process, the controller decides what to do about it, and the output devices manipulate the process to achieve the desired result.

A *programmable logic controller* or *PLC* is a general-purpose controller, applicable to many different types of process control applications. The word “programmable” in its name reveals just why PLCs are so useful: the end-user is able to program, or instruct, the PLC to do virtually any control function imaginable. Unlike PID loop controllers, which are special-purpose devices intended to perform a single type of control function, a PLC may be instructed to do almost anything with the signals it receives from input devices.

PLCs were introduced to industry as electronic replacements for electromechanical relay controls. In applications where relays typically control the starting and stopping of electric motors and other discrete output devices, the reliability of an electronic PLC meant fewer system failures and longer operating life. The re-programmability of a PLC also meant changes could be implemented to the control system strategy much easier than with relay circuits, where re-wiring was the only way to alter the system’s function. Additionally, the computer-based nature of a PLC meant that process control data could now be communicated by the PLC over *networks*, allowing process conditions to be monitored in distant locations, and by multiple operator stations.

The legacy of PLCs as relay-replacements is probably most evident in their traditional programming language: a graphical convention known as a *Ladder Diagram*. Ladder Diagram PLC programs resemble ladder-style electrical schematics, where vertical power “rails” convey control power to a set of parallel “rung” circuits containing switch contacts and relay coils. A human being programming a PLC literally draws the diagram on the screen, using relay-contact symbols to represent instructions to read data bits in the PLC’s memory, and relay-coil symbols to represent instructions writing data bits to the PLC’s memory. This style of programming was developed to make it easier for industrial electricians to adapt to the new technology of PLCs. While Ladder Diagram programming definitely has limitations compared to other computer programming languages, it is relatively easy to learn and diagnose, which is why it remains popular as a PLC programming language today.

Typical devices connecting to a PLC’s inputs include hand switches, process switches, sensors, analog transmitters (4-20 mA), thermocouples, thermistors, and strain gauges. Typical devices

connecting to a PLC's outputs include electric lamps, solenoids, relay coils, motor contactors, analog final control elements (e.g. throttling control valves, variable-speed motor drives), and audible buzzers. While PLCs were originally designed for discrete (on/off) control applications such as conveyor belt management, batch sequencing, and assembly line controls, modern PLCs are equally capable of inputting and outputting analog signals as well. Thus, it is just as likely now to find a PLC performing PID loop control as it is to find a PLC turning discrete devices on and off.

12.1 PLC examples

Programmable logic controllers are essentially nothing more than special-purpose, industrial computers. As such, they are built far more ruggedly than an ordinary personal computer (PC), and designed to run extremely reliable operating system software¹. PLCs as a rule do not contain disk drives, cooling fans, or any other moving parts. This is an intentional design decision, intended to maximize the reliability of the hardware in harsh industrial environments where the PLC chassis may be subjected to temperature extremes, vibration, humidity, and airborne particulates (dust, fibers, and/or fumes).


Large PLC systems consist of a rack into which circuit “cards” are plugged. These cards include processors, input and output (I/O) points, communications ports, and other functions necessary to the operation of a complete PLC system. Such “modular” PLCs may be configured differently according to the specific needs of the application. Individual card failures are also easier to repair in a modular system, since only the failed card need be replaced, not all the cards or the whole card rack.

Small PLC systems consist of a monolithic “brick” containing all processor, I/O, and communication functions. These PLCs are typically far less expensive than their modular cousins, but are also more limited in I/O capability and must be replaced as a whole in the event of failure.

The following photographs show several examples of real PLC systems, some modular and some monolithic. These selections are not comprehensive by any means, as there are many more manufacturers and models of PLC than those I have photographed. They do, however, represent some of the more common brands and models in current (2010) industrial use.

¹There are such things as *soft PLCs*, which consist of special-purpose software running on an ordinary personal computer (PC) with some common operating system. Soft PLCs enjoy the high speed and immense memory capacity of modern personal computers, but do not possess the same ruggedness either in hardware construction or in operating system design. Their applications should be limited to non-critical controls where neither main process production nor safety would be jeopardized by a control system failure.

The first photograph is of a Siemens (Texas Instruments) 505 series PLC, installed in a control panel of a municipal wastewater treatment plant. This is an example of a modular PLC, with individual processor, I/O, and communication cards plugged into a rack. Three racks appear in this photograph (two completely filled with cards, and the third only partially filled):




The power supply and processor card for each rack is located on the left-hand end, with I/O cards plugged into slots in the rest of the rack. Input devices such as switches and sensors connect by wire to terminals on *input* cards, while output devices such as lamps, solenoids, and motor contactor coils connect by wire to terminals on *output* cards.

One of the benefits of modular PLC construction is that I/O cards may be changed out as desired, altering the I/O configuration of the PLC as needed. If, for example, the PLC needs to be configured to monitor a greater number of sensors, more input cards may be plugged into the rack and subsequently wired to those sensors. Or, if the *type* of sensor needs to be changed – perhaps from a 24 volt DC sensor to one operating on 120 volts AC – a different type of input card may be substituted to match the new sensor(s).

In this particular application, the PLC is used to sequence the operation of self-cleaning “trash racks” used to screen large debris such as rags, sticks, and other non-degradable items from municipal wastewater prior to treatment. These trash racks are actuated by electric motors, the captured debris scraped off and transported to a solid waste handling system. The motion of the trash racks, the sensing of wastewater levels and pressures, and the monitoring of any human-operated override controls are all managed by these PLCs. The programming of these PLCs involves timers, counters, sequencers, and other functions to properly manage the continuous operation of the trash racks.

The next photograph shows an Allen-Bradley (Rockwell) PLC-5 system, used to monitor and control the operation of a large natural gas compressor. Two racks appear in this first photograph, with different types of I/O cards plugged into each rack:



Like the Siemens 505 PLC seen previously, this Allen-Bradley PLC-5 system is fully modular and configurable. The types and locations of the I/O cards inserted into the rack may be altered by appropriately skilled technicians to suit any desired application. The programming of the PLC's processor card may also be altered if a change in the control strategy is desired for any reason.

In this particular application, the PLC is tasked with monitoring certain variables on the gas compressor unit, and taking corrective action if needed to keep the machine productive and safe. The automatic control afforded by the PLC ensures safe and efficient start-ups, shut-downs, and handling of emergency events. The networking and data-logging capability of the PLC ensures that critical data on the compressor unit may be viewed by the appropriate personnel. For this particular compressor station, the data gets communicated from Washington state where the compressor is located all the way to Utah state where the main operations center is located. Human operators in Utah are able to monitor the compressor's operating conditions and issue commands to the compressor over digital networks.

Both the Siemens (formerly Texas Instruments) 505 and Allen-Bradley (Rockwell) PLC-5 systems are considered "legacy" PLC systems by modern standards, the two systems in the previous photographs being about 20 years old each. It is not uncommon to find "obsolete" PLCs still in operation, though. Given their extremely rugged construction and reliable design, these control systems may continue to operate without significant trouble for decades.


A newer model of PLC manufactured by Allen-Bradley is the SLC 500 series (often verbally referred to as the “Slick 500”), which is also modular in design like the older PLC-5 system, although the racks and modules of the SLC 500 design are more compact. The SLC 500 rack shown in the next photograph has 7 “slots” for processor and I/O cards to plug in to, numbered 0 through 6 (left to right):



The first three slots of this particular SLC 500 rack (0, 1, and 2) are occupied by the processor card, an analog input card, and a discrete input card, respectively. The slots 3 and 4 are empty (revealing the backplane circuit board and connectors for accepting new cards). The slots 5 and 6 hold discrete output and analog output cards, respectively.


A feature visible on all cards in this system are numerous LED indicators, designed to show the status of each card. The processor card has LED indicators for “Run” mode, “Fault” conditions, “Force” conditions (when either input or output bits have been forced into certain states by the human programmer for testing purposes), and communication network indicators. Each discrete I/O card has indicator LEDs showing the on/off status of each I/O bit, and the analog card has a single LED showing that the card is powered.

A nine-slot SLC 500 system is shown in the next photograph, controlling a high-purity water treatment system for a biopharmaceuticals manufacturing facility. As you can see in this photograph, not all slots in this particular rack are occupied by I/O cards either:



Some of the inputs to this PLC include water level switches, pressure switches, water flow meters, and conductivity meters (to measure the purity of the water, greater electrical conductivity indicating the presence of more dissolved minerals, which is undesirable in this particular process application). In turn, the PLC controls the starting and stopping of water pumps and the switching of water valves to manage the water purification and storage processes.

A modern PLC manufactured by Siemens appears in this next photograph, an S7-300, which is a different design of modular PLC. Instead of individual cards plugging into a rack, this modular PLC design uses individual modules plugging into each other on their sides to form a wider unit:



A modern PLC manufactured by Allen-Bradley (Rockwell) is this ControlLogix 5000 system, shown in this photograph used to control a cereal manufacturing process. The modular design of the ControlLogix 5000 system follows the more traditional scheme of individual cards plugged into a rack of fixed size:




While the Siemens S7 and Rockwell ControlLogix PLC platforms represent large-scale, modular PLC systems, there exist much smaller PLCs available for a fraction of the cost. Perhaps the least expensive PLC on the market at this time of writing is the Koyo “CLICK” PLC series, the processor module (with eight discrete input and six discrete output channels built in) shown in my hand (sold for 69 US dollars in the year 2010, and with free programming software!):




This is a semi-modular PLC design, with a minimum of input/output (I/O) channels built into the processor module, but having the capacity to accept multiple I/O modules plugged in to the side, much like the Siemens S7-300 PLC.

Other semi-modular PLCs expand using I/O cards that plug in to the base unit not unlike traditional rack-based PLC systems. The Koyo DirectLogic DL06 is a good example of this type of semi-modular PLC, the following photograph showing a model DL06 accepting a thermocouple input card in one of its four available card slots:



This photograph shows the PLC base unit with 20 discrete input channels and 16 discrete output channels, accepting an analog input card (this particular card is designed to input signals from thermocouples to measure up to four channels of temperature).


Some low-end PLCs are strictly monolithic, with no ability to accept additional I/O modules. This General Electric Series One PLC (used to monitor a small-scale hydroelectric power generating station) is an example of a purely monolithic design, having no “expansion” slots to accept I/O cards:



A disadvantage of monolithic PLC construction is that damaged I/O cannot be independently replaced. If an I/O channel on one of these PLCs becomes damaged, the entire PLC must be replaced to fix the problem. In a modular system, the damaged I/O card may simply be unplugged from the rack and replaced with a new I/O card. Another disadvantage of monolithic PLCs is the inherently fixed nature of the I/O: the end-user cannot customize the I/O configuration to match the application. For these reasons, monolithic PLCs are usually found on small-scale processes with few I/O channels and limited potential for expansion.


12.2 Input/Output (I/O) capabilities

Every programmable logic controller must have some means of receiving and interpreting signals from real-world sensors such as switches, and encoders, and also be able to effect control over real-world control elements such as solenoids, valves, and motors. This is generally known as *input/output*, or *I/O*, capability. Monolithic (“brick”) PLCs have a fixed amount of I/O capability built into the unit, while modular (“rack”) PLCs use individual circuit board “cards” to provide customized I/O capability.



The advantages of using replaceable I/O cards instead of a monolithic PLC design are numerous. First, and most obvious, is the fact that individual I/O cards may be easily replaced in the event of failure without having to replace the entire PLC. Specific I/O cards may be chosen for custom applications, biasing toward discrete cards for applications using many on/off inputs and outputs, or biasing toward analog cards for applications using many 4-20 mA and similar signals. Some PLCs even offer the feature of *hot-swappable* cards, meaning each card may be removed and a new one inserted without de-energizing power to the PLC processor and rack. Please note that one should not assume any system has hot-swappable cards, because if you attempt to change out a card “live” in a system without this feature, you run the risk of damaging the card and/or the rest of the unit it is plugged in to!

Some PLCs have the ability to connect to processor-less remote racks filled with additional I/O cards or modules, thus providing a way to increase the number of I/O channels beyond the capacity of the base unit. The connection from host PLC to remote I/O racks usually takes the form of a special digital network, which may span a great physical distance:




An alternative scheme for system expansion is to network multiple PLCs together, where each PLC has its own dedicated rack and processor. Through the use of communication instructions, one PLC may be programmed to read data from and/or write data to another PLC, effectively using the other PLC as an extension of its own I/O. Although this method is more expensive than remote I/O (where the remote racks lack their own dedicated processors), it provides the capability of stand-alone control in the event the network connection between PLC processors becomes severed.

Input/output capability for programmable logic controllers comes in three basic varieties: *discrete*, *analog*, and *network*; each type discussed in a following subsection.

12.2.1 Discrete I/O


A “discrete” data point is one with only two states *on* and *off*. Process switches, pushbutton switches, limit switches, and proximity switches are all examples of discrete sensing devices. In order for a PLC to be aware of a discrete sensor’s state, it must receive a signal from the sensor through a *discrete input* channel. Inside each discrete input module is (typically) a set of light-emitting diodes (LEDs) which will be energized when the corresponding sensing device turns on. Light from each LED shines on a photo-sensitive device such as a phototransistor inside the module, which in turn activates a *bit* (a single element of digital data) inside the PLC’s memory. This opto-coupled arrangement makes each input channel of a PLC rather rugged, capable of isolating the sensitive computer circuitry of the PLC from transient voltage “spikes” and other electrical phenomena capable of causing damage:



Energizing an input channel lights the LED inside the optocoupler, turning on the phototransistor, sending a "high" signal to the PLC's microprocessor, setting (1) that bit in the PLC's input register.

The internal schematic diagram for a discrete input module (“card”) shown above reveals the componentry typical for a single input channel on that card. Each input channel has its own optocoupler, writing to its own unique memory register bit inside the PLC’s memory. Discrete input cards for PLCs typically have 4, 8, 16, or 32 channels.

Indicator lamps, solenoid valves, and motor starters (assemblies consisting of contactors and overload protection devices) are all examples of discrete control devices. In a manner similar to discrete inputs, a PLC connects to any number of different discrete final control devices through a *discrete output channel*². Discrete output modules typically use the same form of opto-isolation to allow the PLC's computer circuitry to send electrical power to loads: the internal PLC circuitry driving an LED which then activates some form of photosensitive switching device. Alternatively, small electromechanical relays may be used in lieu of opto-isolating semiconductor switching elements such as transistors (DC) or TRIACs (AC):



Setting a bit (1) in the PLC's output register sends a "high" signal to the LED inside the optocoupler, turning on the photo-TRIAC, sending AC power to the output channel to energize the load.


As with the schematic diagram for a discrete input module shown previously, the schematic diagram shown here for a discrete output module reveals the componentry typical for a single channel on that card. Each output channel has its own optocoupler, driven by its own unique memory register bit inside the PLC's memory. Discrete output cards for PLCs also typically have 4, 8, 16, or 32 channels.

An important concept to master when working with DC discrete I/O is the distinction between *current-sourcing* and *current-sinking* devices. The terms "sourcing" and "sinking" refer to the direction of current (as denoted by conventional flow notation) into or out of a device's control wire³. A device sending (conventional flow) current out of its control terminal to some other device(s) is said to be *sourcing* current, while a device accepting (conventional flow) current into its control terminal is said to be *sinking* current.

²I/O "channels" are often referred to as "points" in industry lingo. Thus, a "32-point input card" refers to an input circuit with 32 separate channels through which to receive signals from on/off switches and sensors.

³By "control wire," I mean the single conductor connecting the I/O card channel to the field device, as opposed to conductors directly common with either the positive or negative lead of the voltage source. If you focus your attention on this one wire, noting the direction of conventional-flow current through it, the task of determining whether a device is sourcing or sinking current becomes much simpler.


To illustrate, the following illustration shows a PLC output channel is *sourcing* current to an indicator lamp, which is *sinking* current to ground:



These terms really only make sense when electric current is viewed from the perspective of conventional flow, where the positive terminal of the DC power supply is envisioned to be the “source” of the current, with current finding its way “down” to ground (the negative terminal of the DC power supply). In every circuit formed by the output channel of a PLC driving a discrete control device, or by a discrete sensing device driving an input channel on a PLC, one element in the circuit must be sourcing current while the other is sinking current.

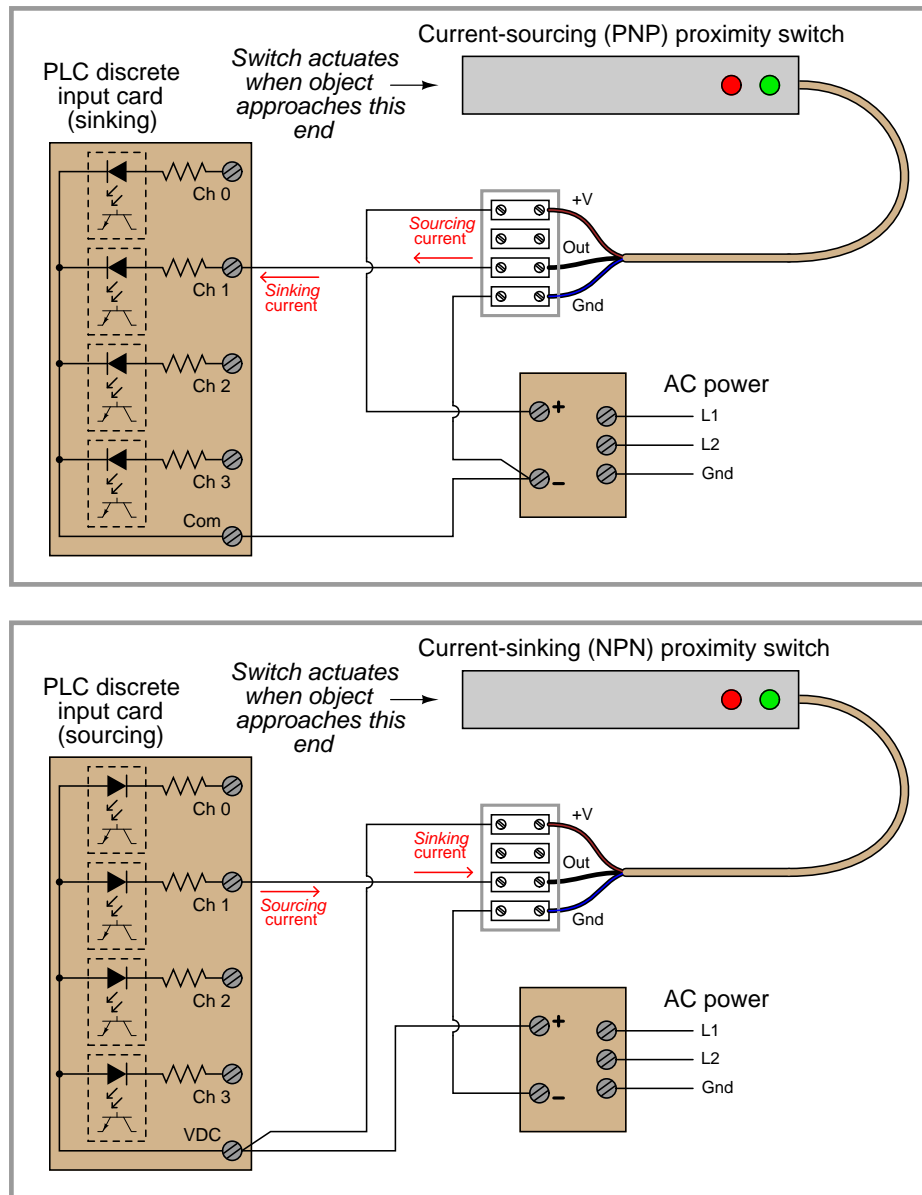
An engineering colleague of mine has a charming way to describe sourcing and sinking: *blowing* and *sucking*. A device that sources current to another “blows” current toward the other device. A device that sinks current “sucks” current from the other device. Many students seem to find these terms helpful in first mastering the distinction between sourcing and sinking despite (or perhaps because of!) their informal nature.

If the discrete device connecting to the PLC is not polarity-sensitive, either type of PLC I/O module will suffice. For example, the following diagrams show a mechanical limit switch connecting to a sinking PLC input and to a sourcing PLC input:




Note the differences in polarity and labeling between the sinking card's common terminal and the sourcing card's common terminal. On the "sinking" card, the input channel terminal is positive while the common ("Com") terminal is negative. On the "sourcing" card, the input channel terminal is negative while the common ("VDC") terminal is positive.

Some discrete sensing devices *are* polarity-sensitive, such as electronic proximity sensors containing transistor outputs. A “sourcing” proximity switch can only interface with a “sinking” PLC input channel, and vice-versa:



In all cases, the “sourcing” device sends current *out of* its signal terminal while the “sinking” device takes current *into* its signal terminal.

Two photographs of a DC (sinking) discrete input module for an Allen-Bradley model SLC 500 PLC are shown here: one with the plastic cover closed over the connection terminals, and the other with the plastic cover opened up for viewing the terminals. A legend on the inside of the cover shows the purpose of each screw terminal: eight input channels (numbered 0 through 7) and two redundant “DC Com” terminals for the negative pole of the DC power supply to connect:



A standard feature found on practically every PLC discrete I/O module is a set of LED indicators visually indicating the status of each bit (discrete channel). On the SLC 500 module, the LEDs appear as a cluster of eight numbered squares near the top of the module.


A photograph of discrete output terminals on another brand of PLC (a Koyo model DL06) shows somewhat different labeling:



Here, each output channel terminal is designated with a letter/number code beginning with the letter “Y”. Several “common” terminals labeled with “C” codes service clusters of output channels. In this particular case, each “common” terminal is common only to four output channels. With sixteen total output channels on this PLC, this means there are four different “common” terminals. While this may seem somewhat strange (why not just have one “common” terminal for all sixteen output channels?), it more readily permits different DC power supplies to service different sets of output channels.

Electrical polarity is not an issue with AC discrete I/O, since the polarity of AC reverses periodically anyway. However, there is still the matter of whether the “common” terminal on a discrete PLC module will connect to the *neutral* (grounded) or *hot* (ungrounded) AC power conductor.

The next photograph shows a discrete AC output module for an Allen-Bradley SLC 500 PLC, using TRIACs as power switching devices rather than transistors as is customary with DC discrete output modules:



This particular eight-channel module provides two sets of TRIACs for switching power to AC loads, each set of four TRIACs receiving AC power from a “hot” terminal (VAC 1 or VAC 2), the other side of the load device being connected to the “neutral” (grounded) conductor of the AC power source.

Fortunately, the hardware reference manual supplied by the manufacturer of every PLC shows diagrams illustrating how to connect discrete input and output channels to field devices. One should always consult these diagrams before connecting devices to the I/O points of a PLC!

12.2.2 Analog I/O

In the early days of programmable logic controllers, processor speed and memory were too limited to support anything but discrete (on/off) control functions. Consequently, the only I/O capability found on early PLCs were discrete in nature⁴. Modern PLC technology, though, is powerful enough to support the measurement, processing, and output of analog (continuously variable) signals.

All PLCs are digital devices at heart. Thus, in order to interface with an analog sensor or control device, some “translation” is necessary between the analog and digital worlds. Inside every analog input module is an *ADC*, or *Analog-to-Digital Converter*, circuit designed to convert an analog electrical signal into a multi-bit binary word. Conversely, every analog output module contains a *DAC*, or *Digital-to-Analog Converter*, circuit to convert the PLC’s digital command words into analog electrical quantities.

Analog I/O is commonly available for modular PLCs for many different analog signal types, including:

- Voltage (0 to 10 volt, 0 to 5 volt)
- Current (0 to 20 mA, 4 to 20 mA)
- Thermocouple (millivoltage)
- RTD (millivoltage)
- Strain gauge (millivoltage)

⁴Some modern PLCs such as the Koyo “CLICK” are also discrete-only. Analog I/O and processing is significantly more complex to engineer and more expensive to manufacture than discrete control, and so low-end PLCs are more likely to lack analog capability.

The following photographs show two analog I/O cards for an Allen-Bradley SLC 500 modular PLC system, an analog input card and an analog output card. Labels on the terminal cover doors indicate screw terminal assignments:



12.2.3 Network I/O

Many different digital network standards exist for PLCs to communicate with, from PLC to PLC and between PLCs and field devices. One of the earliest digital protocols developed for PLC communication was *Modbus*, originally for the Modicon brand of PLC. Modbus was adopted by other PLC and industrial device manufacturers as a *de facto* standard⁵, and remains perhaps the most universal digital protocol available for industrial digital devices today.

Another digital network standard developed by a particular manufacturer and later adopted as a *de facto* standard is *Profibus*, originally developed by Siemens.

For more information on networking standards used in PLC systems, refer to the “Digital electronic instrumentation” chapter, specifically sections on specific network standards such as Modbus and Profibus.

⁵A “de facto” standard is one arising naturally out of legacy rather than by an premeditated agreement between parties. Modbus and Profibus networks are considered “de facto” standards because those networks were designed, built, and marketed by pioneering firms prior to their acceptance as standards for others to conform to. In Latin, *de facto* means “from the fact,” which in this case refers to the fact of pre-existence: a standard agreed upon to conform to something already in popular use. By contrast, a standard intentionally agreed upon before its physical realization is a *de jure* standard (Latin for “from the law”). FOUNDATION Fieldbus is an example of a *de jure* standard, where a committee arrives at a consensus for a network design and specifications prior to that network being built and marketed by any firm.

12.3 Logic programming

Although it seems each model of PLC has its own idiosyncratic standard for programming, there does exist an international standard for controller programming that most PLC manufacturers at least attempt to conform to. This is the IEC 61131-3 standard, which will be the standard presented in this chapter.

One should take solace in the fact that despite differences in the details of PLC programming from one manufacturer to another and from one model to another, the basic principles are largely the same. There exist much greater disparities between different general-purpose programming languages (e.g. C/C++, BASIC, FORTRAN, Pascal, Java, Ada, etc.) than between the programming languages supported by different PLCs, and this fact does not prevent computer programmers from being “multilingual.” I have personally written and/or analyzed programs for over a half-dozen different manufacturers of PLCs (Allen-Bradley, Siemens, Square D, Koyo, Fanuc, Moore Products APACS and QUADLOG, and Modicon), with multiple PLC models within most of those brands, and I can tell you the differences in programming conventions are largely insignificant. After learning how to program one model of PLC, it is quite easy to adapt to programming other makes and models of PLC. If you are learning to program a particular PLC that does not exactly conform to the IEC 61131-3 standard, you will still be able to apply every single principle discussed in this chapter – the fundamental concepts are truly that universal.

The IEC 61131-3 standard specifies five distinct forms of programming language for industrial controllers:

- Ladder Diagram (LD)
- Structured Text (ST)
- Instruction List (IL)
- Function Block Diagram (FBD)
- Sequential Function Chart (SFC)

Not all programmable logic controllers support all five language types, but nearly all of them support Ladder Diagram (LD), which will be the primary focus of this book.

Programming languages for many industrial devices are limited by design. One reason for this is *simplicity*: any programming language simple enough in structure for someone with no formal computer programming knowledge to understand is going to be limited in its capabilities. Another reason for programming limitations is *safety*: the more flexible and unbounded a programming language is, the more potential there will be to unintentionally create complicated “run-time” errors when programming. The ISA safety standard number 84 classifies industrial programming languages as either *Fixed Programming Languages* (FPL), *Limited Variability Languages* (LVL), or *Full Variability Languages* (FVL). Ladder Diagram and Function Block Diagram programming are both considered to be “limited variability” languages, whereas Instruction List (and traditional computer programming languages such as C/C++, FORTRAN, BASIC, etc.) are considered “full variability” languages with all the attendant potential for complex errors.







12.3.1 Relating I/O status to virtual elements

Perhaps the most important yet elusive concept to grasp when learning to program PLCs is the relationship between the electrical status of the PLC's I/O points and the status of variables and other "elements" in its programming. This is especially true for Ladder Diagram (LD) programming, where the program itself resembles an electrical diagram. Making the mental connection between the "real" world of the switches, contactors, and other electrical devices connected to the PLC and the "imaginary" world of the PLC's program consisting of virtual contacts and relay "coils" is most fundamental.

The first fundamental rule one should keep in mind when examining a Ladder Diagram PLC program is that **each virtual contact shown in the program *actuates* whenever it reads a "1" state in its respective bit and will be *at rest* whenever it reads a "0" state in its respective bit** (in the PLC's memory). If the contact is a normally-open (NO) type, it will open when its bit is 0 and close when its bit is 1. If the contact is a normally-closed (NC) type, it will close when its bit is 0 and open when its bit is 1. A 0 bit state causes the contact to be in its "normal" (resting) condition, while a 1 bit state *actuates* the contact, forcing it into its non-normal (actuated) state.

Another rule to remember when examining a Ladder Diagram PLC program is that the programming software offers *color highlighting*⁶ to display the virtual status of each program element: **a colored contact is *closed*, while an un-colored contact is *open***. While the presence or absence of a "slash" symbol marks the *normal* status of a contact, its live color highlighting shown by PLC programming software reveals the "conductive" status of the elements *in real time*.


The following table shows how the two types of contacts in a PLC's Ladder Diagram program respond to bit states, using red coloring to signify each contact's virtual conductivity:

		<i>Contact type</i>		
				
		Normally-open	Normally-closed	
<i>Bit state</i>	0	 <i>(open)</i>	 <i>(closed)</i>	(Normal)
	1	 <i>(closed)</i>	 <i>(open)</i>	(Actuated)


Just as a pressure switch's contacts are actuated by a high pressure condition, and a level switch's contacts are actuated by a high level condition, and a temperature switch's contacts are actuated by a high temperature condition, so a PLC's virtual contact is actuated by a high *bit* condition (1). In the context of any switch, an *actuated* condition is the opposite of its *normal* (resting) condition.

⁶It should be noted that in some situations the programming software will fail to color the contacts properly, especially if their status changes too quickly for the software communication link to keep up, and/or if the bit(s) change state multiple times within one scan of the program. However, for simple programs and situations, this rule holds true and is a great help to beginning programmers as they learn the relationship between real-world conditions and conditions within the PLC's "virtual" world.

The following simplified⁷ illustration shows a small PLC with two of its discrete input channels electrically energized, causing those two bits to have “1” statuses. The color-highlighted contacts in the programming editor software’s display shows a collection of contacts addressed to those input bits in various states (colored = closed ; un-colored = open). As you can see, every contact addressed to a “set” bit (1) is in its actuated state, while every contact addressed to a “cleared” bit (0) is in its normal state:



(Displayed on the editing software)



Remember that a *colored* contact is a *closed* contact. The contacts appearing as colored are either normally-closed contacts with “0” bit states, or normally-open contacts with “1” bit states. It is the combination of bit state and contact type (NO vs. NC) that determines whether the virtual contact will be open (un-colored) or closed (colored) at any given time. Correspondingly, it is a combination of colored highlighting and virtual contact type that indicates the real-world energization status of a particular PLC input at any given time.

⁷The electrical wiring shown in this diagram is incomplete, with the “Common” terminal shown unconnected for simplicity’s sake.

In my teaching experience, the main problem students have comprehending PLC Ladder Diagram programs is that they over-simplify and try to directly associate real-world switches connected to the PLC with their respective contact instructions inside the PLC program. Students mistakenly think the real-world switch connecting to the PLC and the respective virtual switch contact inside the PLC program are one and the same, when this is not the case at all. Rather, the real-world switch sends power to the PLC input, *which in turn controls the state of the virtual contact(s) programmed into the PLC*. Specifically, I see students routinely fall into the following mis-conceptions:

- Students mistakenly think the contact instruction type (NO vs. NC) needs to match that of its associated real-world switch
- Students mistakenly think color highlighting of a contact instruction is equivalent to the electrical status of its associated real-world PLC input
- Students mistakenly think a closed real-world switch must result in a closed contact instruction in the live PLC program

To clarify, here are the fundamental rules one should keep in mind when interpreting contact instructions in Ladder Diagram PLC programs:


- **Each input bit in the PLC's memory will be a "1" when its input channel is powered, and will be a "0" when its input channel is unpowered**
- **Each virtual contact shown in the program *actuates* whenever it reads a "1" state in its respective bit, and will be *at rest* whenever it reads a "0" state in its respective bit**
- **A colored contact is *closed* (passes virtual power in the PLC program), while an un-colored contact is *open* (blocks virtual power in the PLC program)**

In trying to understand PLC Ladder Diagram programs, the importance of these rules cannot be overemphasized. The truth of the matter is a causal chain – rather than a direct equivalence – between the real-world switch and the contact instruction status. The real-world switch controls whether or not electrical power reaches the PLC input channel, which in turn controls whether the input register bit will be a "1" or a "0", which in turn controls whether the contact instruction will actuated or at rest. Virtual contacts inside the PLC program are thus *controlled* by their corresponding real-world switches, rather than simply being *identical* to their real-world counterparts as novices tend to assume. Following these rules, we see that normally-open (NO) contact instructions will mimic what their real-world switches are doing, while normally-closed (NC) contact instructions will act opposite of their real-world counterparts.

The color highlighting of *coil* instructions in a Ladder Diagram PLC program follows similar rules. A coil will be "on" (colored) when all contact instructions prior to it are closed (colored). A colored coil writes a "1" to its respective bit in memory, while an un-colored coil instruction writes a "0" to its respective bit in memory. If these bits are associated with real-world discrete output channels on the PLC, their states will control the real-world energization of devices electrically connected to those channels.

To further illuminate these fundamental concepts, we will examine the operation of a simple PLC system designed to energize a warning lamp in the event that a process vessel experiences a high fluid pressure. The PLC's task is to energize a warning lamp if the process vessel pressure ever exceeds 270 PSI, and keep that warning lamp energized even if the pressure falls below the trip point of 270 PSI. This way, operators will be alerted to both *past* and *present* process vessel overpressure events.

120 volt AC "line" power (L1 and L2) provides electrical energy for the PLC to operate, as well as signal potential for the input switches and power for the warning lamp. Two switches connect to the input of this PLC: one normally-open pushbutton switch acting as the alarm reset (pressing this switch "unlatches" the alarm lamp) and one normally-open pressure switch acting as the sensing element for high process vessel pressure:




The reset pushbutton connects to discrete input X1 of the PLC, while the pressure switch connects to discrete input X4. The warning lamp connects to discrete output Y5. Red indicator LEDs next to each I/O terminal visually indicate the electrical status of the I/O points, while red-shaded

highlighting shows the *virtual power*⁸ status of the “contacts” and “coils” in the PLC’s program, displayed on the screen of a personal computer connected to the PLC through a programming cable.

With no one pressing the reset pushbutton, that switch will be in its normal status, which for a “normally-open” switch is open. Likewise with the pressure switch: with process pressure less than the trip point of 270 PSI, the pressure switch will also be in its normal status, which for a “normally-open” switch is open. Since neither switch is conducting electricity right now, neither discrete input X1 nor X4 will be energized. This means the “virtual” contacts inside the PLC program will likewise be in their own normal states. Thus, any virtual contact drawn as a normally-open will be open (not passing virtual power), and any virtual contact drawn as a normally-closed (a diagonal slash mark through the contact symbol) will be closed. This is why the two normally-open virtual contacts X4 and Y5 have no highlighting, but the normally-closed virtual contact X1 does – the colored highlighting representing the ability to pass virtual power.


⁸For a PLC program contact, the shading represents virtual “conductivity.” For a PLC program coil, the shading represents a set (1) bit.

If the process vessel experiences a high pressure (> 270 PSI), the pressure switch will actuate, closing its normally-open contact. This will energize input $X4$ on the PLC, which will “close” the virtual contact $X4$ in the ladder program. This sends virtual power to the virtual “coil” $Y5$, which in turn latches itself on through virtual contact $Y5$ and also energizes the real discrete output $Y5$ to energize the warning lamp:




⁹It is worth noting the legitimacy of referencing virtual contacts to output bits (e.g. contact Y5), and not just to input bits. A “virtual contact” inside a PLC program is nothing more than an instruction to the PLC’s processor to *read* the status of a bit in memory. It matters not whether that bit is associated with a physical input channel, a physical output channel, or some abstract bit in the PLC’s memory. It would, however, be wrong to associate a virtual coil with an input bit, as coil instructions *write* bit values to memory, and input bits are supposed to be controlled solely by the energization states of their physical input channels.

If now the process pressure falls below 270 PSI, the pressure switch will return to its normal state (open), thus de-energizing discrete input X4 on the PLC. Because of the latching contact Y5 in the PLC's program, however, output Y5 remains on to keep the warning lamp in its energized state:



Thus, the Y5 contact performs a *seal-in* function to keep the Y5 bit set (1) even after the high-pressure condition clears. This is precisely the same concept as the "seal-in" auxiliary contact on a hard-wired motor starter circuit, where the electromechanical contactor keeps itself energized after the "Start" pushbutton switch has been released.

The only way for a human operator to re-set the warning lamp is to press the pushbutton. This will have the effect of energizing input X1 on the PLC, thus opening virtual contact X1 (normally-closed) in the program, thus interrupting virtual power to the virtual coil Y5, thus powering down the warning lamp and un-latching virtual power in the program:



12.3.2 Memory maps and I/O addressing

A wise PLC programmer once told me that the first thing any aspiring programmer should learn about the PLC they intend to program is how the digital memory of that PLC is organized. This is sage advice for any programmer, especially on systems where memory is limited, and/or where I/O has a fixed association with certain locations in the system's memory. Virtually every microprocessor-based control system comes with a published *memory map* showing the organization of its limited memory: how much is available for certain functions, which addresses are linked to which I/O points, how different locations in memory are to be referenced by the programmer.

Discrete input and output channels on a PLC correspond to individual *bits* in the PLC's memory array. Similarly, analog input and output channels on a PLC correspond to multi-bit *words* (contiguous blocks of bits) in the PLC's memory. The association between I/O points and memory locations is by no means standardized between different PLC manufacturers, or even between different PLC models designed by the same manufacturer. This makes it difficult to write a general tutorial on PLC addressing, and so my ultimate advice is to consult the engineering references for the PLC system you intend to program.

The most common brand of PLC in use in the United States at the time of this writing (2010) is Allen-Bradley (Rockwell), which happens to use a unique form of I/O addressing¹⁰ students tend to find confusing. For these two reasons (popularity and confusion), I will focus on Allen-Bradley addressing conventions for the bulk of this section.

¹⁰The most modern Allen-Bradley PLCs have all but done away with fixed-location I/O addressing, opting instead for *tag name* based I/O addressing. However, enough legacy Allen-Bradley PLC systems still exist in industry to warrant coverage of these addressing conventions.

The following table shows a partial memory map for an Allen-Bradley SLC 500 PLC¹¹:


File number	File type	Logical address range
0	Output image	O:0 to O:30
1	Input image	I:0 to I:30
2	Status	S:0 to S:n
3	Binary	B3:0 to B3:255
4	Timers	T4:0 to T4:255
5	Counters	C5:0 to C5:255
6	Control	R6:0 to R6:255
7	Integer	N7:0 to N7:255
8	Floating-point	F8:0 to F8:255
9	Network	x9:0 to x9:255
10 through 255	User defined	x10:0 to x255:255

Note that Allen-Bradley's use of the word "file" differs from personal computer parlance. In the SLC 500 controller, a "file" is a block of random-access memory used to store a particular type of data. By contrast, a "file" in a personal computer is a contiguous collection of data bits with collective meaning (e.g. a word processing file or a spreadsheet file), usually stored on the computer's hard disk drive. Within each of the Allen-Bradley PLC's "files" are multiple "elements," each element consisting of a set of bits (8, 16, 24, or 32) representing data. Elements are addressed by number following the colon after the file designator, and individual bits within each element addressed by a number following a slash mark. For example, the first bit (bit 0) of the second element in file 3 (Binary) would be addressed as B3:2/0.

In Allen-Bradley PLCs such as the SLC 500 and PLC-5 models, files 0, 1, and 2 are exclusively reserved for discrete outputs, discrete inputs, and status bits, respectively. Thus, the letter designators O, I, and S (file types) are redundant to the numbers 0, 1, and 2 (file numbers). Other file types such as B (binary), T (timers), C (counters), and others have their own default file numbers (3, 4, and 5, respectively), but may also be used in some of the user-defined file numbers (10 and above). For example, file 7 in an Allen-Bradley controller is reserved for data of the "integer" type (N), but integer data may also be stored in any file numbered 10 or greater at the user's discretion. Thus, file numbers and file type letters for data types other than output (O), input (I), and status (S) always appear together. You would not typically see an integer word addressed as N:30 (integer word 30 in the PLC's memory) for example, but rather as N7:30 (integer word 30 *in file 7* of the PLC's memory) to distinguish it from other integer word 30's that may exist in other files of the PLC's memory.

¹¹Also called the *data table*, this map shows the addressing of memory areas reserved for programs entered by the user. Other areas of memory exist within the SLC 500 processor, but these other areas are inaccessible to the technician writing PLC programs.

This file-based addressing notation bears further explanation. When an address appears in a PLC program, special characters are used to separate (or “delimit”) different fields from each other. The general scheme for Allen-Bradley SLC 500 PLCs is shown here:



Not all file types need to distinguish individual words and bits. Integer files (N), for example, consist of one 16-bit word for each element. For instance, N7:5 would be the 16-bit integer word number five held in file seven. A discrete input file type (I), though, needs to be addressed as individual bits because each separate I/O point refers to a single bit. Thus, I:3/7 would be bit number seven residing in input element three. The “slash” symbol is necessary when addressing discrete I/O bits because we do not wish to refer to all sixteen bits in a word when we just mean a single input or output point on the PLC. Integer numbers, by contrast, are collections of 16 bits each in the SLC 500 memory map, and so are usually addressed as entire words rather than bit-by-bit¹².

Certain file types such as timers are more complex. Each timer “element¹³” consists of *two* different 16-bit words (one for the timer’s accumulated value, the other for the timer’s target value) in addition to no less than *three* bits declaring the status of the timer (an “Enabled” bit, a “Timing” bit, and a “Done” bit). Thus, we must make use of both the decimal-point and slash separator symbols when referring to data within a timer. Suppose we declared a timer in our PLC program with the address T4:2, which would be timer number two contained in timer file four. If we wished to address that timer’s current value, we would do so as T4:2.ACC (the “Accumulator” word of timer number two in file four). The “Done” bit of that same timer would be addressed as T4:2/DN (the “Done” bit of timer number two in file four)¹⁴.


¹²This is not to say one *cannot* specify a particular bit in an otherwise whole word. In fact, this is one of the powerful advantages of Allen-Bradley’s addressing scheme: it gives you the ability to precisely specify portions of data, even if that data is not generally intended to be portioned into smaller pieces!

¹³Programmers familiar with languages such as C and C++ might refer to an Allen-Bradley “element” as a *data structure*, each type with a set configuration of words and/or bits.

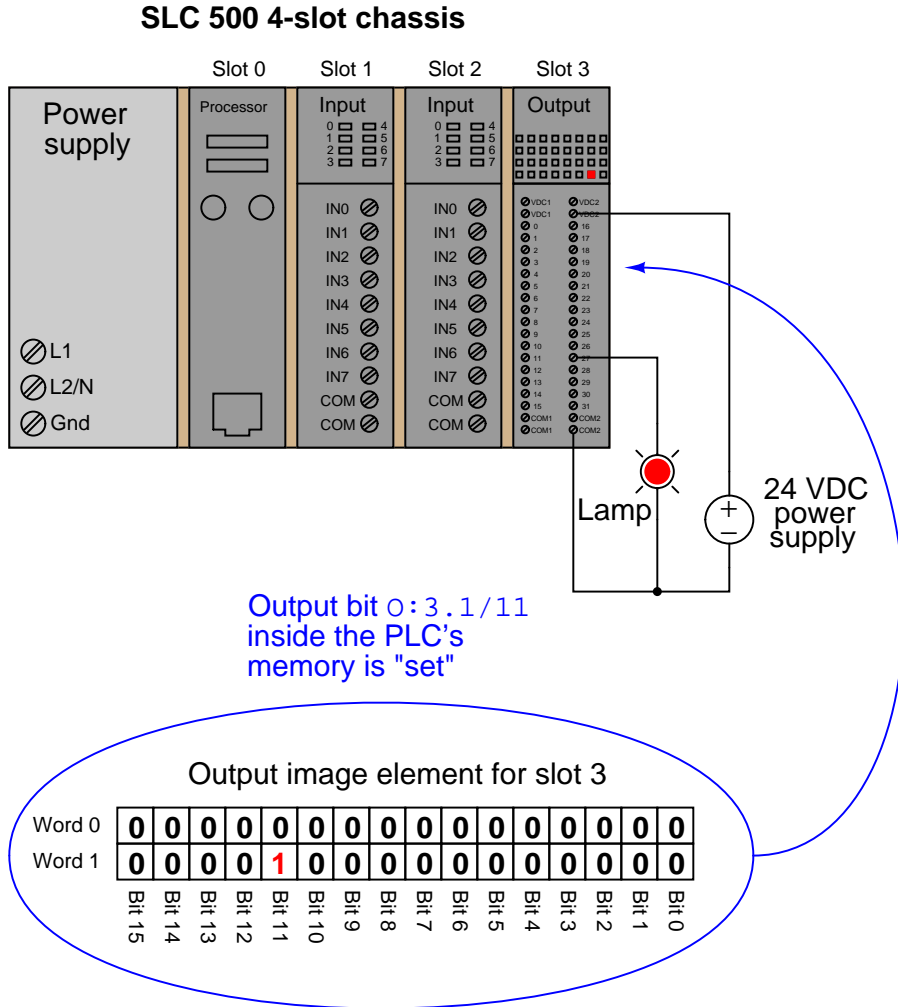
¹⁴Referencing the Allen-Bradley engineering literature, we see that the accumulator word may alternatively be addressed by number rather than by mnemonic, T4:2.2 (word 2 being the accumulator word in the timer data structure), and that the “done” bit may be alternatively addressed as T4:2.0/13 (bit number 13 in word 0 of the timer’s data structure). The mnemonics provided by Allen-Bradley are certainly less confusing than referencing word and bit numbers for particular aspects of a timer’s function!

A hallmark of the SLC 500's addressing scheme common to many legacy PLC systems is that the address labels for input and output bits explicitly reference the physical locations of the I/O channels. For instance, if an 8-channel discrete input card were plugged into slot 4 of an Allen-Bradley SLC 500 PLC, and you wished to specify the second bit (bit 1 out of a 0 to 7 range), you would address it with the following label: I:4/1. Addressing the seventh bit (bit number 6) on a discrete output card plugged into slot 3 would require the label O:3/6. In either case, the numerical structure of that label tells you exactly where the real-world input signal connects to the PLC.

To illustrate the relationship between physical I/O and bits in the PLC's memory, consider this example of an Allen-Bradley SLC 500 PLC, showing one of its discrete input channels energized (the switch being used as a "Start" switch for an electric motor):



If an input or output card possesses more than 16 bits – as in the case of the 32-bit discrete output card shown in slot 3 of the example SLC 500 rack – the addressing scheme further subdivides each element into *words* and bits (each “word” being 16 bits in length). Thus, the address for bit number 27 of a 32-bit input module plugged into slot 3 would be I:3.1/11 (since bit 27 is equivalent to bit 11 of word 1 – word 0 addressing bits 0 through 15 and word 1 addressing bits 16 through 31):



A close-up photograph of a 32-bit DC input card for an Allen-Bradley SLC 500 PLC system shows this multi-word addressing:



The first sixteen input points on this card (the left-hand LED group numbered 0 through 15) are addressed I:X.0/0 through I:X.0/15, with “X” referring to the slot number the card is plugged into. The next sixteen input points (the right-hand LED group numbered 16 through 31) are addressed I:X.1/0 through I:X.1/15.

Legacy PLC systems typically reference each one of the I/O channels by labels such as “I:1/3” (or equivalent¹⁵) indicating the actual location of the input channel terminal on the PLC unit. The IEC 61131-3 programming standard refers to this channel-based addressing of I/O data points as *direct addressing*. A synonym for direct addressing is *absolute addressing*.

Addressing I/O bits directly by their card, slot, and/or terminal labels may seem simple and elegant, but it becomes very cumbersome for large PLC systems and complex programs. Every time a technician or programmer views the program, they must “translate” each of these I/O labels to some real-world device (e.g. “Input I:1/3 is actually the *Start* pushbutton for the middle tank mixer motor”) in order to understand the function of that bit. A later effort to enhance the clarity of PLC programming was the concept of addressing variables in a PLC’s memory by arbitrary names rather than fixed codes. The IEC 61131-3 programming standard refers to this as *symbolic addressing* in contrast to “direct” (channel-based) addressing, allowing programmers arbitrarily

¹⁵Some systems such as the Texas Instruments 505 series used “X” labels to indicate discrete input channels and “Y” labels to indicate discrete output channels (e.g. input X9 and output Y14). This same labeling convention is still used by Koyo in its DirectLogic and “CLICK” PLC models. Siemens continues a similar tradition of I/O addressing by using the letter “I” to indicate discrete inputs and the letter “Q” to indicate discrete outputs (e.g. input channel I0.5 and output Q4.1).

name I/O channels in ways that are meaningful to the system as a whole. To use our simple motor “Start” switch example, it is now possible for the programmer to designate input I:1/3 (an example of a *direct address*) as “Motor_start_switch” (an example of a *symbolic address*) within the program, thus greatly enhancing the readability of the PLC program. Initial implementations of this concept maintained direct addresses for I/O data points, with symbolic names appearing as supplements to the absolute addresses.

The modern trend in PLC addressing is to avoid the use of direct addresses such as I:1/3 altogether, so they do not appear anywhere in the programming code. The Allen-Bradley “Logix” series of programmable logic controllers is the most prominent example of this new convention at the time of this writing. Each I/O point, regardless of type or physical location, is assigned a *tag name* which is meaningful in a real-world sense, and these tag names (or *symbols* as they are alternatively called) are referenced to absolute I/O channel locations by a database file. An important requirement of tag names is that they contain no space characters between words (e.g. instead of “Motor start switch”, a tag name should use hyphens or underscore marks as spacing characters: “Motor_start_switch”), since spaces are generally assumed by computer programming languages to be delimiters (separators between different variables).


Having introduced Allen-Bradley’s addressing notation for SLC 500 model PLCs, I will now abandon it in favor of the modern convention of symbolic addressing throughout the rest of this chapter, so as to avoid making the programming examples brand- or model-specific. Each data point within my PLC programs will bear its own tag name rather than a direct (channel-based) address label.

12.4 Ladder Diagram (LD) programming

In the United States, the most common language used to program PLCs is *Ladder Diagram* (LD), also known as *Relay Ladder Logic* (RLL). This is a graphical language showing the logical relationships between inputs and outputs as though they were contacts and coils in a hard-wired electromechanical relay circuit. This language was invented for the express purpose of making PLC programming feel “natural” to electricians familiar with relay-based logic and control circuits. While Ladder Diagram programming has many shortcomings, it remains extremely popular and so will be the primary focus of this chapter.

Every Ladder Diagram program is arranged to resemble an electrical diagram, making this a graphical (rather than text-based) programming language. Ladder diagrams are to be thought of as *virtual circuits*, where virtual “power” flows through virtual “contacts” (when closed) to energize virtual “relay coils” to perform logical functions. None of the contacts or coils seen in a Ladder Diagram PLC program are real; rather, they act on bits in the PLC’s memory, the logical inter-relationships between those bits expressed in the form of a diagram *resembling* a circuit.

The following computer screenshot shows a typical Ladder Diagram program¹⁶ being edited on a personal computer:



Contacts appear just as they would in an electrical relay logic diagram – as short vertical line segments separated by a horizontal space. Normally-open contacts are empty within the space between the line segments, while normally-closed contacts have a diagonal line crossing through that space. Coils are somewhat different, appearing as either circles or pairs of parentheses. Other instructions appear as rectangular boxes.

Each horizontal line is referred to as a *rung*, just as each horizontal step on a stepladder is called a “rung.” A common feature among Ladder Diagram program editors, as seen on this screenshot, is

¹⁶This particular program and editor is for the Koyo “CLICK” series of micro-PLCs.

the ability to color-highlight those virtual “components” in the virtual “circuit” ready to “conduct” virtual “power.” In this particular editor, the color used to indicate “conduction” is light blue. Another form of status indication seen in this PLC program is the values of certain variables in the PLC’s memory, shown in red text.

For example, you can see coil T2 energized at the upper-right corner of the screen (filled with light blue coloring), while coil T3 is not. Correspondingly, each normally-open T2 contact appears colored, indicating its “closed” status, while each normally-closed T2 contact is uncolored. By contrast, each normally-open T3 contact is uncolored (since coil T3 is unpowered) while each normally-closed T3 contact is shown colored to indicate its conductive status. Likewise, the current count values of timers T2 and T3 are shown as 193 and 0, respectively. The output value of the math instruction box happens to be 2400, also shown in red text.

Color-highlighting of Ladder Diagram components only works, of course, when the computer running the program editing software is connected to the PLC and the PLC is in the “run” mode (and the “show status” feature of the editing software is enabled). Otherwise, the Ladder Diagram is nothing more than black symbols on a white background. Not only is status highlighting very useful in de-bugging PLC programs, but it also serves an invaluable diagnostic purpose when a technician analyzes a PLC program to check the status of real-world input and output devices connected to the PLC. This is especially true when the program’s status is viewed remotely over a computer network, allowing maintenance staff to investigate system problems without even being near the PLC!


12.4.1 Contacts and coils

The most elementary objects in Ladder Diagram programming are *contacts* and *coils*, intended to mimic the contacts and coils of electromechanical relays. Contacts and coils are *discrete* programming elements, dealing with Boolean (1 and 0; on and off; true and false) variable states. Each contact in a Ladder Diagram PLC program represents the *reading* of a single bit in memory, while each coil represents the *writing* of a single bit in memory.

Discrete input signals to the PLC from real-world switches are read by a Ladder Diagram program by contacts referenced to those input channels. In legacy PLC systems, each discrete input channel has a specific address which must be applied to the contact(s) within that program. In modern PLC systems, each discrete input channel has a tag name created by the programmer which is applied to the contact(s) within the program. Similarly, discrete output channels – referenced by coil symbols in the Ladder Diagram – must also bear some form of address or tag name label.

To illustrate, we will imagine the construction and programming of a redundant flame-sensing system to monitor the status of a burner flame using three sensors. The purpose of this system will be to indicate a “lit” burner if at least two out of the three sensors indicate flame. If only one sensor indicates flame (or if no sensors indicate flame), the system will declare the burner to be un-lit. The burner’s status will be visibly indicated by a lamp that human operators can readily see inside the control room area.

Our system's wiring is shown in the following diagram:




Each flame sensor outputs a DC voltage signal indicating the detection of flame at the burner, either on (24 volts DC) or off (0 volts DC). These three discrete DC voltage signals are sensed by the first three channels of the PLC's discrete input card. The indicator lamp is a 120 volt light bulb, and so must be powered by an AC discrete output card, shown here in the PLC's last slot.

To make the ladder program more readable, we will assign tag names (symbolic addresses) to each input and output bit in the PLC, describing its real-world device in an easily-interpreted format¹⁷. We will tag the first three discrete input channels as `IN_sensor_A`, `IN_sensor_B`, and `IN_sensor_C`, and the output as `OUT_burner_lit`.

¹⁷If this were a legacy Allen-Bradley PLC system using absolute addressing, we would be forced to address the three sensor inputs as `I:1/0`, `I:1/1`, and `I:1/2` (slot 1, channels 0 through 2), and the indicator lamp output as `O:2/0` (slot 2, channel 0). If this were a newer Logix5000 Allen-Bradley PLC, the default tag names would be `Local:1:I.Data.0`, `Local:1:I.Data.1`, and `Local:1:I.Data.2` for the three inputs, and `Local:2:O.Data.0` for the output. However, in either system we have the ability to assign symbolic addresses so we have a way to reference the I/O channels without having to rely on these cumbersome labels. The programs showing in this book exclusively use tag names rather than absolute addresses, since this is the more modern programming convention.

A ladder program to determine if at least two out of the three sensors detect flame is shown here, with the tag names referencing each contact and coil:




Series-connected contacts in a Ladder Diagram perform the logical AND function, while parallel contacts perform the logical OR function. Thus, this two-out-of-three flame-sensing program could be verbally described as:

“Burner is lit if either A and B, or either B and C, or either A and C”


An alternate way to express this is to use the notation of *Boolean algebra*, where multiplication represents the AND function and addition represents the OR function:

$$\text{Burner_lit} = AB + BC + AC$$

Yet another way to represent this logical relationship is to use logic gate symbols:



To illustrate how this program would work, we will consider a case where flame sensors B and C detect flame, but sensor A does not¹⁸. This represents a two-out-of-three-good condition, and so we would expect the PLC to turn on the “Burner lit” indicator light as programmed. From the perspective of the PLC’s rack, we would see the indicator LEDs for sensors B and C lit up on the discrete input card, as well as the indicator LED for the lamp’s output channel:




Those two energized input channels “set” bits (1 status) in the PLC’s memory representing the status of flame sensors B and C. Flame sensor A’s bit will be “clear” (0 status) because its corresponding input channel is de-energized. The fact that the output channel LED is energized (and the “Burner lit” indicator lamp is energized) tells us the PLC program has “set” that corresponding bit in the PLC’s output memory register to a “1” state.

¹⁸The most likely reason why one out of two flame sensors might not detect the presence of a flame is some form of misalignment or fouling of the flame sensor. In fact, this is a good reason for using a 2-out-of-3 flame detection system rather than a simplex (1-out-of-1) detector scheme: to make the system more tolerant of occasional sensor problems without compromising burner safety.

A display of input and output register bits shows the “set” and “reset” states for the PLC at this moment in time:

Input register						Output register									
0	0	0	0	0	1	1	0	0	0	0	0	0	0	1	
IN7	IN6	IN5	IN4	IN3	IN2	IN1	IN0	OUT7	OUT6	OUT5	OUT4	OUT3	OUT2	OUT1	OUT0

Examining the Ladder Diagram program with status indication enabled, we see how only the middle contact pair is passing “virtual power” to the output coil:




Recall that the purpose of a contact in a PLC program is to *read* the status of a bit in the PLC's memory. These six “virtual contacts” read the three input bits corresponding to the three flame sensors. Each normally-open “contact” will “close” when its corresponding bit has a value of 1, and will “open” (go to its normal state) when its corresponding bit has a value of 0. Thus, we see here that the two contacts corresponding to sensor A appear without highlighting (representing no “conductivity” in the virtual relay circuit) because the bit for that input is reset (0). The two contacts corresponding to sensor B and the two contacts corresponding to sensor C all appear highlighted (representing “conductivity” in the virtual circuit) because their bits are both set (1).

Recall also that the purpose of a coil in a PLC program is to *write* the status of a bit in the PLC's memory. Here, the “energized” coil sets the bit for the PLC output 0 to a “1” state, thus activating the real-world output and sending electrical power to the “Burner lit” lamp.

Note that the color highlighting does *not* indicate a virtual contact is *conducting* virtual power, but merely that it is *able* to conduct power. Color highlighting around a virtual coil, however, *does* indicate the presence of virtual “power” at that coil.

Contacts and relays are not just useful for implementing simple logic functions, but they may also perform *latching* functions as well. A very common application of this in industrial PLC systems is a latching start/stop program for controlling electric motors by means of momentary-contact pushbutton switches. As before, this functionality will be illustrated by means of an hypothetical example circuit and program:




In this system, two pushbutton switches are connected to discrete inputs on a PLC, and the PLC in turn energizes the coil of a motor contactor relay by means of one of its discrete outputs¹⁹.

¹⁹The particular input and output channels chosen for this example are completely arbitrary. There is no particular

An overload contact is wired directly in series with the contactor coil to provide motor overcurrent protection, even in the event of a PLC failure where the discrete output channel remains energized²⁰.

The ladder program for this motor control system would look like this:



Pressing the “Start” pushbutton energizes discrete input channel 6 on the PLC, which “closes” the virtual contact in the PLC program labeled `IN_switch_Start`. The normally-closed virtual contact for input channel 7 (the “Stop” pushbutton) is already closed by default when the “Stop” button is not being pressed, and so the virtual coil will receive “power” when the “Start” pushbutton is pressed and the “Stop” pushbutton is not.

Note the *seal-in* contact bearing the exact same label as the coil: `OUT_contactor`. At first it may seem strange to have both a contact and a coil in a PLC program labeled identically²¹, since contacts are most commonly associated with inputs and coils with outputs, but this makes perfect sense if you realize the true meaning of contacts and coils in a PLC program: as *read* and *write* operations on bits in the PLC’s memory. The coil labeled `OUT_contactor` *writes* the status of that bit, while the contact labeled `OUT_contactor` *reads* the status of that same bit. The purpose of this contact, of course, is to latch the motor in the “on” state after a human operator has released his or her finger from the “Start” pushbutton.


This programming technique is known as *feedback*, where an output variable of a function (in this case, the feedback variable is `OUT_contactor`) is also an input to that same function. The path of feedback is *implicit* rather than *explicit* in Ladder Diagram programming, with the only indication of feedback being the common name shared by coil and contact. Other graphical programming languages (such as Function Block) have the ability to show feedback paths as connecting lines between function outputs and inputs, but this capacity does not exist in Ladder Diagram.

reason to choose input channels 6 and 7, or output channel 2, as I have shown in the wiring diagram. Any available I/O channels will suffice.

²⁰While it is possible to wire the overload contact to one of the PLC’s discrete input channels and then program a *virtual* overload contact in series with the output coil to stop the motor in the event of a thermal overload, this strategy would rely on the PLC to perform a safety function which is probably better performed by hard-wired circuitry.

²¹A very common misconception among students first learning PLC Ladder Diagram programming is to always associate contacts with PLC inputs and coils with PLC outputs, thus it seems weird to have a contact bear the same label as an output. However, this is a false association. In reality, contacts and coils are *read* and *write* instructions, and thus it is possible to have the PLC read one of its own output bits as a part of some logic function. What *would* be truly strange is to label a coil with an input bit address or tag name, since the PLC is not electrically capable of setting the real-world energization status of any input channels.

A step-by-step sequence showing the operation and status of this simple program illustrates how the seal-in contact functions, through a start-up and shut-down cycle of the motor:




This sequence helps illustrate the *order of evaluation* or *scan order* of a Ladder Diagram program. The PLC reads a Ladder Diagram from left to right, top to bottom, in the same general order as

a human being reads sentences and paragraphs written in English. However, according to the IEC 61131-3 standard, a PLC program must evaluate (read) all inputs (contacts) to a function before determining the status of a function's output (coil or coils). In other words, the PLC does not make any decision on how to set the state of a coil until all contacts providing power to that coil have been read. Once a coil's status has been written to memory, any contacts bearing the same tag name will update with that status on subsequent rungs in the program.

Step 5 in the previous sequence is particularly illustrative. When the human operator presses the "Stop" pushbutton, the input channel for `IN_switch_Stop` becomes activated, which "opens" the normally-closed virtual contact `IN_switch_Stop`. Upon the next scan of this program rung, the PLC evaluates all input contacts (`IN_switch_Start`, `IN_switch_Stop`, and `OUT_contactor`) to check their status before deciding what status to write to the `OUT_contactor` coil. Seeing that the `IN_switch_Stop` contact has been forced open by the activation of its respective discrete input channel, the PLC writes a "0" (or "False") state to the `OUT_contactor` coil. However, the `OUT_contactor` feedback contact does not update until the next scan, which is why you still see it color-highlighted during step 5.


A potential problem with this system as it is designed is that the human operator loses control of the motor in the event of an "open" wiring failure in either pushbutton switch circuit. For instance, if a wire fell off a screw contact for the "Start" pushbutton switch circuit, the motor could not be started if it was already stopped. Similarly, if a wire fell off a screw contact for the "Stop" pushbutton switch circuit, the motor could not be stopped if it was already running. In either case, a broken wire connection acts the same as the pushbutton switch's "normal" status, which is to keep the motor in its present state. In some applications, this failure mode would not be a severe problem. In many applications, though, it is quite dangerous to have a running motor that cannot be stopped. For this reason, it is customary to design motor start/stop systems a bit differently from what has been shown here.

In order to build a “fail-stop” motor control system with our PLC, we must first re-wire the pushbutton switch to use its normally-closed (NC) contact:



This keeps discrete input channel 7 activated when the pushbutton is unpressed. When the operator presses the “Stop” pushbutton, the switch’s contact will be forced open, and input channel 7 will de-energize. If a wire happens to fall off a screw terminal in the “Stop” switch circuit, input channel 7 will de-energize just the same as if someone pressed the “Stop” pushbutton, which will automatically shut off the motor.

In order for the PLC program to work properly with this new switch wiring, the virtual contact for `IN_switch_Stop` must be changed from a normally-closed (NC) to a normally-open (NO):



As before, the `IN_switch_Stop` virtual contact is in the “closed” state when no one presses the “Stop” switch, enabling the motor to start any time the “Start” switch is pressed. Similarly, the `IN_switch_Stop` virtual contact will open any time someone presses the “Stop” switch, thus stopping virtual “power” from flowing to the `OUT_contactor` coil.

Although this is a very common way to build PLC-controlled motor start/stop systems – with an NC pushbutton switch and an NO “Stop” virtual contact – students new to PLC programming often find this logical reversal confusing²². Perhaps the most common reason for this confusion is a mis-understanding of the “normal” concept for switch contacts, be they real or virtual. The `IN_switch_Stop` virtual contact is programmed to be normally-open (NO), but yet it is *typically* found in the closed state. Recall that the “normal” status of any switch is its status while in a resting condition of no stimulation, *not* necessarily its status while the process is in a “normal” operating mode. The “normally-open” virtual contact `IN_switch_Stop` is typically found in the closed state because its corresponding input channel is typically found energized, owing to the normally-closed pushbutton switch contact, which passes real electrical power to the input channel while no one presses the switch. Just because a switch is configured as normally-open does not necessarily mean it will be *typically* found in the open state! The status of any switch contact, whether real or virtual, is a function of its configuration (NO versus NC) and the stimulus applied to it.

Another concern surrounding real-world wiring problems is what this system will do if the motor contactor coil circuit opens for any reason. An open circuit may develop as a result of a wire falling off a screw terminal, or it may occur because the thermal overload contact tripped open due to an over-temperature event. The problem with our motor start/stop system as designed is that it is not “aware” of the contactor’s real status. In other words, the PLC “thinks” the contactor will be


²²In an effort to alleviate this confusion, the Allen-Bradley corporation (Rockwell) uses the terms *examine if closed* (XIC) and *examine if open* (XIO) to describe “normally open” and “normally closed” virtual contacts, respectively, in their Ladder Diagram programming. The idea here is that a virtual contact drawn as a normally-open symbol will be “examined” (declared “true”) by the PLC’s processor if its corresponding input channel is energized (powered by a real-life contact in the *closed* state). Conversely, a virtual contact drawn as a normally-closed symbol (with a slash mark through the middle) will be “examined” by the PLC’s processor if its corresponding input channel is de-energized (if the real-life contact sending power to that terminal is in the open state). In my experience, I have found this nomenclature to be even more confusing to students than simply calling these virtual contacts “normally open” and “normally closed” like other PLC manufacturers do. The foundational concept for students to grasp here is that *the virtual contact is not a direct representation of the real-life electrical switch contact – rather, it is a read instruction for the bit set by power coming from the real-life electrical switch contact.*

energized any time discrete output channel 2 is energized, but that may not actually be the case if there is an open fault in the contactor's coil circuit.

This may lead to a dangerous condition if the open fault in the contactor's coil circuit is later cleared. Imagine an operator pressing the "Start" switch but noticing the motor does not actually start. Wondering why this may be, he or she goes to look at the overload relay to see if it is tripped. If it is tripped, and the operator presses the "Reset" button on the overload assembly, the motor will immediately start because the PLC's discrete output has remained energized all the time following the pressing of the "Start" switch. Having the motor start up as soon as the thermal overload is reset may come as a surprise to operations personnel, and this could be quite dangerous if anyone happens to be near the motor-powered machinery when it starts.


What would be safer is a motor control system that refuses to "latch" on unless the contactor actually energizes when the "Start" switch is pressed. For this to be possible, the PLC must have some way of sensing the contactor's status.

In order to make the PLC “aware” of the contactor’s real status, we may connect the auxiliary switch contact to one of the unused discrete input channels on the PLC, like this:



Now, the PLC is able to sense the real-time status of the contactor via input channel 5.


We may modify the PLC program to recognize this status by assigning a new tag name to this input (`IN_contactor_aux`) and using a normally-open virtual contact of this name as the seal-in contact instead of the `OUT_contactor` bit:



Now, if the contactor fails to energize for any reason when the operator presses the “Start” switch, the PLC’s output will fail to latch when the “Start” switch is released. When the open fault in the contactor’s coil circuit is cleared, the motor will *not* immediately start up, but rather wait until the operator presses the “Start” switch again, which is a much safer operating characteristic than before.

A special class of virtual “coil” used in PLC ladder programming that bears mentioning is the “latching” coil. These usually come in two forms: a *set* coil and a *reset* coil. Unlike a regular “output” coil that positively writes to a bit in the PLC’s memory with every scan of the program, “set” and “reset” coils only write to a bit in memory when energized by virtual power. Otherwise, the bit is allowed to retain its last value.

A very simple motor start/stop program could be written with just two input contacts and two of these latching coils (both bearing the same tag name, writing to the same bit in memory):



Note the use of a normally-open (NO) pushbutton switch contact (again!), with no auxiliary contact providing status indication of the contactor to the PLC. This is a very minimal program, shown for the strict purpose of illustrating the use of “set” and “reset” latching coils in Ladder Diagram PLC programming.

“Set” and “Reset” coils²³ are examples of what is known in the world of PLC programming as *retentive instructions*. A “retentive” instruction *retains* its value after being virtually “de-energized” in the Ladder Diagram “circuit.” A standard output coil is *non-retentive*, which means it does not “latch” when de-energized. The concept of retentive and non-retentive instructions will appear again as we explore PLC programming, especially in the area of *timers*.


Ordinarily, we try to avoid multiple coils bearing the same label in a PLC Ladder Diagram program. With each coil representing a “write” instruction, multiple coils bearing the same name represents multiple “write” operations to the same bit in the PLC’s memory. Here, with latching coils, there is no conflict because each of the coils only writes to the `OUT_contactor` bit when its

²³Referred to as “Latch” and “Unlatch” coils by Allen-Bradley.

respective contact is energized. So long as only one of the pushbutton switches is actuated at a time, there is no conflict between the identically-named coils.

This raises the question: what would happen if *both* pushbutton switches were simultaneously pressed? What would happen if *both* “Set” and “Reset” coils were “energized” at the same time? The result is that the `OUT_contactor` bit would first be “set” (written to a value of 1) then “reset” (written to a value of 0) in that order as the two rungs of the program were scanned from top to bottom. PLCs typically do not typically update their discrete I/O registers while scanning the Ladder Diagram program (this operation takes place either before or after each program scan), so the real discrete output channel status will be whatever the *last* write operation told it to be, in this case “reset” (0, or off).

Even if the discrete output is not “confused” due to the conflicting write operations of the “Set” and “Reset” coils, other rungs of the program written between the “Set” and “Reset” rungs might be. Consider for example a case where there were other program rungs following the “Set” and “Reset” rungs reading the status of the `OUT_contactor` bit for some purpose. Those other rungs *would* indeed become “confused” because they would see the `OUT_contactor` bit in the “set” state while the actual discrete output of the PLC (and any rungs following the “Reset” rung) would see the `OUT_contactor` bit in the “reset” state:




Multiple (non-retentive) output coils with the same memory address are almost always a programming *faux pas* for this reason, but even retentive coils which are designed to be used in matched pairs can cause trouble if the implications of simultaneous energization are not anticipated. Multiple *contacts* with identical addresses are no problem whatsoever, because multiple “read” operations to the same bit in memory will never cause a conflict.


The IEC 61131-3 PLC programming standard specifies *transition-sensing* contacts as well as the more customary “static” contacts. A transition-sensing contact will “actuate” only for a duration of one program scan, even if its corresponding bit remains active. Two types of transition-sensing Ladder Diagram contacts are defined in the IEC standard: one for *positive* transitions and another

for *negative* transitions. The following example shows a wiring diagram, Ladder Diagram program, and a timing diagram demonstrating how each type of transition-sensing contact functions when stimulated by a real (electrical) input signal to a discrete channel:


Real-world I/O wiring



PLC program



Timing diagram



When the pushbutton switch is pressed and the discrete input energized, the first test lamp will blink “on” for exactly one scan of the PLC’s program, then return to its off state. The positive-transition contact (with the letter “P” inside) activates the coil `OUT_test1` only during the scan it sees the status of `IN_test` transition from “false” to “true,” even though the input remains energized for many scans after that transition. Conversely, when the pushbutton switch is released and the discrete input de-energizes, the second test lamp will blink “on” for exactly one scan of the PLC’s

program then return to its off state. The negative-transition contact (with the letter “N” inside) activates the coil `OUT_test2` only during the scan it sees the status of `IN_test` transition from “true” to “false,” even though the input remains de-energized for many scans after that transition:


It should be noted that the duration of a single PLC program scan is typically very short: measured in milliseconds. If this program were actually tested in a real PLC, you would probably not be able to see either test lamp light up, since each pulse is so short-lived. Transitional contacts are typically used any time it is desired to execute an instruction just one time following a “triggering” event, as opposed to executing that instruction over and over again so long as the event status is maintained “true.”

Contacts and coils represent only the most basic of instructions in the Ladder Diagram PLC programming language. Many other instructions exist, which will be discussed in the following subsections.

12.4.2 Counters


A *counter* is a PLC instruction that either increments (counts up) or decrements (counts down) an integer number value when prompted by the transition of a bit from 0 to 1 (“false” to “true”). Counter instructions come in three basic types: *up* counters, *down* counters, and *up/down* counters. Both “up” and “down” counter instructions have single inputs for triggering counts, whereas “up/down” counters have two trigger inputs: one to make the counter increment and one to make the counter decrement.

To illustrate the use of a counter instruction, we will analyze a PLC-based system designed to count objects as they pass down a conveyor belt:



In this system, a continuous (unbroken) light beam causes the light sensor to close its output contact, energizing discrete channel IN4. When an object on the conveyor belt interrupts the light beam from source to sensor, the sensor’s contact opens, interrupting power to input IN4. A pushbutton switch connected to activate discrete input IN5 when pressed will serve as a manual “reset” of the count value. An indicator lamp connected to one of the discrete output channels will serve as an indicator of when the object count value has exceeded some pre-set limit.


We will now analyze a simple Ladder Diagram program designed to increment a counter instruction each time the light beam breaks:




This particular counter instruction (CTU) is an incrementing counter, which means it counts “up” with each off-to-on transition input to its “CU” input. The normally-closed virtual contact (IN_sensor_object) is typically held in the “open” state when the light beam is continuous, by virtue of the fact the sensor holds that discrete input channel energized while the beam is continuous. When the beam is broken by a passing object on the conveyor belt, the input channel de-energizes, causing the virtual contact IN_sensor_object to “close” and send virtual power to the “CU” input of the counter instruction. This increments the counter just as the leading edge of the object breaks the beam. The second input of the counter instruction box (“R”) is the *reset* input, receiving virtual power from the contact IN_switch_reset whenever the reset pushbutton is pressed. If this input is activated, the counter immediately resets its current value (CV) to zero.

Status indication is shown in this Ladder Diagram program, with the counter’s preset value (PV) of 25 and the counter’s current value (CV) of 0 shown highlighted in blue. The preset value is something programmed into the counter instruction before the system put into service, and it serves as a threshold for activating the counter’s output (Q), which in this case turns on the count indicator lamp (the OUT_counts_reached coil). According to the IEC 61131-3 programming standard, this counter output should activate whenever the current value is equal to or greater than the preset value (Q is active if $CV \geq PV$).

This is the status of the same program after thirty objects have passed by the sensor on the conveyor belt. As you can see, the current value of the counter has increased to 30, exceeding the preset value and activating the discrete output:




If all we did not care about maintaining an accurate total count of objects past 25 – but merely wished the program to indicate when 25 objects had passed by – we could also use a *down* counter instruction preset to a value of 25, which turns on an output coil when the count reaches zero:



Here, a “load” input causes the counter’s current value to equal the preset value (25) when activated. With each sensor pulse received, the counter instruction decrements. When it reaches zero, the Q output activates.


A potential problem in either version of this object-counting system is that the PLC cannot discriminate between forward and reverse motion on the conveyor belt. If, for instance, the conveyor belt were ever reversed in direction, the sensor would continue to count objects that had already passed by before (in the forward direction) as those objects retreated on the belt. This would be a problem because the system would “think” more objects had passed along the belt (indicating greater production) than actually did.

One solution to this problem is to use an up/down counter, capable of both incrementing (counting up) and decrementing (counting down), and equip this counter with two light-beam sensors capable of determining direction of travel. If two light beams are oriented parallel to each other, closer than the width of the narrowest object passing along the conveyor belt, we will have enough information to determine direction of object travel:



This is called *quadrature* signal timing, because the two pulse waveforms are approximately 90° (one-quarter of a period) apart in phase. We can use these two phase-shifted signals to increment or decrement an up/down counter instruction, depending on which pulse leads and which pulse lags.

A Ladder Diagram PLC program designed to interpret the quadrature pulse signals is shown here, making use of negative-transition contacts as well as standard contacts:



The counter will increment (count up) when sensor B de-energizes only if sensor A is already in the de-energized state (i.e. light beam A breaks before B). The counter will decrement (count down) when sensor A de-energizes only if sensor B is already in the de-energized state (i.e. light beam B breaks before A).

Note that the up/down counter has both a “reset” (R) input and a “load” input (“LD”) to force the current value. Activating the reset input forces the counter’s current value (CV) to zero, just as we saw with the “up” counter instruction. Activating the load input forces the counter’s current value to the preset value (PV), just as we saw with the “down” counter instruction. In the case of an up/down counter, there are two Q outputs: a QU (output up) to indicate when the current value is equal to or greater than the preset value, and a QD (output down) to indicate when the current value is equal to or less than zero.

Note how the current value (CV) of each counter shown is associated with a tag name of its own, in this case `parts_counted`. The integer number of a counter’s current value (CV) is a variable in the PLC’s memory just like boolean values such as `IN_sensor_A` and `IN_switch_reset`, and may be associated with a tag name or symbolic address just the same²⁴. This allows other instructions in a PLC program to read (and sometimes write!) values from and to that memory location.


²⁴This represents the IEC 61131-3 standard, where each variable within an instruction may be “connected” to its own arbitrary tag name. Other programming conventions may differ somewhat. The Allen-Bradley Logix5000 series of controllers is one of those that differs, following a convention reminiscent of structure element addressing in the C programming language: each counter is given a tag name, and variables in each counter are addressed as elements within that structure. For example, a Logix5000 counter instruction might be named `parts_count`, with the accumulated count value (equivalent to the IEC’s “current value”) addressed as `parts_count.ACC` (each element within the counter specified as a suffix to the counter’s tag name).

12.4.3 Timers

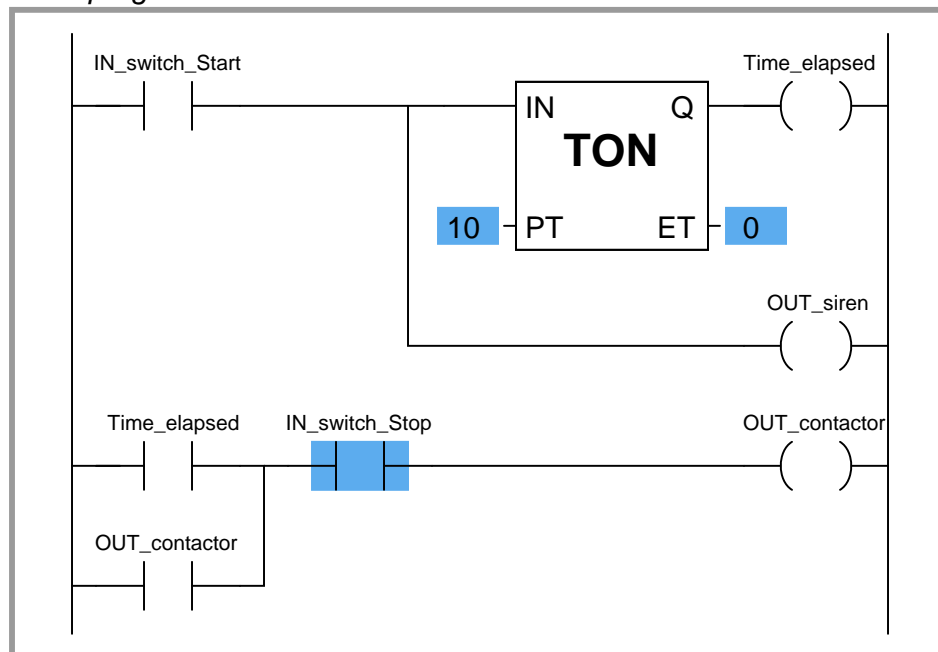
A *timer* is a PLC instruction measuring the amount of time elapsed following an event. Timer instructions come in two basic types: *on-delay* timers and *off-delay* timers. Both “on-delay” and “off-delay” timer instructions have single inputs triggering the timed function.

An “on-delay” timer activates an output only when the input has been active for a minimum amount of time. Take for instance this PLC program, designed to sound an audio alarm siren prior to starting a conveyor belt. To start the conveyor belt motor, the operator must press and hold the “Start” pushbutton for 10 seconds, during which time the siren sounds, warning people to clear away from the conveyor belt that is about to start. Only after this 10-second start delay does the motor actually start (and latch “on”):

Real-world I/O wiring



PLC program




Similar to an “up” counter, the on-delay timer’s elapsed time (ET) value increments once per second until the preset time (PT) is reached, at which time its output (Q) activates. In this program, the preset time value is 10 seconds, which means the Q output will not activate until the “Start” switch has been depressed for 10 seconds. The alarm siren output, which is not activated by the timer, energizes immediately when the “Start” pushbutton is pressed.

An important detail regarding this particular timer’s operation is that it be *non-retentive*. This means the timer instruction should *not* retain its elapsed time value when the input is de-activated. Instead, the elapsed time value should reset back to zero every time the input de-activates. This ensures the timer resets itself when the operator releases the “Start” pushbutton. A *retentive* on-delay timer, by contrast, maintains its elapsed time value even when the input is de-activated. This makes it useful for keeping “running total” times for some event.


Most PLCs provide retentive and non-retentive versions of on-delay timer instructions, such that the programmer may choose the proper form of on-delay timer for any particular application. The IEC 61131-3 programming standard, however, addresses the issue of retentive versus non-retentive timers a bit differently. According to the IEC 61131-3 standard, a timer instruction may be specified with an additional *enable* input (EN) that causes the timer instruction to behave non-retentively when activated, and retentively when de-activated. The general concept of the enable (EN) input is that the instruction behaves “normally” so long as the enable input is active (in this case, non-retentive timing action is considered “normal” according to the IEC 61131-3 standard), but the instruction “freezes” all execution whenever the enable input de-activates. This “freezing” of operation has the effect of retaining the current time (CT) value even if the input signal de-activates.

For example, if we wished to add a retentive timer to our conveyor control system to record total run time for the conveyor motor, we could do so using an “enabled” IEC 61131-3 timer instruction like this:



When the motor’s contactor bit (`OUT_contactor`) is active, the timer is enabled and allowed to time. However, when that bit de-activates (becomes “false”), the timer instruction as a whole is disabled, causing it to “freeze” and retain its current time (CT) value²⁵. This allows the motor to be started and stopped, with the timer maintaining a tally of total motor run time.

If we wished to give the operator the ability to manually reset the total run time value to zero, we could hard-wire an additional switch to the PLC’s discrete input card and add “reset” contacts to the program like this:




Whenever the “Reset” switch is pressed, the timer is enabled (EN) but the timing input (IN) is disabled, forcing the timer to (non-retentively) reset its current time (CT) value to zero.


²⁵The “enable out” (ENO) signal on the timer instruction serves to indicate the instruction’s status: it activates when the enable input (EN) activates and de-activates when either the enable input de-activates or the instruction generates an error condition (as determined by the PLC manufacturer’s internal programming). The ENO output signal serves no useful purpose in this particular program, but it is available if there were any need for other rungs of the program to be “aware” of the run-time timer’s status.

The other major type of PLC timer instruction is the *off-delay* timer. This timer instruction differs from the on-delay type in that the timing function begins as soon as the instruction is de-activated, not when it is activated. An application for an off-delay timer is a cooling fan motor control for a large industrial engine. In this system, the PLC starts an electric cooling fan as soon as the engine is detected as rotating, and keeps that fan running for two minutes following the engine's shut-down to dissipate residual heat:

Real-world I/O wiring




PLC program



When the input (IN) to this timer instruction is activated, the output (Q) immediately activates (with no time delay at all) to turn on the cooling fan motor contactor. This provides the engine with cooling as soon as it begins to rotate (as detected by the speed switch connected to the PLC's discrete input). When the engine stops rotating, the speed switch returns to its normally-open position, de-activating the timer's input signal which starts the timing sequence. The Q output remains active while the timer counts from 0 seconds to 120 seconds. As soon as it reaches 120 seconds, the output de-activates (shutting off the cooling fan motor) and the elapsed time value remains at 120 seconds until the input re-activates, at which time it resets back to zero.

The following timing diagrams compare and contrast on-delay with off-delay timers:




While it is common to find on-delay PLC instructions offered in both retentive and non-retentive forms within the instruction sets of nearly every PLC manufacturer and model, it is almost unheard of to find retentive off-delay timer instructions. Typically, off-delay timers are non-retentive only²⁶.

²⁶The enable (EN) input signals specified in the IEC 61131-3 programming standard make retentive off-delay timers possible (by de-activating the enable input while maintaining the “IN” input in an inactive state), but bear in mind that most PLC implementations of timers do not have separate EN and IN inputs. This means (for most PLC timer instructions) the only input available to activate the timer is the “IN” input, in which case it is *impossible* to create a retentive off-delay timer (since such a timer’s elapsed time value would be immediately re-set to zero each time the input re-activates).

12.4.4 Data comparison instructions

As we have seen with counter and timers, some PLC instructions generate digital values other than simple Boolean (on/off) signals. Counters have current value (CV) registers and timers have elapsed time (ET) registers, both of which are typically integer number values. Many other PLC instructions are designed to receive and manipulate non-Boolean values such as these to perform useful control functions.


The IEC 61131-3 standard specifies a variety of *data comparison* instructions for comparing two non-Boolean values, and generating Boolean outputs. The basic comparative operations of “less than” ($<$), “greater than” ($>$), “less than or equal to” (\leq), “greater than or equal to” (\geq), “equal to” ($=$), and “not equal to” (\neq) may be found as a series of “box” instructions in the IEC standard:




The Q output for each instruction “box” activates whenever the evaluated comparison function is “true” and the enable input (EN) is active. If the enable input remains active but the comparison function is false, the Q output de-activates. If the enable input de-de-activates, the Q output retains its last state.

A practical application for a comparative function is something called *alternating motor control*, where the run-times of two redundant electric motors²⁷ are monitored, with the PLC determining which motor to turn on next based on which motor has run the least:

Real-world I/O wiring



PLC program




In this program, two retentive on-delay timers keep track of each electric motor's total run time, storing the run time values in two registers in the PLC's memory: `Motor_A_runtime` and

²⁷Perhaps two pumps performing the same pumping function, one serving as a backup to the other. Alternating motor control ensures the two motors' run times are matched as closely as possible.

Motor_B_runtime. These two integer values are input to the “greater than” instruction box for comparison. If motor A has run longer than motor B, motor B will be the one enabled to start up next time the “start” switch is pressed. If motor A has run less time or the same amount of time as motor B (the scenario shown by the blue-highlighted status indications), motor A will be the one enabled to start. The two series-connected virtual contacts `OUT_motor_A` and `OUT_motor_B` ensure the comparison between motor run times is not made until both motors are stopped. If the comparison were continually made, a situation might arise where *both* motors would start if someone happened to press the Start pushbutton with one motor is already running.


12.4.5 Math instructions

The IEC 61131-3 standard specifies several dedicated ladder instructions for performing arithmetic calculations. Some of them are shown here:



As with the data comparison instructions, each of these math instructions must be enabled by an “energized” signal to the enable (EN) input. Input and output values are linked to each math instruction by tag name.

An example showing the use of such instructions is shown here, converting a temperature measurement in units of degrees Fahrenheit to units of degrees Celsius. In this particular case, the program inputs a temperature measurement of 138 °F and calculates the equivalent temperature of 58.89 °C:




Note how two separate math instructions were required to perform this simple calculation, as well as a dedicated variable (X) used to store the intermediate calculation between the subtraction and the division “boxes.”

Although not specified in the IEC 61131-3 standard, many programmable logic controllers support Ladder Diagram math instructions allowing the direct entry of arbitrary equations. Rockwell

(Allen-Bradley) Logix5000 programming, for example, has the “Compute” (CPT) function, which allows any typed expression to be computed in a single instruction as opposed to using several dedicated math instructions such as “Add,” “Subtract,” etc. General-purpose math instructions dramatically shorten the length of a ladder program compared to the use of dedicated math instructions for any applications requiring non-trivial calculations.

For example, the same Fahrenheit-to-Celsius temperature conversion program implemented in Logix5000 programming only requires a single math instruction and no declarations of intermediate variables:


Rockwell Logix5000 PLC program



12.4.6 Sequencers


Many industrial processes require control actions to take place in certain, predefined sequences. Batch processes are perhaps the most striking example of this, where materials for making a batch must be loaded into the process vessels, parameters such as temperature and pressure controlled during the batch processing, and then discharge of the product monitored and controlled. Before the advent of reliable programmable logic devices, this form of sequenced control was usually managed by an electromechanical device known as a *drum sequencer*. This device worked on the principle of a rotating cylinder (drum) equipped with tabs to actuate switches as the drum rotated into certain positions. If the drum rotated at a constant speed (turned by a clock motor), those switches would actuate according to a timed schedule²⁸.

The following photograph shows a drum sequencer with 30 switches. Numbered tabs on the circumference of the drum mark the drum's rotary position in one of 24 increments. With this number of switches and tabs, the drum can control up to thirty discrete (on/off) devices over a series of twenty-four sequenced steps:



²⁸The operation of the drum is not unlike that of an old *player piano*, where a strip of paper punched with holes caused hammers in the piano to automatically strike their respective strings as the strip was moved along at a set speed, thus playing a pre-programmed song.

A typical application for a sequencer is to control a *Clean In Place (CIP)* system for a food processing vessel, where a process vessel must undergo a cleaning cycle to purge it of any biological matter between food processing cycles. The steps required to clean the vessel are well-defined and must always occur in the same sequence in order to ensure hygienic conditions. An example timing chart is shown here:



In this example, there are nine discrete outputs – one for each of the nine final control elements (pumps and valves) – and seventeen steps to the sequence, each one of them timed. In this particular sequence, the only input is the discrete signal to commence the CIP cycle. From the initiation of the CIP to its conclusion two and a half hours (150 minutes) later, the sequencer simply steps through the programmed routine.


Another practical application for a sequencer is to implement a *Burner Management System* (BMS), also called a *flame safety system*. Here, the sequencer manages the safe start-up of a combustion burner: beginning by “purging” the combustion chamber with fresh air to sweep out any residual fuel vapors, waiting for the command to light the fire, energizing a spark ignition system on command, and then continuously monitoring for presence of good flame and proper fuel supply pressure once the burner is lit.

In a general sense, the operation of a drum sequencer is that of a *state machine*: the output of the system depends on the condition of the machine’s internal state (the drum position), not just the conditions of the input signals. Digital computers are very adept at implementing state functions, and so the general function of a drum sequencer should be (and is) easy to implement in a PLC. Other PLC functions we have seen (“latches” and timers in particular) are similar, in that the PLC’s output at any given time is a function of both its present input condition(s) and its past input condition(s). Sequencing functions expand upon this concept to define a much larger number of possible states (“positions” of a “drum”), some of which may even be timed.

Unfortunately, despite the utility of drum sequence functions and their ease of implementation in digital form, there seems to be very little standardization between PLC manufacturers regarding sequencing instructions. Sadly, the IEC 61131-3 standard (at least at the time of this writing, in 2009) does not specifically define a sequencing function suitable for Ladder Diagram programming. PLC manufacturers are left to invent sequencing instructions of their own design. What follows here is an exploration of some different sequencer instructions offered by PLC manufacturers.

Koyo “drum” instructions

The *drum* instruction offered in Koyo PLCs is a model of simplicity itself. This instruction is practically self-explanatory, as shown in the following example:




The three-by-three grid of squares represent steps in the sequence and bit states for each step. Rows represent steps, while columns represent output bits written by the drum instruction. In this particular example, a three-step sequence proceeds at the command of a single input (X001), and the drum instruction’s advance from one step to the next proceeds strictly on the basis of elapsed time (a *time base* orientation). When the input is active, the drum proceeds through its timed sequence. When the input is inactive, the drum halts wherever it left off, and resumes timing as soon as the input becomes active again.

Being based on time, each step in the drum instruction has a set time duration for completion. The first step in this particular example has a duration of 10 seconds, the second step 15 seconds, and the third step 18 seconds. At the first step, only output bit Y001 is set. In the second step, only output bit Y002 is set. In the third step, output bits Y002 and Y003 are set (1), while bit Y001 is reset (0). The colored versus uncolored boxes reveal which output bits are set and reset with each step. The current step number is held in memory register DS1, while the elapsed time (in seconds) is stored in timer register TD1. A “complete” bit is set at the conclusion of the three-step sequence.

Koyo drum instructions may be expanded to include more than three steps and more than three output bits, with each of those step times independently adjustable and each of the output bits arbitrarily assigned to any writable bit addresses in the PLC’s memory.

This next example of a Koyo drum instruction shows how it may be set up to trigger on *events* rather than on elapsed times. This orientation is called an *event base*:

Koyo CLICK PLC program




Here, a three-step sequence proceeds when enabled by a single input (X001), with the drum instruction's advance from one step to the next proceeding only as the different event condition bits become set. When the input is active, the drum proceeds through its sequence when each event condition is met. When the input is inactive, the drum halts wherever it left off regardless of the event bit states.

For example, during the first step (when only output bit Y001 is set), the drum instruction waits for the first condition input bit X002 to become set (1) before proceeding to step 2, with time being irrelevant. When this happens, the drum immediately advances to step 2 and waits for input bit X003 to be set, and so forth. If all three event conditions were met simultaneously (X002, X003, and X004 all set to 1), the drum would skip through all steps as fast as it could (one step per PLC program scan) with no appreciable time elapsed for each step. Conversely, the drum instruction will wait as long as it must for the right condition to be met before advancing, whether that event takes place in milliseconds or in days.

Allen-Bradley sequencer instructions

Rockwell (Allen-Bradley) PLCs use a more sophisticated set of instructions to implement sequences. The closest equivalent to Koyo's *drum* instruction is the Allen-Bradley *SQO* (Sequencer Output) instruction, shown here:



You will notice there are no colored squares inside the SQO instruction box to specify when certain bits are set or reset throughout the sequence, in contrast to the simplicity of the Koyo PLC's drum instruction. Instead, the Allen-Bradley SQO instruction is told to read a set of 16-bit words beginning at a location in the PLC's memory arbitrarily specified by the programmer, one word at a time. It steps to the next word in that set of words with each new position (step) value. This means Allen-Bradley sequencer instructions rely on the programmer already having pre-loaded an area of the PLC's memory with the necessary 1's and 0's defining the sequence. This makes the Allen-Bradley sequencer instruction more challenging for a human programmer to interpret because the bit states are not explicitly shown inside the SQO instruction box, but it also makes the sequencer far more flexible in that these bits are not fixed parameters of the SQO instruction and therefore may be dynamically altered as the PLC runs. With the Koyo drum instruction, the assigned output states are part of the instruction itself, and are therefore fixed once the program is downloaded to the PLC (i.e. they cannot be altered without editing and re-loading the PLC's program). With the Allen-Bradley, the on-or-off bit states for the sequence may be freely altered²⁹ during run-time. This is a very useful feature in recipe-control applications, where the recipe is subject to change at the whim of production personnel, and they would rather not have to rely on a technician or an engineer to re-program the PLC for each new recipe.

The "Length" parameter tells the SQO instruction how many words will be read (i.e. how many steps are in the entire sequence). The sequencer advances to each new position when its

²⁹Perhaps the most practical way to give production personnel access to these bits without having them learn and use PLC programming software is to program an HMI panel to write to those memory areas of the PLC. This way, the operators may edit the sequence at any time simply by pressing "buttons" on the screen of the HMI panel, and the PLC need not have its program altered in any "hard" way by a technician or engineer.

enabling input transitions from inactive to active (from “false” to “true”), just like a count-up (CTU) instruction increments its accumulator value with each new false-to-true transition of the input. Here we see another important difference between the Allen-Bradley SQR instruction and the Koyo drum instruction: the Allen-Bradley instruction is fundamentally *event-driven*, and does not proceed on its own like the Koyo drum instruction is able to when configured for a *time* base.

Sequencer instructions in Allen-Bradley PLCs use a notation called *indexed addressing* to specify the locations in memory for the set of 16-bit words it will read. In the example shown above, we see the “File” parameter specified as `#B3:0`. The “#” symbol tells the instruction that this is a *starting* location in memory for the first 16-bit word, when the instruction’s position value is zero. As the position value increments, the SQR instruction reads 16-bit words from successive addresses in the PLC’s memory. If `B3:0` is the word referenced at position 0, then `B3:1` will be the memory address read at position 1, `B3:2` will be the memory address read at position 2, etc. Thus, the “position” value causes the SQR instruction to “point” or “index” to successive memory locations.

The bits read from each indexed word in the sequence are compared against a static mask³⁰ specifying which bits in the indexed word are relevant. At each position, only these bits are written to the destination address.

As with most other Allen-Bradley instructions, the sequencer requires the human programmer to declare a special area in memory reserved for the instruction’s internal use. The “R6” file exists just for this purpose, each element in that file holding bit and integer values associated with a sequencer instruction (e.g. the “enable” and “done” bits, the array length, the current position, etc.).

³⁰In this particular example, the mask value is FFFF hexadecimal, which means all 1’s in a 16-bit field. This mask value tells the sequencer instruction to regard *all* bits of each `B3` word that is read. To contrast, if the mask were set to a value of 000F hexadecimal instead, the sequencer would only pay attention to the four least-significant bits of each `B3` word that is read, while ignoring the 12 more-significant bits of each 16-bit word. The mask allows the SQR instruction to only write to selected bits of the destination word, rather than always writing all 16 bits of the indexed word to the destination word.

To illustrate, let us examine a set of bits held in the B3 file of an Allen-Bradley SLC 500 PLC, showing how each row (element) of this data file would be read by an SQO instruction as it stepped through its positions:

Data File B3 (bin) -- BINARY

B3:0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	If File = #B3:0, then . . .	
B3:1	0	1	1	1	1	0	0	0	0	0	0	0	0	0	1	← Read at position = 1	
B3:2	1	0	0	0	0	0	1	1	0	0	0	0	0	1	0	← Read at position = 2	
B3:3	0	0	0	1	0	0	0	0	0	1	0	1	0	1	0	← Read at position = 3	
	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0	

The sequencer’s position number is added to the file reference address as an *offset*. Thus, if the data file is specified in the SQO instruction box as #B3:0, then B3:1 will be the row of bits read when the sequencer’s position value is 1, B3:2 will be the row of bits read when the position value is 2, and so on.

The *mask* value specified in the SQO instruction tells the instruction which bits out of each row will be copied to the destination address. A mask value of FFFFh (FFFF in *hexadecimal* format) means all 16 bits of each B3 word will be read and written to the destination. A mask value of 0001h means only the first (least-significant) bit will be read and written, with the rest being ignored.

Let’s see what would happen with an SQO instruction having a mask value of 000Fh, starting from file index #B3:0, and writing to a destination that is output register 0:0.0, given the bit array values in file B3 shown above:


B3:0	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0		
B3:1	0	1	1	1	1	0	0	0	0	0	0	0	0	0	1		
B3:2	1	0	0	0	0	0	1	1	0	0	0	0	0	1	0	← Read at position = 2	
B3:3	0	0	0	1	0	0	0	0	0	1	0	1	0	1	0		
	Bit 15	Bit 14	Bit 13	Bit 12	Bit 11	Bit 10	Bit 9	Bit 8	Bit 7	Bit 6	Bit 5	Bit 4	Bit 3	Bit 2	Bit 1	Bit 0	

Mask	0	0	0	0	0	0	0	0	0	0	0	0	0	1	1	1	1	000Fh
------	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	-------

O:0.0	X	X	X	X	X	X	X	X	X	X	X	X	X	0	0	1	0	Output register as written at sequencer position = 2
-------	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---	---------------------------------------------------------

When this SQO instruction is at position 2, it reads the bit values 0010 from B3:2 and writes only those four bits to 0:0.0. The “X” symbols shown in the illustration mean that all the other bits in that output register are untouched – the SQO instruction does not write to those bits because they are “masked off” from being written. You may think of the mask’s zero bits inhibiting source bits from being written to the destination word in the same sense that *masking tape* prevents paint from being applied to a surface.

The following Allen-Bradley SLC 500 PLC program shows how a pair of SQO instructions plus an on-delay timer instruction may be used to duplicate the exact same functionality as the “time base” Koyo drum instruction presented earlier:




The first SQO instruction reads bits in the B3 file array, sending only the three least-significant of them to the output register 0:0.0 (as specified by the 0007h mask value). The second SQO instruction reads integer number values from elements of the N7 integer file and places them into the “preset” register of timer T4:0, so as to dynamically update the timer’s preset value with each step of the sequence. The timer, in turn, counts off each of the time delays and then enables both sequencers to advance to the next position when the specified time has elapsed. Here we see a tremendous benefit of the SQO instruction’s indexed memory addressing: the fact that the SQO instruction reads its bits from arbitrarily-specified memory addresses means we may use SQO instructions to sequence *any type of data existing in the PLC’s memory!* We are not limited to turning on and off individual bits as we are with the Koyo drum instruction, but rather are free to index whole integer numbers, ASCII characters, or any other forms of binary data resident in the PLC’s memory.

Data file windows appear on the computer screen showing the bit array held in the B3 file as well as the timer values held in the N7 file. In this live screenshot, we see both sequencer instructions at position 2, with the second SQO instruction having loaded a value of 15 seconds from register N7:2 to the timer’s preset register T4:0.PRE.

Note how the enabling contact address for the second SQO instruction is the “enable” bit of the first instruction, ensuring both instructions are enabled simultaneously. This keeps the two separate sequencers synchronized (on the same step).

Event-based transitions may be implemented in Allen-Bradley PLCs using a complementary sequencing instruction called SQC (Sequencer Compare). The SQC instruction is set up very similar to the SQO instruction, with an indexed file reference address to read from, a reserved memory structure for internal use, a set length, and a position value. The purpose of the SQC instruction is to read a data register and compare it against another data register, setting a “found” (FD) bit if the two match. Thus, the SQC instruction is ideally suited for detecting when certain conditions have been met, and thus may be used to enable an SQO instruction to proceed to the next step in its sequence.

The following program example shows an Allen-Bradley MicroLogix 1100 PLC programmed with both an SQO and an SQC instruction:



The three-position SQO (Sequencer Output) instruction reads data from B3:1, B3:2, and B3:3, writing the four least-significant of those bits to output register O:0.0. The three-position SQC (Sequencer Compare) instruction reads data from B3:6, B3:7, and B3:8, comparing the four least-significant of those bits against input bits in register I:0.0. When the four input bit conditions match the selected bits in the B3 file, the SQC instruction’s FD bit is set, causing both the SQO instruction and the SQC instruction to advance to the next step.

Lastly, Allen-Bradley PLCs offer a third sequencing instruction called *Sequencer Load* (SQL), which performs the opposite function as the Sequencer Output (SQO). An SQL instruction takes data from a designated source and writes it into an indexed register according to a position count value, rather than reading data from an indexed register and sending it to a designated destination as does the SQO instruction. SQL instructions are useful for reading data from a live process and storing it in different registers within the PLC’s memory at different times, such as when a PLC is used for *datalogging* (recording process data).

12.5 Structured Text (ST) programming

(Will be addressed in future versions of this book)

12.6 Instruction List (IL) programming

(Will be addressed in future versions of this book)

12.7 Function Block Diagram (FBD) programming

(Will be addressed in future versions of this book)

12.8 Sequential Function Chart (SFC) programming


(Will be addressed in future versions of this book)

12.9 Human-Machine Interfaces

Programmable logic controllers are built to input various signal types (discrete, analog), execute control algorithms on those signals, and then output signals in response to control processes. By itself, a PLC generally lacks the capability of displaying those signal values and algorithm variables to human operators. A technician or engineer with access to a personal computer and the requisite software for editing the PLC's program may connect to the PLC and view the program's status "online" to monitor signal values and variable states, but this is not a practical way for operations personnel to monitor what the PLC is doing on a regular basis. In order for operators to monitor and adjust parameters inside the PLC's memory, we need a different sort of interface allowing certain variables to be read and written without compromising the integrity of the PLC by exposing too much information or allowing any unqualified person to alter the program itself.

One solution to this problem is a dedicated computer display programmed to provide selective access to certain variable's in the PLC's memory, generally referred to as *Human³¹-Machine Interface*, or *HMI*.

HMIs may take the form of general-purpose ("personal") computers running special graphic software to interface with a PLC, or as special-purpose computers designed to be mounted in sheet metal panel fronts to perform no task but the operator-PLC interface. This first photograph shows an example of an ordinary personal computer (PC) with HMI software running on it:




The display shown here happens to be for monitoring a vacuum swing adsorption (VSA) process for purifying oxygen extracted from ambient air. Somewhere, a PLC (or collection of PLCs) is monitoring and controlling this VSA process, with the HMI software acting as a "window" into the PLC's memory to display pertinent variables in an easy-to-interpret form for operations personnel. The personal computer running this HMI software connects to the PLC(s) via digital network cables such as Ethernet.

³¹An older term for an operator interface panel was the "Man-Machine Interface" or "MMI." However, this fell out of favor due to its sexist tone.

This next photograph shows an example of a special-purpose HMI panel designed and built expressly to be used in industrial operating environments:




These HMI panels are really nothing more than “hardened” personal computers built ruggedly and in a compact format to facilitate their use in industrial environments. Most industrial HMI panels come equipped with touch-sensitive screens, allowing operators to press their fingertips on displayed objects to change screens, view details on portions of the process, etc.



Technicians and/or engineers program HMI displays to read and write data via a digital network to one or more PLCs. Graphical objects arrayed on the display screen of an HMI often mimic real-world indicators and switches, in order to provide a familiar interface for operations personnel. A “pushbutton” object on the face of an HMI panel, for example, would be configured to *write* one bit of data to the PLC, in a manner similar to a real-world switch writing one bit of data to the PLC’s input register.

Modern HMI panels and software are almost exclusively tag-based, with each graphic object on the screen associated with at least one data tag name, which in turn is associated to data points (bits, or words) in the PLC by way of a tag name database file resident in the HMI. Graphic objects on the HMI screen either accept (read) data from the PLC to present useful information to the operator, send (write) data to the PLC from operator input, or both. The task of programming an HMI unit consists of building a tag name database and then drawing screens to illustrate the process to as good a level of detail as operators will need to run it.

An example screenshot of a tag name database table for a modern HMI is shown here:



The screenshot shows a window titled "Tag Name Database" with a close button (X) in the top right corner. Below the title bar, there is a "Device Name:" dropdown menu set to "PLC_01" and a "Fill Down" button. Below these are two groups of checkboxes: "Display System Tags" (checked) and "Display Internal Tags" (checked), each with a "Highlight" checkbox (unchecked). To the right of these is a "Highlight Unused Tags" checkbox (unchecked). Below the controls is a table with the following data:

No.	Tag Name	Data Type	PLC Address	Device Name	Attribute
1	START_PUSHBUTTON	Discrete	10024	PLC_01	R
2	STOP_PUSHBUTTON	Discrete	10031	PLC_01	R
3	MOTOR_RUN	Discrete	00005	PLC_01	R/w
4	MOTOR_RUN_TIMER	Unsigned int 32	40010	PLC_01	R/w
5	START_COUNTER	Unsigned int 16	40242	PLC_01	R/w
6	MOTOR_TEMPERATURE	Floating PT 32	30008	PLC_01	R
7	ERROR_MESSAGE	Ascii String	40560	PLC_01	R/w
8	MOTOR_SPEED	Floating PT 32	30017	PLC_01	R

The tag name database is accessed and edited using the same software to create graphic images in the HMI. In this particular example you can see several tag names (e.g. `START_PUSHBUTTON`, `MOTOR_RUN_TIMER`, `ERROR_MESSAGE`, `MOTOR_SPEED`) associated with data points within the PLC's memory (in this example, the PLC addresses are shown in Modbus register format). In many cases the tag name editor will be able to display corresponding PLC memory points in the same manner as they appear in the PLC programming editor software (e.g. `I:5/10`, `SM0.4`, `C11`, etc.).

An important detail to note in this tag name database display is the read/write attributes of each tag. Note in particular how four of the tags shown are *read-only*: this means the HMI only has permission to read the values of those four tags from the PLC's memory, and not to write (alter) those values. The reason for this in the case of these four tags is that those tags refer to PLC input data points. The `START_PUSHBUTTON` tag, for instance, refers to a discrete input in the PLC energized by a real pushbutton switch. As such, this data point gets its state from the energization of the discrete input terminal. If the HMI were to be given *write* permission for this data point, there would likely be a conflict. Suppose input terminal on the PLC were energized (setting the `START_PUSHBUTTON` bit to a "1" state) and the HMI simultaneously attempted to write a "0" state to the same tag. One of these two data sources would win, and other would lose, possibly resulting in unexpected behavior from the PLC program. For this reason, data points in the PLC linked to real-world inputs should always be limited as "read-only" permission in the HMI's database, so the

HMI cannot possibly generate a conflict.

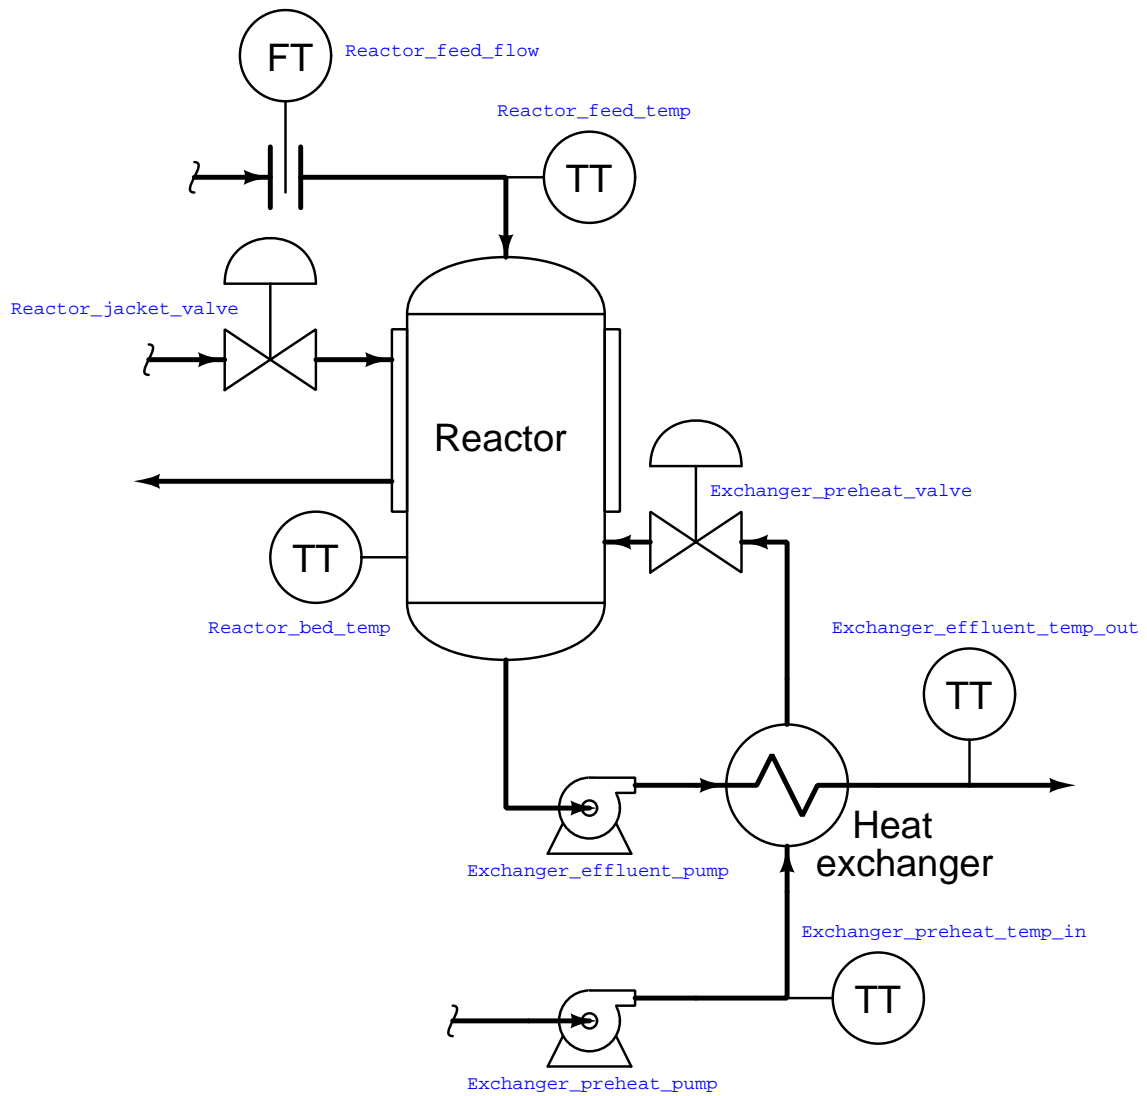
The potential for data conflict also exists for some of the other points in the database, however. A good example of this is the `MOTOR_RUN` bit, which is the bit within the PLC program telling the real-world motor to run. Presumably, this bit gets its data from a coil in the PLC's Ladder Diagram program. However, since it also appears in the HMI database with *read/write* permission, the potential exists for the HMI to over-write (i.e. conflict) that same bit in the PLC's memory. Suppose someone programmed a toggling "pushbutton" screen object in the HMI linked to this tag: pushing this virtual "button" on the HMI screen would attempt to set the bit (1), and pushing it again would attempt to reset the bit (0). If this same bit is being written to by a coil in the PLC's program, however, there exists the distinct possibility that the HMI's "pushbutton" object and the PLC's coil will conflict, one trying to tell the bit to be a "0" while the other tries to tell that bit to be a "1". This situation is quite similar to the problem experienced when multiple coils in a Ladder Diagram program are addressed to the same bit.

The general rule to follow here is *never allow more than one element to write to any data point*. In my experience teaching PLC and HMI programming, this is one of the more common errors students make when first learning to program HMIs: they will try to have both the HMI and the PLC writing to the same memory locations, with weird results.

One of the lessons you will learn when programming large, complex systems is that it is very beneficial to define all the necessary tag names *before* beginning to lay out graphics in an HMI. The same goes for PLC programming: it makes the whole project go faster with less confusion if you take the time to define all the necessary I/O points (and tag names, if the PLC programming software supports tag names in the programming environment) before you begin to create any code specifying how those inputs and outputs will relate to each other.

Maintaining a consistent convention for tag names is important, too. For example, you may wish to begin the tag name of every hard-wired I/O point as either `INPUT` or `OUTPUT` (e.g. `INPUT_PRESSURE_SWITCH_HIGH`, `OUTPUT_SHAKER_MOTOR_RUN`, etc.). The reason for maintaining a strict naming convention is not obvious at first, since the whole point of tag names is to give the programmer the freedom to assign *arbitrary names* to data points in the system. However, you will find that most tag name editors list the tags in alphabetical order, which means a naming convention organized in this way will present all the input tags contiguously (adjacent) in the list, all the output tags contiguously in the list, and so on.

Another way to leverage the alphabetical listing of tag names to your advantage is to begin each tag name with a word describing its association to a major piece of equipment. Take for instance this example of a process with several data points defined in a PLC control system and displayed in an HMI:



If we list all these tags in alphabetical order, the association is immediately obvious:

- `Exchanger_effluent_pump`
- `Exchanger_effluent_temp_out`
- `Exchanger_preheat_pump`
- `Exchanger_preheat_temp_in`
- `Exchanger_preheat_valve`
- `Reactor_bed_temp`
- `Reactor_feed_flow`
- `Reactor_feed_temp`
- `Reactor_jacket_valve`

As you can see from this tag name list, all the tags directly associated with the heat exchanger are located in one contiguous group, and all the tags directly associated with the reactor are located in the next contiguous group. In this way, judicious naming of tags serves to group them in hierarchical fashion, making them easy for the programmer to locate at any future time in the tag name database.

You will note that all the tag names shown here lack space characters between words (e.g. instead of “Reactor bed temp”, a tag name should use hyphens or underscore marks as spacing characters: “Reactor_bed_temp”), since spaces are generally assumed by computer programming languages to be delimiters (separators between different variable names).

Like programmable logic controllers themselves, the capabilities of HMIs have been steadily increasing while their price decreases. Modern HMIs support graphic trending, data archival, advanced alarming, and even web server ability allowing other computers to easily access certain data over wide-area networks. The ability of HMIs to log data over long periods of time relieves the PLC of having to do this task, which is very memory-intensive. This way, the PLC merely “serves” current data to the HMI, and the HMI is able to keep a record of current and past data using its vastly larger memory reserves³².

Some modern HMI panels even have a PLC built inside the unit, providing control and monitoring in the same device. Such panels provide terminal strip connection points for discrete and even analog I/O, allowing all control and interface functions to be located in a single panel-mount unit.

³²If the HMI is based on a personal computer platform (e.g. Rockwell RSView, Wonderware, FIX/Intellution software), it may even be equipped with a hard disk drive for enormous amounts of historical data storage.

12.10 How to teach yourself PLC programming

First and foremost, you need to get your very own PLC to work with. Computer programming of any kind is not a spectator sport, and can only be learned by significant investment of time and effort at the keyboard. In many ways, learning to program is like learning a new spoken or written language: there is new vocabulary and new grammatical rules to master, and many ways to make mistakes.

Fortunately, many low-cost PLCs exist on the market for individuals to purchase. My own personal favorites are the “CLICK” PLC models manufactured by Koyo and marketed through Automation Direct, and also the Allen-Bradley MicroLogix series of PLC (especially the 1000 and 1100 models).


The first document you should read once you get your PLC is something called a *Getting Started* guide. Every PLC manufacturer publishes a document with this name (or something similar such as *Quick Start* or *Getting Results*). This manual will step you through all the basic procedures for entering a simple program into your PLC and getting it to run. It is generally *far* easier to learn programming by copying and adapting a worked example than it is to start from a “blank page” on your own, just as it is easiest to learn a spoken or written language by practicing sentences spoken in that language by other people before constructing your own sentences from scratch.

In order to work with your PLC, you will need a convenient way to simulate different input conditions coming from discrete (switch) devices. Any set of hand-operated switches will do, my recommendation being household light switches (very inexpensive and rugged). Attaching an array of these switches to a wooden board along with the PLC and interconnecting terminal blocks forms what is often called a PLC *trainer*. The following photograph shows one such trainer³³, using an Allen-Bradley MicroLogix 1000 PLC:



³³This particular trainer was partially constructed from recycled materials – the wooden platform, light switches, and power cord – to minimize cost.

Another example of a student-built PLC trainer is this unit, housed inside of an attaché case. Not only does this trainer contain an Allen-Bradley MicroLogix 1100 PLC along with input switches and output indicator lights, but it also includes an HMI touch-screen panel on a fold-down bracket:



The educational value of building your own PLC trainer is difficult to overstate when learning about PLCs. Learning how to build properly-functioning I/O circuits is every bit as important to a working technician as learning how to develop PLC programs. Additionally, the experience gained in general wiring layout and fabrication are valuable skills for any instrumentation practitioner.

Once you have learned the basic steps for entering, running, and saving a PLC program, you are ready to begin building your knowledge of the language's vocabulary and grammar. In computer programming (of all types), there are different *functions* of the language one must become familiar with in order to do useful tasks. A great way to learn how to use these functions is to create your own "demonstration" programs illustrating the use of each function.

For example, if you open up the pages of almost any computer programming book, somewhere near the beginning you will find a demonstration program called "Hello World!" The purpose of a "Hello World!" program is to do nothing more than display the words *Hello World!* on the computer screen. It is an entirely useless program to run, but it is highly useful for gaining teaching the programmer the basics of program construction and text message functionality.

By the same token, you may learn the basics of each programming function by writing simple "Hello World"-type programs illustrating each one of those functions. These demonstration programs will not serve any useful purpose (other than to help you learn), and should be kept as simple as possible in order to minimize confusion.

For example, *every* PLC provides instructions to perform the following tasks:

- Turn discrete outputs on and off
- Count discrete events
- Time events
- Control events in a specific sequence
- Compare numerical values (greater than, less than, equal, not equal)
- Perform arithmetic functions

Just as every spoken or written language has verbs, nouns, adjectives, and adverbs to describe actions and things, every PLC programming language has specific functions to perform useful tasks. The details of how to perform each function will vary somewhat between PLC manufacturers and models, but the overall functions are quite similar. The reference manuals provided for your PLC will describe in detail how to use each function. Your task is to write simple demonstration programs for each function, allowing you to directly explore how each function works, and to gain an understanding of each function by observing its behavior and also by making (inevitable) mistakes.

After writing each demonstration program, you should add a lot of comments to it, so you will be able to understand what you did later when you go back to your demonstration program for reference. These comments should cover the following points:

- Proper use of the function
- A verbal description of what the function does
- A list of possible (practical) uses for the function
- Idiosyncrasies of the function (i.e. odd or unexpected behavior, tricky points to watch out for)
- Mistakes you may have made (and thus might make again!) in using the function

Years ago when I was teaching myself how to program using the *C* language, I wrote a set of “tutorial” programs demonstrating common programming functions and techniques. The following is a partial list of these tutorial programs, which I still keep to this day:

- Program that accepts and then prints alphanumeric characters (including their equivalent numerical values)
- Program demonstrating how to use command-line arguments to the `main()` function
- Program demonstrating basic “curses” commands for plotting characters at arbitrary locations on the screen
- Program illustrating the declaration and use of *data structures*
- Program illustrating how to prototype and then call *functions* (subroutines)
- Program executing an infinite loop
- Program illustrating how to return a *pointer* from a function

Each one of these tutorial programs is heavily commented, to explain to myself in my own words how they work and what they are doing. Not only did they help me learn how to write programs in C, but they also serve as a handy reference for me any time in the future I need to refresh my knowledge. The act of writing tutorial programs is akin to *journaling* as a way to work through complex problems in life – in a way, it is like having a conversation with yourself.

12.11 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader’s view of this chapter’s concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **“Normal” switch status:** the “normal” status of a switch contact as defined by the manufacturer is its *resting* condition (minimum stimulus).
- **“Seal-in” circuit:** when an electrical relay uses one of its own switch contacts to continue its own coil energization after the initial triggering event has passed. Relevant to all manner of relay control circuits.
- **Sourcing versus sinking:** whether the electronic device in question is outputting (conventional flow) current or inputting current. Relevant to the proper connection of discrete DC input and output cards.

References

“1758 PLC-5 Programmable Controllers Addressing Reference Manual”, Publication 5000-6.4.4, Allen-Bradley Company, Inc., Milwaukee, WI, 1995.

“Allen-Bradley I/O Modules Wiring Diagrams”, Publication CIG-WD001A-EN-P, Rockwell Automation, Inc., Milwaukee, WI, 2005.

IEC 61131-3, “International Standard, Programmable Controllers – Part 3: Programming Languages”, Edition 2.0, International Electrotechnical Commission, Geneva, Switzerland, 2003.

“Logix5000 Controllers I/O and Tag Data”, Publication 1756-PM004B-EN-P, Rockwell Automation, Inc., Milwaukee, WI, 2008.

“Programming with STEP 7”, Siemens AG, Nürnberg, Germany, 2006.

“S7-200 Programmable Controller System Manual”, Order Number 6ES7298-8FA24-8BH0, Edition 09/2007, Siemens AG, Nürnberg, Germany, 2007.

“SLC 500 Family of Programmable Controllers Addressing Reference Manual”, Publication 5000-6.4.23, Allen-Bradley Company, Inc., Milwaukee, WI, 1995.

“SLC 500 Modular Hardware Style User Manual”, Publication 1747-UM011E-EN-P, Rockwell Automation, Inc., Milwaukee, WI, 2004.

Chapter 13

Analog electronic instrumentation


An “analog” electronic signal is a voltage or current proportionate to the value of some physical measurement or control quantity. An instrument is often classified as being “analog” simply by virtue of using an analog signal standard to communicate information, even if the internal construction and design of the instrument may be mostly digital in nature. This is to distinguish such instruments from those making use of no analog electronic signals at all (e.g. wireless or Fieldbus instruments).

13.1 4 to 20 mA analog current signals

The most popular form of signal transmission used in modern industrial instrumentation systems (as of this writing) is the 4 to 20 milliamp DC standard. This is an *analog* signal standard, meaning that the electric current is used to proportionately represent measurements or command signals. Typically, a 4 milliamp current value represents 0% of scale, a 20 milliamp current value represents 100% of scale, and any current value in between 4 and 20 milliamps represents a commensurate percentage in between 0% and 100%. The following table shows the corresponding current and percentage values for each 25% increment between 0% and 100%. Every instrument technician tasked with maintaining 4-20 mA instruments commits these values to memory, because they are referenced so often:

Current value	% of scale
4 mA	0%
8 mA	25%
12 mA	50%
16 mA	75%
20 mA	100%


For example, if we were to calibrate a 4-20 mA temperature transmitter for a measurement range of 50 to 250 degrees C, we could relate the current and measured temperature values on a graph like this:



This is not unlike 3-15 pounds per square inch (PSI) pneumatic signal standard, where a varying air pressure signal proportionately represents some process variable. Both 3-15 PSI and 4-20 mA signal standards are referred to as *live zero* because their ranges begin with a non-zero value. This “live” zero provides a simple means of discriminating between a legitimate 0% signal value and a failed signal (e.g. leaking tube or severed cable)¹.

¹Not all industrial measurement and control signals are “live zero” like the 3-15 PSI and 4-20 mA standards. 0 to

An important concept to grasp with all analog instrumentation is that instruments sending and receiving analog signals must be compatibly ranged in order to properly represent the desired variable. To illustrate, let us consider a temperature measurement system consisting of a thermocouple², a temperature transmitter, a 250 ohm resistor (to convert the 4-20 mA analog signal into a 1-5 volt analog signal), and a special voltmeter functioning as a temperature indicator:



Note how the output range of each sending device matches the input range of its corresponding receiving device. If we view this system as a path for information to flow from the thermocouple's tip to the transmitter to the resistor and finally to the voltmeter/indicator, we see that the analog output range of each device must correspond to the analog input range of the *next* device, or else the real-world meaning of the analog signal will be lost.


This correspondence does not happen automatically, but must be established by the instrument technician building the system. In this case, it would be the technician's responsibility to properly adjust the range of the temperature transmitter, and also to ensure the indicator's display scale was properly labeled. Both the thermocouple and the resistor are non-adjustable devices, their input/output characteristics being fixed by physical laws.

10 volts DC is a common "dead zero" signal standard, although far more common in environmental (building heating and cooling) control systems than industrial control systems. I once encountered an old analog control system using -10 volts to +10 volts as its analog signal range, which meant 0 volts represented a 50% signal! A failed signal path in such a system could have been very misleading indeed, as a 50% signal value is not suspicious in the least.

²This is a temperature sensing element consisting of two different metal wires joined together, which generate a small voltage proportional to temperature. The correspondence between junction temperature and DC millivoltage is very well established by scientific testing, and so we may use this principle to sense process temperature.

DC current signals are also used in control systems to command the positioning of a final control element, such as a control valve or a variable-speed motor drive (VSD). In these cases, the milliamp value does not directly represent a process measurement, but rather how the degree to which the final control element influences the process. Typically (but not always!), 4 milliamps commands a closed (shut) control valve or a stopped motor, while 20 milliamps commands a wide-open valve or a motor running at full speed. Final control elements often are equipped with adjustable ranges so that an accurate correspondence between the analog signal and the desired control action may be ensured.


Thus, most industrial control systems use at least *two* different 4-20 mA signals: one to represent the process variable (PV) and one to represent the command signal to the final control element (the “manipulated variable” or MV):



The relationship between these two signals depends entirely on the response of the controller. There is no reason to ever expect the PV and MV current signals to be equal to each other except by chance, for they represent entirely different variables. In fact, if the controller is reverse-acting, it is entirely normal for the two current signals to be inversely related: as the PV signal increases going to a reverse-acting controller, the output signal will decrease. If the controller is placed into “manual” mode by a human operator, the output signal will have no automatic relation to the PV signal at all, instead being entirely determined by the operator’s whim.

13.2 Relating 4 to 20 mA signals to instrument variables

A 4 to 20 mA current signal represents some signal along a 0 to 100 percent scale. Usually, this scale is linear as shown by this graph:



Being a linear function, we may use the standard slope-intercept linear equation to relate signal percentage to current values:

$$y = mx + b$$

Where,

y = Output from instrument

x = Input to instrument


m = Slope

b = y -intercept point (i.e. the *live zero* of the instrument's range)

Once we determine suitable values for m and b , we may then use this linear equation to predict any value for y given x , and vice-versa. This is very useful for predicting the 4-20 mA signal output of a process transmitter, or the expected stem position of a 4-20 mA controlled valve, or any other correspondence between a 4-20 mA signal and some physical variable.

Before we may use this equation for any practical purpose, we must determine the slope (m) and intercept (b) values appropriate for the instrument we wish to apply the equation to. Next, we will see some examples of how to do this.

For the linear function shown, we may determine the slope value (m) by dividing the line's *rise* by its *run*. Two sets of convenient points we may use in calculating rise over run are 4 and 20 milliamps (for the rise), and 0 and 100 percent (for the run):



$$m = \frac{\text{Rise}}{\text{Run}} = \frac{(20 - 4)}{(100 - 0)} = \frac{16}{100}$$

$$y = \left(\frac{16}{100} \right) x + b$$

To calculate the y-intercept (b), all we need to do is solve for b at some known coordinate pair of x and y . Again, we find convenient points³ for this task at 0 percent and 4 milliamps:

$$4 = \left(\frac{16}{100} \right) 0 + b$$

$$4 = 0 + b$$

$$b = 4$$

³We could have just as easily chosen 100 percent for x and 20 milliamps for y , for it would have yielded the same result of $b = 4$.

Now we have a complete formula for converting a percentage value into a milliamp value:

$$y = \left(\frac{16}{100} \right) x + 4$$

We may now use this formula to calculate how many milliamps represent any given percentage of signal. For example, suppose we needed to convert a percentage of 34.7% into a corresponding 4-20 mA current. We would do so like this:

$$y = \left(\frac{16}{100} \right) 34.7 + 4$$


$$y = 5.552 + 4$$

$$y = 9.552$$

Thus, 34.7% is equivalent to 9.552 milliamps in a 4-20 mA signal range.

The slope-intercept formula for linear functions may be applied to *any* linear instrument, as illustrated in the following examples.

13.2.1 Example calculation: controller output to valve



An electronic loop controller outputs a signal of 8.55 mA to a direct-responding control valve (where 4 mA is shut and 20 mA is wide open). How far open should the control valve be at this MV signal level?

To solve for percentage of stem travel (x) at 8.55 milliamps of signal current (y), we may use the linear equation developed previously to predict current in milliamps (y) from signal value in percent (x):

$$y = \left(\frac{16}{100} \right) x + 4$$

$$8.55 = \left(\frac{16}{100} \right) x + 4$$


$$4.55 = \left(\frac{16}{100} \right) x$$

$$\left(\frac{100}{16} \right) 4.55 = x$$

$$x = 28.4$$

Therefore, we should expect the valve to be 28.4% open at an applied MV signal of 8.55 milliamps.

13.2.2 Example calculation: flow transmitter



A flow transmitter is ranged 0 to 350 gallons per minute, 4-20 mA output, direct-responding. Calculate the current signal value at a flow rate of 204 GPM.

One way we could solve for the amount of signal current is to convert the flow value of 204 GPM into a ratio of the flowmeter's full-flow value, then apply the same formula we used in the previous example relating percentage to milliamps. Converting the flow rate into a "per unit" ratio is a matter of simple division, since the flow measurement range is zero-based:

$$\frac{204 \text{ GPM}}{350 \text{ GPM}} = 0.583 \text{ per unit}$$

Converting a "per unit" ratio into percent merely requires multiplication by 100, since "percent" literally means "per 100":

$$0.583 \text{ per unit} \times 100\% = 58.3\%$$

Next, we plug this percentage value into the formula:


$$y = \left(\frac{16}{100} \right) 58.3 + 4$$

$$y = 9.33 + 4$$

$$y = 13.33$$

Therefore, the transmitter should output a PV signal of 13.3 mA at a flow rate of 204 GPM.

An alternative approach is to set up a linear equation specifically for this flowmeter given its measurement range (0 to 350 GPM) and output signal range (4 to 20 mA). We will begin this process by sketching a simple graph relating flow rate to current:



The slope (m) for this equation is rise over run, in this case 16 milliamps of rise for 350 GPM of run:

$$y = \left(\frac{20 - 4}{350 - 0} \right) x + b = \left(\frac{16}{350} \right) x + b$$

The y-intercept for this equation is 4, since the current output will be 4 milliamps at zero flow:

$$y = \left(\frac{16}{350} \right) x + 4$$

Now that the linear equation is set up for this particular flowmeter, we may plug in the 204 GPM value for x and solve for current:


$$y = \left(\frac{16}{350} \right) 204 + 4$$

$$y = 9.33 + 4$$

$$y = 13.33$$


Just as before, we arrive at a current of 13.33 milliamps representing a flow rate of 204 GPM.

13.2.3 Example calculation: temperature transmitter



An electronic temperature transmitter is ranged 50 to 140 degrees Fahrenheit and has a 4-20 mA output signal. Calculate the current output by this transmitter if the measured temperature is 79 degrees Fahrenheit.

First, we will set up a linear equation describing this temperature transmitter's function:



Calculating and substituting the slope (m) value for this equation, using the full rise-over-run of the linear function:

$$y = \left(\frac{20 - 4}{140 - 50} \right) x + b = \left(\frac{16}{90} \right) x + b$$

The y-intercept value will be different⁴ for this example than it was for previous examples, since the measurement range is not zero-based. However, the procedure for finding this value is the same – plug any corresponding x and y values into the equation and solve for b . In this case, I will use the values of 4 mA for y and 50 °F for x :

$$4 = \left(\frac{16}{90}\right) 50 + b$$

$$4 = 8.89 + b$$

$$b = -4.89$$

Therefore, our customized linear equation for this temperature transmitter is as follows:

$$y = \left(\frac{16}{90}\right) x - 4.89$$

At a sensed temperature of 79 °F, the transmitter’s output current will be 9.16 mA:

$$y = \left(\frac{16}{90}\right) 79 - 4.89$$

$$y = 14.04 - 4.89$$

$$y = 9.16$$

⁴A common misconception for people learning to apply the slope-intercept formula to linear instrument ranges is that they tend to assume b will always be equal to the lower-range value (LRV) of the instrument’s output range. For example, given a transmitter with a 4-20 mA output range, the false assumption is that $b = 4$. This does happen to be true *only if the instrument possesses a “dead-zero” input range*, but it will not be true for instruments with a live-zero input range such in this case here where the temperature input range is 50 to 140 degrees.

We may apply the same alternative method of solution to this problem as we did for the flowmeter example: first converting the process variable into a simple “per unit” ratio or percentage of measurement range, then using that percentage to calculate current in milliamps. The “tricky” aspect of this example is the fact the temperature measurement range does not begin at zero.

Converting 79 °F into a percentage of a 50-to-140 °F range requires that we first subtract the live-zero value, then divide by the span:

$$\text{Per unit ratio} = \frac{79 - 50}{140 - 50} = 0.3222$$

$$\text{Percentage} = 0.3222 \text{ per unit} \times 100\% = 32.22\%$$

Next, plugging this percentage value into our standard linear equation for 4-20 mA signals:

$$y = \left(\frac{16}{100} \right) 32.22 + 4$$

$$y = 5.16 + 4$$


$$y = 9.16$$

Again, we arrive at the exact same figure for transmitter output current: 9.16 milliamps at a measured temperature of 79 °F.

The choice to calculate transmitter current by first setting up a “customized” linear equation for the transmitter in question or by converting the measured value into a percentage and using a “standard” linear equation for current is arbitrary. Either method will produce accurate results, although it could be argued that the “customized equation” approach may save time if many different current values must be calculated.


Certainly, if you are programming a computer to convert a received milliamp signal value into a measurement range (such as degrees Fahrenheit), it makes more sense to have the computer evaluate a single equation rather than perform multiple steps of calculations as we do when using percentage values as an intermediate step between the input and output value calculations. Evaluating one equation rather than two saves processing time.

13.2.4 Example calculation: pH transmitter



A pH transmitter has a calibrated range of 4 pH to 10 pH, with a 4-20 mA output signal. Calculate the pH sensed by the transmitter if its output signal is 11.3 mA.

First, we will set up a linear equation describing this temperature transmitter's function:



Note how we are free to set up 4-20 mA as the independent variable (x axis) and the pH as the dependent variable (y axis). We could arrange current on the y axis and the process measurement on the x axis as before, but this would force us to manipulate the linear equation to solve for x .

Calculating and substituting the slope (m) value for this equation, using the full rise-over-run of the linear function:

$$y = \left(\frac{10 - 4}{20 - 4} \right) x + b = \left(\frac{6}{16} \right) x + b$$

Solving for the y-intercept value using the coordinate values of 4 pH and 4 mA, we see again that this is an application where $b \neq 4$ mA. This is due to the fact that the instrument's input range (i.e. the domain of the $y = mx + b$ function) does *not* begin at zero:

$$4 = \left(\frac{6}{16}\right)4 + b$$

$$4 = 1.5 + b$$

$$b = 2.5$$

Therefore, our customized linear equation for this pH transmitter is as follows:

$$y = \left(\frac{6}{16}\right)x + 2.5$$

Calculating the corresponding pH value for an output current signal of 11.3 mA now becomes a very simple matter:


$$y = \left(\frac{6}{16}\right)11.3 + 2.5$$

$$y = 4.24 + 2.5$$

$$y = 6.74$$


Therefore, the transmitter's 11.3 mA output signal reflects a measured pH value of 6.74 pH.

13.2.5 Example calculation: reverse-acting I/P transducer signal



A current-to-pressure transducer is used to convert a 4-20 mA electronic signal into a 3-15 PSI pneumatic signal. This particular transducer is configured for **reverse action** instead of direct, meaning that its pressure output at 4 mA should be 15 PSI and its pressure output at 20 mA should be 3 PSI. Calculate the necessary current signal value to produce an output pressure of 12.7 PSI.

Reverse-acting instruments are still linear, and therefore still follow the slope-intercept line formula $y = mx + b$, albeit with a negative slope:



Calculating and substituting the slope (m) value for this equation, using the full rise-over-run of the linear function. Note how the “rise” is actually a “fall” from 20 milliamps down to 4 milliamps, yielding a negative value for m :

$$y = \left(\frac{4 - 20}{15 - 3} \right) x + b = \left(\frac{-16}{12} \right) x + b = \left(-\frac{16}{12} \right) x + b$$

Solving for the y-intercept value using the coordinate values of 3 PSI and 20 mA:

$$20 = \left(-\frac{16}{12} \right) 3 + b$$

$$20 = -4 + b$$

$$b = 24$$

Therefore, our customized linear equation for this I/P transducer is as follows:

$$y = \left(-\frac{16}{12} \right) x + 24$$

Calculating the corresponding current signal for an output pressure of 12.7 PSI:

$$y = \left(-\frac{16}{12} \right) 12.7 + 24$$


$$y = -16.93 + 24$$

$$y = 7.07$$

Therefore, a current signal of 7.07 mA is necessary to drive the output of this reverse-acting I/P transducer to a pressure of 12.7 PSI.


13.2.6 Example calculation: PLC analog input scaling

Allen-Bradley SLC 500 controller



An Allen-Bradley SLC500 programmable logic controller (PLC) uses a 16-bit analog-to-digital converter in its model 1746-NI4 analog input card to convert 4-20 mA signals into digital number values ranging from 3277 (at 4 mA) to 16384 (at 20 mA). However, these raw numbers from the PLC's analog card must be mathematically scaled inside the PLC to represent real-world units of measurement, in this case 0 to 700 GPM of flow. Formulate a scaling equation to program into the PLC so that 4 mA of current registers as 0 GPM, and 20 mA of current registers as 700 GPM.

We are already given the raw number values from the analog card's analog-to-digital converter (ADC) circuit for 4 mA and 20 mA: 3277 and 16384, respectively. These values define the domain of our linear graph:



Calculating and substituting the slope (m) value for this equation, using the full rise-over-run of the linear function:

$$y = \left(\frac{700 - 0}{16384 - 3277} \right) x + b = \left(\frac{700}{13107} \right) x + b$$

Solving for the y-intercept value using the coordinate values of 0 GPM and 3277 ADC counts:

$$0 = \left(\frac{700}{13107} \right) 3277 + b$$

$$0 = 175 + b$$

$$b = -175$$

Therefore, our PLC scaling equation for this particular flowmeter is as follows:

$$y = \left(\frac{700}{13107} \right) x - 175$$


This type of scaling calculation is so common in PLC applications that Allen-Bradley has provided a special SCL (“scale”) instruction just for this purpose. Instead of “slope” (m) and “intercept” (b), the instruction prompts the human programmer to enter “rate” and “offset” values, respectively. Furthermore, the rate in Allen-Bradley’s SCL instruction is expressed as the numerator of a fraction where the denominator is fixed at 10000, allowing fractional (less than one) slope values to be specified using integer numbers. Aside from these details, the concept is exactly the same.

Expressing our slope of $\frac{700}{13107}$ as a fraction with 10000 as the denominator is a simple matter of solving for the numerator using cross-multiplication and division:

$$\frac{700}{13107} = \frac{r}{10000}$$

$$r = 534$$


Thus, the SCL instruction would be configured as follows⁵




⁵The “Source” and “Dest” parameters shown in this instruction box refer to special addresses in the PLC’s memory where the input (ADC count) and output (scaled flowrate) values will be found. You need not concern yourself with the meanings of I:4.2 and N7:15, because these addresses are unimportant to the task of deriving a scaling formula.

13.2.7 Graphical interpretation of signal ranges

An illustration some students find helpful in understanding analog signal ranges is to consider the signal range as a *length* expressed on a number line. For example, the common 4-20 mA analog current signal range would appear as such:



If one were to ask the percentage corresponding to a 14.4 mA signal on a 4-20 mA range, it would be as simple as determining the length of a line segment stretching from the 4 mA mark to the 14.4 mA mark:




As a percentage, this thick line is 10.4 mA long (the distance between 14.4 mA and 4 mA) over a total (possible) length of 16 mA (the total span between 20 mA and 4 mA). Thus:

$$\text{Percentage} = \left(\frac{14.4 \text{ mA} - 4 \text{ mA}}{20 \text{ mA} - 4 \text{ mA}} \right) 100\%$$

$$\text{Percentage} = \left(\frac{10.4 \text{ mA}}{16 \text{ mA}} \right) 100\%$$

$$\text{Percentage} = 65\%$$


This same “number line” approach may be used to visualize any conversion from one analog scale to another. Consider the case of an electronic pressure transmitter calibrated to a pressure range of -5 to $+25$ PSI, having an (obsolete) current signal output range of 10 to 50 mA. The appropriate current signal value for an applied pressure of $+12$ PSI would be represented on the number line as such:



Finding the “length” of this line segment in units of milliamps is as simple as setting up a proportion between the length of the line in units of PSI over the total (span) in PSI, to the length of the line in units of mA over the total (span) in mA:

$$\frac{17 \text{ PSI}}{30 \text{ PSI}} = \frac{? \text{ mA}}{40 \text{ mA}}$$

Solving for the unknown (?) current by cross-multiplication and division yields a value of 22.67 mA. Of course, this value of 22.67 mA only tells us the length of the line segment on the number line; it does not directly tell us the current signal value. To find that, we must add the “live zero” offset of 10 mA, for a final result of 32.67 mA.




Thus, an applied pressure of $+12$ PSI to this transmitter should result in a 32.67 mA output signal.

13.2.8 Thinking in terms of per unit quantities


Although it is possible to generate a “custom” linear equation in the form of $y = mx + b$ for any linear-responding instrument relating input directly to output, a more general approach may be used to relate input to output values by translating all values into (and out of) *per unit* quantities. A “per unit” quantity is simply a ratio between a given quantity and its maximum value. A half-full glass of water could thus be described as having a fullness of 0.5 *per unit*. The concept of percent (“per one hundred”) is very similar, the only difference between per unit and percent being the base value of comparison: half-full glass of water has a fullness of 0.5 per unit (i.e. $\frac{1}{2}$ of the glass’s full capacity), which is the same thing as 50 percent (i.e. 50 on a scale of 100, with 100 representing complete fullness).

Let’s now apply this concept to a realistic 4-20 mA signal application. Suppose you were given a liquid level transmitter with an input measurement range of 15 to 85 inches and an output range of 4 to 20 milliamps, respectively, and you desired to know how many milliamps this transmitter should output at a measured liquid level of 32 inches. Both the measured level and the milliamp signal may be expressed in terms of *per unit* ratios, as shown by the following graphs:



So long as we choose to express process variable and analog signal values as a per unit ratios ranging from 0 to 1, we see how m (the slope of the line) is simply equal to the span of the process variable or analog signal range, and b is simply equal to the lower-range value (LRV) of the process variable or analog signal range. The advantage of thinking in terms of “per unit” is the ability to quickly and easily write linear equations for any given range. In fact, this is so easy that we don’t even have to use a calculator to compute m in most cases, and we *never* have to calculate b because the LRV is explicitly given to us. The instrument’s input equation is $y = 70x + 15$ because the span of the 15-to-85 inch range is 70, and the LRV is 15. The instrument’s output equation is $y = 16x + 4$ because the span of the 4-to-20 milliamp range is 16, and the LRV is 4.

If we manipulate each of the $y = mx + b$ equations to solve for x (per unit of span), we may express the relationship between the input and output of any linear instrument as a pair of fractions with the per unit value serving as the proportional link between input and output:



$$\frac{\text{Input} - \text{LRV}_{\text{in}}}{\text{URV}_{\text{in}} - \text{LRV}_{\text{in}}} = \begin{array}{c} \text{Per unit of span} \\ \text{(from 0 to 1 inclusive)} \end{array} = \frac{\text{Output} - \text{LRV}_{\text{out}}}{\text{URV}_{\text{out}} - \text{LRV}_{\text{out}}}$$

The question remains, how do we apply these equations to our example problem: calculating the milliamp value corresponding to a liquid level of 32 inches for this instrument? The answer to this question is that we must perform a *two-step* calculation: first, convert 32 inches into a per unit ratio, then convert that per unit ratio into a milliamp value.

First, the conversion of inches into a per unit ratio, knowing that 32 is the value of y and we need to solve for x :

$$32 = 70x + 15$$

$$32 - 15 = 70x$$

$$\frac{32 - 15}{70} = x$$

$$x = 0.2429 \text{ per unit (i.e. 24.29\%)}$$

Next, converting this per unit ratio into a corresponding milliamp value, knowing that y will now be the current signal value using m and b constants appropriate for the 4-20 milliamp range:

$$y = 16x + 4$$

$$y = 16(0.2429) + 4$$

$$y = 3.886 + 4$$

$$y = 7.886 \text{ mA}$$

Instead of deriving a single custom $y = mx + b$ equation directly relating input (inches) to output (milliamps) for every instrument we encounter, we may use two simple and generic linear equations to do the calculation in two steps with “per unit” being the intermediate result. Expressed in general form, our linear equation is:

$$y = mx + b$$

$$\text{Value} = (\text{Span})(\text{Per unit}) + \text{LRV}$$

$$\text{Value} = (\text{URV} - \text{LRV})(\text{Per unit}) + \text{LRV}$$

Thus, to find the per unit ratio we simply take the value given to us, subtract the LRV of its range, and divide by the span of its range. To find the corresponding value we take this per unit ratio, multiply by the span of the other range, and then add the LRV of the other range.

Example: Given a pressure transmitter with a measurement range of 150 to 400 PSI and a signal range of 4 to 20 milliamps, calculate the applied pressure corresponding to a signal of 10.6 milliamps.

Solution: Take 10.6 milliamps and subtract the LRV (4 milliamps), then divide by the span (16 milliamps) to arrive at 41.25% (0.4125 per unit). Take this number and multiply by the span of the pressure range (400 PSI – 150 PSI, or 250 PSI) and lastly add the LRV of the pressure range (150 PSI) to arrive at a final answer of 253.125 PSI.

Example: Given a temperature transmitter with a measurement range of –88 degrees to +145 degrees and a signal range of 4 to 20 milliamps, calculate the proper signal output at an applied temperature of +41 degrees.


Solution: Take 41 degrees and subtract the LRV (–88 degrees) which is the same as *adding* 88 to 41, then divide by the span (145 degrees – (–88) degrees, or 233 degrees) to arrive at 55.36% (0.5536 per unit). Take this number and multiply by the span of the current signal range (16 milliamps) and lastly add the LRV of the current signal range (4 milliamps) to arrive at a final answer of 12.86 milliamps.

Example: Given a pH transmitter with a measurement range of 3 pH to 11 pH and a signal range of 4 to 20 milliamps, calculate the proper signal output at 9.32 pH.

Solution: Take 9.32 pH and subtract the LRV (3 pH), then divide by the span (11 pH – 3 pH, or 8 pH) to arrive at 79% (0.79 per unit). Take this number and multiply by the span of the current signal range (16 milliamps) and lastly add the LRV of the current signal range (4 milliamps) to arrive at a final answer of 16.64 milliamps.

13.3 Controller output current loops

The simplest form of 4-20 mA current loop is the type used to represent the output of a process controller, sending a command signal to a final control element. Here, the controller supplies both the electrical power and signal information to the final control element, which acts as an electrical load. To illustrate, consider the example of a controller sending a 4-20 mA signal to an I/P (current-to-pressure) signal converter, which then pneumatically drives a control valve:



This particular controller has two digital displays, one for process variable (PV) and one for setpoint (SP), with a bargraph for displaying the output value (Out). One pushbutton provides the operator with a way to switch between Automatic and Manual modes (A/M), while two other pushbuttons provide means to decrement and increment either the setpoint value (in Automatic mode) or the Output value (in Manual mode).


Inside the controller, a *dependent current source* provides the 4-20 mA DC current signal to the I/P transducer. Like all current sources, its purpose is to maintain current in the “loop” circuit regardless of circuit resistance or any external voltage sources. Unlike a constant current source, a “dependent” current source (represented by a diamond shape instead of a circle shape) varies its current value according to some external stimulus. In this case, either the mathematical function of the controller (Automatic mode) or the arbitrary action of the human operator (Manual mode) tells the current source how much DC current it should maintain in the circuit.

For example, if the operator happened to switch the controller into Manual mode and set the output value at 50%, the proper amount of DC current for this signal percentage would be 12 mA (exactly half-way between 4 mA and 20 mA). If everything is working properly, the current in the “loop” circuit to the I/P transducer should remain exactly at 12 mA regardless of slight changes in wire resistance, I/P coil resistance, or anything else: the current source inside the controller will “fight” as hard as it has to in order to maintain this set amount of current. This current, as it flows

through the wire coil of the I/P transducer mechanism, creates a magnetic field inside the I/P to actuate the pneumatic mechanism and produce a 9 PSI pressure signal output to the control valve (9 PSI being exactly half-way between 3 PSI and 15 PSI in the 3-15 PSI signal standard range). This should move the control valve to the half-way position.

The details of the controller's internal current source are not terribly important. Usually, it takes the form of an operational amplifier circuit driven by the voltage output of a DAC (Digital-to-Analog Converter). The DAC converts a binary number (either from the controller's automatic calculations, or from the human operator's manual setting) into a small DC voltage, which then commands the opamp circuit to regulate output current at a proportional value.

The scenario is much the same if we replace the I/P and control valve with a variable-speed motor drive. From the controller's perspective, the only difference it sees is a resistive load instead of an inductive load. The input resistance of the motor drive circuit converts the 4-20 mA signal into an analog voltage signal (typically 1-5 V, but not always). This voltage signal then commands the motor drive circuitry, telling it to modulate the power going to the electric motor in order to drive it at the desired speed:




Here, the variable-speed motor drive is a high-power electronic circuit that takes 3-phase AC power in and converts it to DC power at a variable voltage (that voltage level controlled by the 4-20 mA signal from the controller output). Another style of motor "drive" is one that takes in AC power and outputs 3-phase AC power at variable voltage and frequency to power an AC induction motor. This latter style is usually called a *variable frequency drive* or *VFD*, but it looks the same to the controller output: a fixed resistive load accepting the 4-20 mA output signal.

In some process applications the final control element is “reverse-acting” in that a controller output current value of 4 mA tells the FCE to go to its “full” (100%) capacity and a controller output current value of 20 mA commands the FCE to go to its minimum (0%) capacity. Such is the case with an I/P transducer driving an “air-to-close” pneumatic control valve, where the valve’s spring works to push the throttling mechanism open and air pressure on the valve diaphragm works to push the throttling mechanism shut. A practical reason for designing a control system like this is if the dictates of process safety require the valve go wide-open if ever the instrument air supply fails or the 4-20 mA output signal circuit fails.

In such applications, we need to configure the controller in such a way that the output display (digital read-out and/or bargraph) is *reverse-indicating* so as to avoid confusing any human operator using the controller. Since 4 mA represents a wide-open control valve and 20 mA represents a fully shut control valve, a reverse-indicating controller will display 0% output when the current signal is 20 mA and 100% output when the current signal is 4 mA (i.e. the controller display is a direct representation of *control valve stem position*, not of the current signal itself):


Air-to-Close valve with controller configured for “reverse indication” on its output display



It should be noted that direct or reverse *indication* on the controller’s faceplate is unrelated to direct or reverse *control action* of the controller’s algorithm. The former is merely the way in which the output signal is graphically represented to the human operator while the latter is the relationship between the process variable (PV) signal and the controller’s output (MV) signal necessary for negative-feedback control. For the sake of argument, this controller’s automatic action could still be *direct* (i.e. greater PV signal = greater output signal) even though its output bargraph indication is *reverse* in order to faithfully show valve stem position to the human operator.

13.4 4-wire (“self-powered”) transmitter current loops


DC electric current signals may also be used to communicate process measurement information from transmitters to controllers, indicators, recorders, alarms, and other input devices. Recall that the purpose of a *transmitter* is to sense some physical variable (e.g. pressure, temperature, flow) and then report that quantity in the form of a signal, in this case a 4 to 20 milliamp DC current proportional to that measured quantity. The simplest form of 4-20 mA measurement loop is one where the transmitter has two terminals for the 4-20 mA signal wires to connect, and two more terminals where a power source connects. These transmitters are called “4-wire” or “self-powered” units. The current signal from the transmitter connects to the *process variable input* terminals of the controller to complete the loop:



Some process controllers are not equipped to directly accept milliamp input signals, but rather can only interpret DC voltage signals. In such cases we must connect a precision resistor across the controller’s input terminals to convert the 4-20 mA transmitter signal into a standardized analog voltage signal the controller can understand. A voltage signal range of 1 to 5 volts is standard, although some models of controller use different voltage ranges and therefore require different precision resistor values. If the voltage range is 1-5 volts and the current range is 4-20 mA, the precision resistor value must be 250 ohms according to Ohm’s Law.

Since this is a digital controller, the input voltage at the controller terminals is interpreted by an analog-to-digital converter (ADC) circuit, which converts the measured voltage into a digital number the controller’s microprocessor can interpret.


In some installations, transmitter power is supplied through additional wires in the cable from a power source located near the controller:



The obvious disadvantage of this scheme is the requirement of two more conductors in the cable. More conductors means the cable will be larger-diameter and more expensive for a given length. Cables with more conductors will require larger electrical conduit to fit in to, and all field wiring panels will have to contain more terminal blocks to marshal the additional conductors. If no suitable electrical power source exists at the transmitter location, though, a 4-wire cable is necessary to service a 4-wire transmitter.

13.5 2-wire (“loop-powered”) transmitter current loops

It is possible to convey electrical power *and* communicate analog information over the same two wires using 4 to 20 milliamps DC, if we design the transmitter to be *loop-powered*. A loop-powered transmitter connects to a process controller with only two wires, which is why loop-powered transmitters are synonymously known as *2-wire transmitters*:




Here, the transmitter is not really a current *source* in the sense that a 4-wire transmitter is. Instead, a 2-wire transmitter’s circuitry is designed to act as a current *regulator*, limiting current in the series loop to a value representing the process measurement, while relying on a remote source of power to motivate the electric current. Please note the direction of the arrow in the transmitter’s dependent current source symbol, and how it relates to the voltage polarity marks. Refer back to the illustration of a 4-wire transmitter circuit for comparison. The current “source” in this loop-powered transmitter actually behaves as an electrical *load*⁶, while the current source in the 4-wire transmitter functioned as a true electrical source.

A loop-powered transmitter gets its operating power from the minimum terminal voltage and current available at its two terminals. With the typical source voltage being 24 volts DC, and the maximum voltage dropped across the controller’s 250 ohm resistor being 5 volts DC, the transmitter should always have at least 19 volts available at its terminals. Given the lower end of the 4-20 mA signal range, the transmitter should always have at least 4 mA of current to function on. Thus, the transmitter will always have a certain minimum amount of electrical power available on which to operate, while regulating current to signal the process measurement to the receiving instrument.

⁶Some of my students have referred to such a circuit as a *smart load*, since it functions as a load but nevertheless exerts control over the circuit current.

Internally, the electronic hardware of a 2-wire transmitter circuitry resembles the following (simplified) diagram. Note that everything shown within the shaded rectangle is represented by the “2-wire transmitter” circle in the previous diagram:



All sensing, scaling, and output conditioning circuitry inside the transmitter must be designed to operate on less than 4 mA of DC current, and at a modest terminal voltage. In order to create loop currents exceeding 4 mA – as the transmitter must do in order to span the entire 4 to 20 milliamp signal range – the transmitter circuitry uses a transistor to shunt (bypass) extra current from one terminal to the other as needed to make the total current indicative of the process measurement. For example, if the transmitter’s internal operating current is only 3.8 mA, and it must regulate loop current at a value of 16 mA to represent a condition of 75% process measurement, the shunt transistor will be driven by the opamp to bypass exactly 12.2 mA of current (because 3.8 mA + 12.2 mA = 16.0 mA).


The very low amount of electrical power available at a 2-wire transmitter’s terminals limits its functionality. If the transmitter requires more electrical power than can be delivered with 4 milliamps and 19 volts (minimum each), the only solution is to go with a 4-wire transmitter where the power conductors are separate from the signal conductors. An example of a process transmitter that must be 4-wire is a chemical analyzer such as a chromatograph, requiring enough power to operate an electrical heater, solenoid valves, and an on-board computer to process the sensor data. There is simply no way to operate a machine as complex and power-hungry as a 2010-era chromatograph on 4 milliamps and 19 volts!

Early current-based industrial transmitters were not capable of operating on such low levels of electrical power, and so used a different current signal standard: 10 to 50 milliamps DC. Loop power supplies for these transmitters ranged upwards of 90 volts to provide enough power for the transmitter. Safety concerns made the 10-50 mA standard unsuitable for some industrial

installations, and modern microelectronic circuitry with its reduced power consumption made the 4-20 mA standard practical for nearly all types of process transmitters.

13.6 4-wire "passive" versus "active" output transmitters

Some self-powered (4-wire) analog electronic transmitters are designed to behave as electrical loads rather than as electrical sources. Such transmitters are commonly referred to as having *passive* or *sinking* 4-20 mA outputs, as opposed to the *active* or *sourcing* 4-wire transmitters previously described:



At first this seems needlessly confusing. Why build a self-powered transmitter requiring a second power supply in the circuit to drive the 4-20 mA signal? The reason for this type of transmitter's existence is the sheer popularity of loop-powered 2-wire 4-20 mA transmitters. Loop-powered field instruments have become so popular in industry that many control systems, PLCs, indicators, and other receiving devices have their own loop power supply built in so that these systems can *only* connect to loads and are therefore incompatible with current-sourcing 4-wire transmitters. Thus, "passive" or "sinking" 4-wire transmitters were developed to accommodate control systems designed to work (only) with loop-powered instruments.


Some 4-wire transmitters are configurable for either passive or active (i.e. sinking versus sourcing) operation, requiring the installing technician to pay close attention in order for the circuit to properly function.

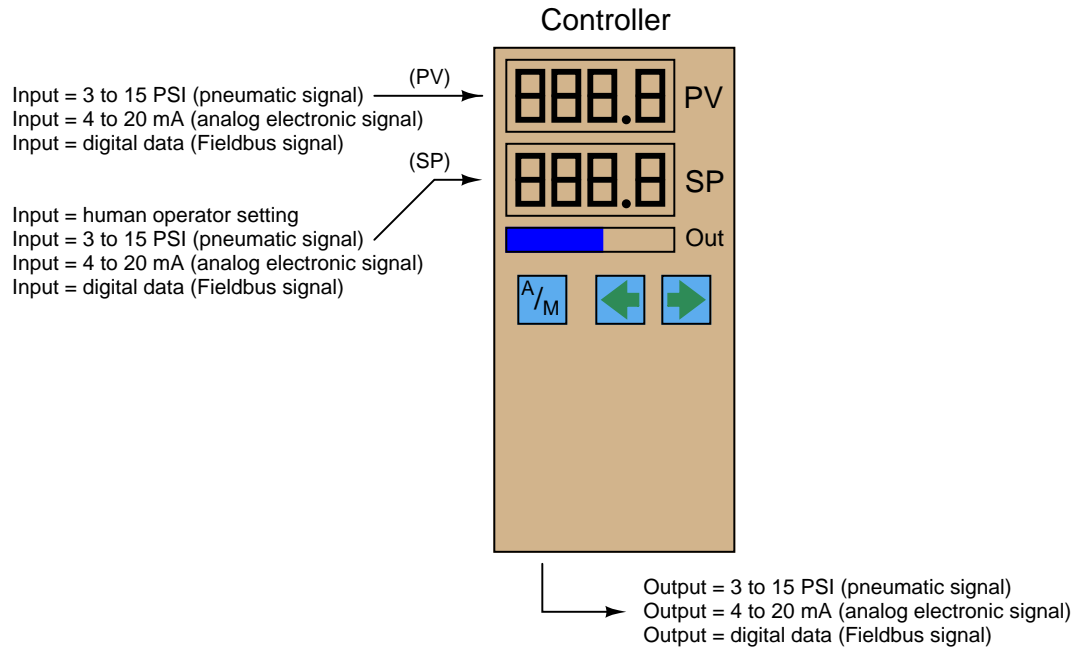
13.7 Troubleshooting current loops

A fundamental principle in instrumentation system troubleshooting is that every instrument has at least one input and at least one output, and that the output(s) should accurately correspond to the input(s). If an instrument's output is not properly corresponding to its input according to the instrument's design function, there must be something wrong with that instrument.

Consider the inputs and outputs of several common instruments: transmitters, controllers, indicators, and control valves. Each of these instruments takes in (input) data in some form, and generates (output) data in some form. In any instrument "loop," the output of one instrument feeds into the input of the next, such that information passes from one instrument to another. By intercepting the data communicated between components of an instrument system, we are able to locate and isolate faults. In order to properly understand the intercepted data, we must understand the inputs and outputs of the respective instruments and the basic functions of those instruments.

The following illustrations highlight inputs and outputs for instruments commonly found in control systems:






In order to check for proper correspondence between instrument inputs and outputs, we must be able to use appropriate test equipment to intercept the signals going into and out of those instruments. For 4-20 mA analog signal-based instruments, this means we must be able to use electrical meters capable of accurately measuring current and voltage.

13.7.1 Using a standard milliammeter to measure loop current

Since the signal of interest is represented by an electric current in an instrumentation current “loop” circuit, the obvious tool to use for troubleshooting is a multimeter capable of accurately measuring DC milliamperes. Unfortunately, though, there is a major disadvantage to the use of a milliammeter: the circuit must be “broken” at some point to connect the meter in series with the current, and this means the current will fall to 0 mA until the meter is connected (then fall to 0 mA when the meter is removed from the circuit). Interrupting the current means interrupting the flow of information conveyed by that current, be it a process measurement or a command signal to a final control element. This *will* have adverse effects on a control system unless certain preparatory steps are taken.

Before “breaking the loop” to connect your meter, one must first warn all appropriate personnel that the signal will be interrupted at least twice, falling to a value of -25% each time. If the signal to be interrupted is coming from a process transmitter to a controller, the controller should be placed in Manual mode so it will not cause an upset in the process (by moving the final control element in response to the sudden loss of PV signal). Also, process alarms should be temporarily disabled so they do not cause panic. If this current signal also drives process shutdown alarms, these should be temporarily disabled so that nothing shuts down upon interruption of the signal.

If the current signal to be interrupted is a command signal from a controller to a final control element, the final control element either needs to be manually overridden so as to hold a fixed setting while the signal varies, or it needs to be bypasses completely by some other device(s). If the final control element is a control valve, this typically takes the form of opening a bypass valve and closing at least one block valve:



Since the manually-operated bypass valve now performs the job the automatic control valve used to do, a human operator must remain posted at the bypass valve to carefully throttle it and maintain control of the process.

Block and bypass valves for a large gas flow control valve may be seen in the following photograph:



In consideration of the labor necessary to safely interrupt the current signal to a control valve in a live process, we see that the seemingly simple task of connecting a milliammeter in series with a 4-20 mA current signal is not as easy as it may first appear. Better ways must exist, no?

13.7.2 Using a clamp-on milliammeter to measure loop current


One better way to measure a 4-20 mA signal without interrupting it is to do so magnetically, using a clamp-on milliammeter. Modern Hall-effect sensors are sensitive and accurate enough to monitor the weak magnetic fields created by the passage of small DC currents in wires. Ammeters using Hall-effect sensors have are completely non-intrusive because they merely clamp around the wire, with no need to “break” the circuit. An example of a such a clamp-on current meter is the Fluke model 771, shown in this photograph:




Note how this milliammeter not only registers loop current (3.98 mA as shown in the photograph), but it also converts the milliamp value into a percentage of range, following the 4 to 20 mA signal standard. One disadvantage to be aware of for clamp-on milliammeters is the susceptibility to error from strong external magnetic fields. Steady magnetic fields (from permanent magnets or DC-powered electromagnets) may be compensated for by performing a “zero” adjustment with the instrument held in a similar orientation prior to measuring loop current through a wire.

13.7.3 Using “test” diodes to measure loop current

Another way to measure a 4-20 mA signal without interrupting it involves the use of a rectifying diode, originally installed in the loop circuit when it was commissioned. A “test” diode may be placed anywhere in series within the loop in such a way that it will be forward-biased. During normal operation, the diode will drop approximately 0.7 volts, as is typical for any silicon rectifying diode when forward biased. The following schematic diagram shows such a diode installed in a 2-wire transmitter loop circuit:



If someone connects a milliammeter in parallel with this diode, however, the very low input resistance of the ammeters “shorts past” the diode and prevents any substantial voltage drop from forming across it. Without the necessary forward voltage drop, the diode effectively turns off and conducts 0 mA, leaving the entire loop current to pass through the ammeter:



When the milliammeter is disconnected, the requisite 0.7 volt drop appears to turn on the diode, and all loop current flows through the diode again. At no time is the loop current ever interrupted, which means a technician may take current measurements this way and never have to worry about generating false process variable indications, setting off alarms, or upsetting the process.

Such a diode may be installed at the nearest junction box, between terminals on a terminal strip, or even incorporated into the transmitter itself. Some process transmitters have an extra pair of terminals labeled “Test” for this exact purpose. A diode is already installed in the transmitter, and these “test” terminals serve as points to connect the milliammeter across.

The following photograph shows an example of this on a Rosemount model 3051 differential pressure transmitter:




Note the two test points labeled “TEST” below and to the right of the main screw terminals where the loop wiring attaches. Connecting an ammeter to these two test points allows for direct measurement of the 4-20 mA current signal without having to un-do any wire connections in the circuit.

Transmitters equipped with analog meter movements for direct visual indication of the 4-20 mA signal usually connect the analog milliammeter in parallel with just such a diode. The reason for doing this is to maintain loop continuity in the event the fine-wire coil inside the milliammeter movement were to accidentally break open.

13.7.4 Using shunt resistors to measure loop current

A similar method for non-invasively measuring current in a 4-20 mA instrumentation circuit is to install a precision resistor in series. If the resistance value is precisely known, the technician merely needs to measure voltage across it with a voltmeter and use Ohm's Law to calculate current:




In electronics, such a precision resistor used for measuring current is often referred to as a *shunt* resistor. Shunt resistor values are commonly very small, for their purpose is to assist in current measurement without imposing undue voltage drop within a circuit. It is rare to find a 250 ohm resistor used strictly as a diagnostic shunt resistor, because the extra voltage drop (1 to 5 volts, depending on the current signal level) may “starve” loop-powered instruments of voltage necessary to operate. Shunt resistor values as low as 1 ohm may be found installed in 4-20 mA current loops at strategic locations where technicians may need to measure loop current⁷.

⁷Of course, a 1 ohm resistor would drop 4 mV at 4 mA loop current, and drop 20 mV at 20 mA loop current. These small voltage values necessitate a highly accurate DC voltmeter for field measurement!

13.7.5 Troubleshooting current loops with voltage measurements

If neither component (diode nor shunt resistor) is pre-installed in the circuit, and if a Hall-effect (clamp-on) precision milliammeter is unavailable, a technician may still perform useful troubleshooting measurements using nothing but a DC voltmeter. Here, however, one must be careful of how to interpret these voltage measurements, for they may not directly correspond to the loop current as was the case with measurements taken in parallel with the precision resistor.

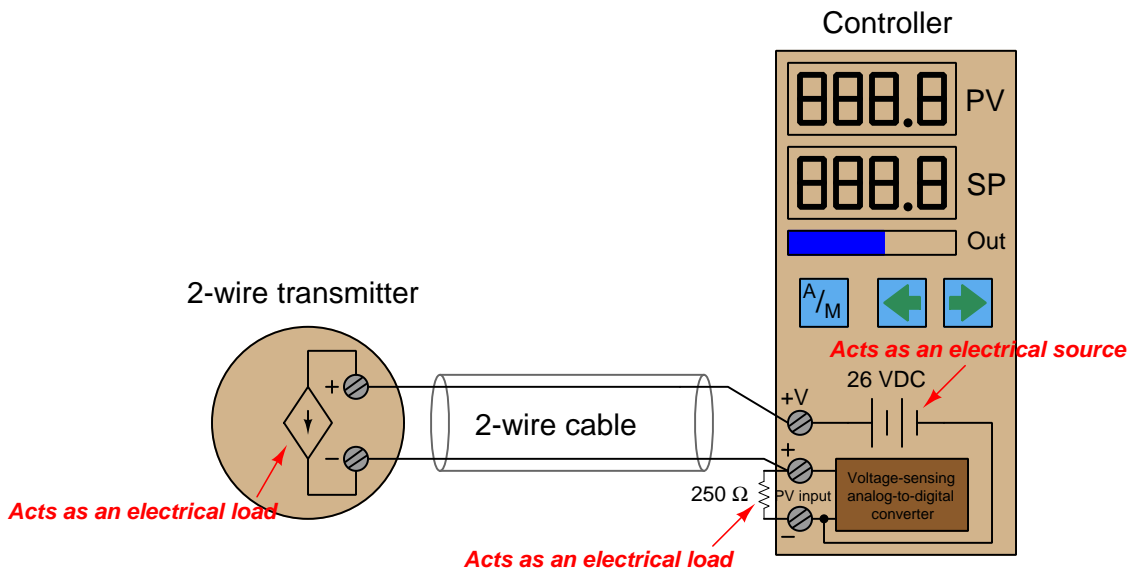
Take for example this 4-20 mA loop where a controller sends a command signal to an I/P transducer:



There is no standardized resistance value for I/P transducer coils, and so the amount of voltage dropped across the I/P terminals for any given amount of loop current will be unique for every different model of I/P. The Fisher model 567 I/P transducer built for 4-20 mA signals has a normal coil resistance of 176 ohms. Thus, we would expect to see a voltage drop of approximately 0.7 volts at 4 mA and a drop of approximately 3.5 volts at 20 mA across the I/P terminals. Since the controller output terminals are directly in parallel with the I/P terminals, we would expect to see approximately the same voltage there as well (slightly greater due to wire resistance). The lack of known precision in the I/P coil resistance makes it difficult to tell exactly how much current is in the loop for any given voltage measurement we take with a voltmeter. However, if we do know the approximate coil resistance of the I/P, we can at least obtain an estimate of loop current, which is usually good enough for diagnostic purposes.

If the I/P coil resistance is completely unknown, voltage measurements become useless for quantitative determination of loop current. Voltage measurements would be useful only for qualitatively determining loop continuity (i.e. whether there is a break in the wiring between the controller and I/P).

Another example for consideration is this loop-powered 4-20 mA transmitter and controller circuit, where the controller supplies DC power for the loop:



It is very common to find controllers with their own built-in loop power supplies, due to the popularity of loop-powered (2-wire) 4-20 mA transmitters. If we know the transmitter requires a DC voltage source somewhere in the circuit to power it up, it makes sense to include one in the controller, right?

The only voltage measurement that directly and accurately corresponds to loop current is the voltage directly across the 250 ohm precision resistor. A loop current of 4 mA will yield a voltage drop of 1 volt, 12 mA will drop 3 volts, 20 mA will drop 5 volts, etc.


A voltage measurement across the transmitter terminals will show us the *difference* in voltage between the 26 volt power supply and the voltage dropped across the 250 ohm resistor. In other words, the transmitter's terminal voltage is simply what is left over from the source voltage of 26 volts after subtracting the resistor's voltage drop. This makes the transmitter terminal voltage inversely proportional to loop current: the transmitter sees approximately 25 volts at 4 mA loop current (0% signal) and approximately 21 volts at 20 mA loop current (100% signal).

The use of the word "approximate" is very intentional here, for loop power supplies are usually non-regulated. In other words, the "26 volt" rating is approximate and subject to change! One of the advantages of the loop-powered transmitter circuit is that the source voltage is largely irrelevant, so long as it exceeds the minimum value necessary to ensure adequate power to the transmitter. If the source voltage drifts for any reason, it will have no impact on the measurement signal at all, because the transmitter is built as a *current regulator*, regulating current in the loop to whatever value represents the process measurement, regardless of slight changes in loop source voltage, wire resistance, etc. This rejection of power supply voltage changes means the loop power supply need not be regulated, and so in practice it rarely is.

This brings us to a common problem in loop-powered 4-20 mA transmitter circuits: maintaining sufficient operating voltage at the transmitter terminals. Recall that a loop-powered transmitter

relies on the voltage dropped across its terminals (combined with a current of less than 4 mA) to power its internal workings. This means the terminal voltage must not be allowed to dip below a certain minimum value, or else the transmitter will not have enough electrical power to continue its normal operation. This makes it possible to “starve” the transmitter of voltage if the loop power supply voltage is insufficient, and/or if the loop resistance is excessive.

To illustrate how this can be a problem, consider the following 4-20 mA measurement loop, where the controller supplies only 20 volts DC to power the loop, and an indicator is included in the circuit to provide operators with a field-mounted indication of the transmitter’s measurement:



The indicator contains its own 250 ohm resistor to provide a 1-5 volt signal for the meter mechanism to sense. This means the total loop resistance has now risen from 250 ohms to 500 ohms (plus any wire resistance). At full current (20 mA), this total circuit resistance will drop (at least) 10 volts, leaving 10 volts or less at the transmitter terminals to power the transmitter’s internal workings. 10 volts may not be enough for the transmitter to successfully operate, though. The Rosemount model 3051 pressure transmitter, for example, requires a minimum of 10.5 volts at the terminals to operate.

However, the transmitter *will* operate just fine at lower loop current levels. When the loop current is only 4 mA, for example, the combined voltage drop across the two 250 ohm resistors will be only 2 volts, leaving about 18 volts at the transmitter terminals: more than enough for practically any model of 4-20 mA loop-powered transmitter to successfully operate. Thus, the problem of insufficient supply voltage only manifests itself when the process measurement nears 100% of range. This could be a difficult problem to diagnose, since it appears only during certain process conditions and not others. A technician looking only for wiring faults (loose connections, corroded terminals, etc.) would never find the problem.


When a loop-powered transmitter is starved of voltage, its behavior becomes erratic. This is especially true of “smart” transmitters with built-in microprocessor circuitry. If the terminal voltage

dips below the required minimum, the microprocessor circuit shuts down. When the circuit shuts down, the current draw decreases accordingly. This causes the terminal voltage to rise again, at which point the microprocessor has enough voltage to start up. As the microprocessor “boots” back up again, it increases loop current to reflect the near-100% process measurement. This causes the terminal voltage to sag, which subsequently causes the microprocessor to shut down again. The result is a slow on/off cycling of the transmitter’s current, which makes the process controller think the process variable is surging wildly. The problem disappears, though, as soon as the process measurement decreases enough that the transmitter is allowed enough terminal voltage to operate normally.

13.7.6 Using loop calibrators

Special-purpose electronic test instruments called *loop calibrators* exist for the express purpose of 4-20 mA current loop circuit troubleshooting. These versatile instruments are generally capable of not only measuring current, but also *sourcing* current to unpowered devices in a loop, and also *simulating* loop-powered 4-20 mA transmitters.

A very popular loop calibrator unit is the Altek model 334A, a battery-powered, hand-held unit with a rotary knob for current adjustment and toggle switches for mode setting. The following illustration shows how this calibrator would be used to measure current in a functioning input signal loop⁸:




Here, the loop wiring is broken at the negative terminal of the loop-powered transmitter, and the calibrator connected in series to measure current. If this loop had a test diode installed, the calibrator could be connected in parallel with the diode to achieve the same function. Note the polarity of the calibrator's test leads in relation to the circuit being tested: the calibrator is acting

⁸In the following illustrated examples, the transmitter is assumed to be a pressure transmitter with a calibrated range of 0 to 750 inches of water column, 4-20 mA. The controller's PV (process variable) display is appropriately ranged to display 0 to 750 as well.


as a passive device (i.e. as a *load* rather than as a *source*), with the more positive loop terminal connected to the calibrator's red test lead and the more negative terminal connected to the black test lead.

The same loop calibrator may be used to *source* (or *drive*) a 4-20 mA signal into an indicating instrument to test the function of that instrument independently. Here, we see the Altek calibrator used as a current source to send a 16.00 mA signal to the PV (process variable) input of the controller, in order to check that the controller properly senses and displays the analog current signal:



No transmitter need be included in this illustration, because the calibrator takes its place. Note how the calibrator functions here as an active *source* of current rather than a passive load as it was in the last example. Not only does it supply the information (i.e. regulate the current), but it also provides the energy in the circuit. The DC power source inside the controller is not used for loop power, because the calibrator in "source" mode provides the necessary power to drive current through the 250 ohm resistor.

A very common use of a loop calibrator in “source” mode is to test a control valve for proper calibration, quick response, and to measure friction. Here, the loop calibrator takes place of the loop controller output, serving as the sole source of current to the I/P transducer:




This circuit configuration is extremely useful to any instrument technician testing the response of a control valve, because it allows the signal to be finely adjusted while in the direct presence of the valve to monitor its motion. If a control valve is suspected of having excessive friction in its moving parts, for instance, a technician may test the valve by incrementing and decrementing the loop calibrator’s source current in progressively smaller steps. Large step-changes in current should cause the valve to overcome friction and move, but small step-changes will fail to move the valve mechanism when frictional forces exceed the incremental forces produced by the changing pressure.

A photograph showing this very use of a loop calibrator in a valve rebuild shop appears here:



In this particular example, the loop calibrator connects to a device on the control valve called a *positioner*, which is a more sophisticated device than an I/P transducer. In addition to converting a 4-20 mA signal into an air pressure, the positioner also actively monitors the valve stem's position to ensure it goes to the correct position for any given 4-20 mA command signal. Here, the technician is using the loop calibrator to verify the control valve faithfully obeys the command signal through the entire 4 to 20 milliamp signal range.

An alternative method of sending a known current signal into an indicating instrument providing loop power is to set the loop calibrator such that it mimics (or *simulates*) the behavior of a loop-powered (2-wire) transmitter. In this mode, the calibrator regulates loop current at a user-determined value, but provides no motivating voltage to drive this current. Instead, it passively relies on the loop's regular voltage source to provide the necessary power:



Note the polarity of the calibrator's test leads: current entering the red lead and exiting the black lead, behaving as an electrical *load* just the same as a loop-powered transmitter. Like a 2-wire transmitter, the calibrator in simulate mode regulates the circuit current while depending on an external voltage source for energy.

A loop calibrator's *simulate transmitter* mode is especially useful for testing the transmitter cable and controller input to ensure any 4-20 mA signal sent by a transmitter will be correctly received and displayed by the controller. This sort of test is commonly performed on newly-installed control systems as part of the commissioning procedure, prior to start-up of the controlled process, in order to verify the controller's process variable input, 24 VDC power supply, and transmitter wiring are all properly functioning. Typically an instrument technician would simulate several different current values (e.g. 4 mA, 8 mA, 12 mA, 16 mA, 20 mA) with the calibrator in "simulate" mode while someone else monitors the controller's PV display and alarms to check for proper function.

A legacy loop calibrator still familiar to many instrument technicians at the time of this writing is the classic Transmation model 1040:



Other examples of vintage loop calibrator technology include the Nassau model 8060 (left) and the Biddle Versa-Cal (right):



A modern loop calibrator manufactured by Fluke is the model 705:



With this calibrator, the *measure*, *source*, and *simulate* modes are accessed by repeatedly pushing a button, with the current mode displayed on the screen:



Note the dual-numeric display, showing both loop current and percentage (assuming a 4-20 mA range).

13.7.7 NAMUR signal levels

One of the intrinsic benefits of a “live zero” analog signal standard such as 4-20 mA is that a wire break (open fault) can immediately be detected by the absence of current in the circuit. If the signal scale started at zero (e.g. 0-20 mA), there would be no way to electrically distinguish between a broken wire and a legitimate 0% signal value. In other words, the “live” LRV point of a 4-20 mA signal provides us with a way to indicate a certain type of circuit fault in addition to indicating an analog measurement during normal operation.

The *NAMUR* signal standard takes this philosophy one step further by defining specific diagnostic meaning to values of current lying outside the 4-20 mA range:

Signal level	Fault condition
Output \leq 3.6 mA	Sensing transducer failed low
3.6 mA < Output < 3.8 mA	Sensing transducer failed (detected) low
3.8 mA \leq Output < 4.0 mA	Measurement under-range
21.0 > Output \geq 20.5 mA	Measurement over-range
Output \geq 21.0 mA	Sensing transducer failed high

NAMUR-compliant transmitters are designed to limit their output signals between 3.8 mA and less than 21 mA when functioning properly. Signals lying outside this range indicate some form of failure has occurred within the transmitter or the circuit wiring.

NAMUR-compliant control systems will recognize these errant milliamp values as fault states, and may be programmed to take specific action upon receiving these signal values. Such actions include forcing controllers into manual mode, initiating automatic shutdown procedures, or taking some other form of safe action appropriate to the knowledge of a failed process transmitter.

13.8 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Linear equations:** any function represented by a straight line on a graph may be represented symbolically by the slope-intercept formula $y = mx + b$. Relevant to instrument input/output scaling.
- **Electrical sources versus loads:** electrical power sources output current (conventional flow) on their positive terminals and input current on their negative terminals (e.g. batteries and generators). Electrical loads do the opposite (e.g. resistors). Relevant to determining voltage drops and current directions in analog current loop circuits, as well as matching polarities between field instruments and controllers.
- **Voltage versus current sources:** voltage sources try to maintain constant voltage with variable current, while current sources try to maintain constant current with variable voltage. Relevant to the operation of 4-20 mA signaling circuits: loop transmitter act as current sources (or in some cases as current regulators), dropping as much or as little voltage as needed to maintain the desired amount of current in the circuit.
- **Self-balancing opamp circuits:** all self-balancing operational amplifier circuits work on the principle of negative feedback maintaining a nearly zero differential input voltage to the opamp. Making the "simplifying assumption" that the opamp's differential input voltage is exactly zero assists in circuit analysis, as does the assumption that the input terminals draw negligible current.

References

"Designing a 4-20mA Current Loop Using the MAX1459 Sensor Signal Conditioner" application note 1064, Maxim Integrated Products, 2005.

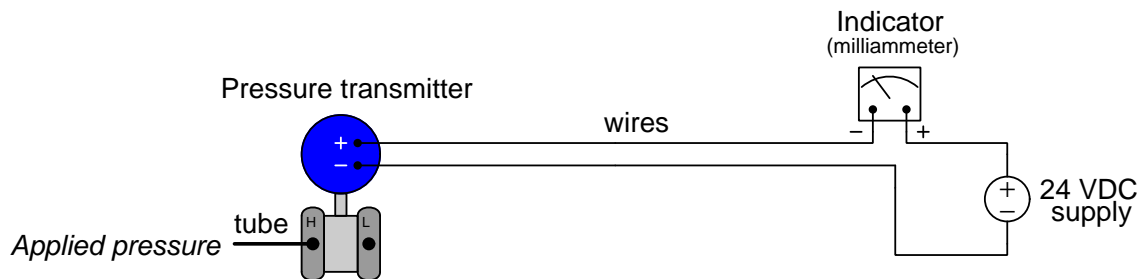
Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Software and Digital Networks*, Third Edition, CRC Press, New York, NY, 2002.

"NAMUR" whitepaper, Emerson Process Management, 2007.

Chapter 14


Pneumatic instrumentation

While electricity is commonly used as a medium for transferring energy across long distances, it is also used in instrumentation to transfer *information*. A simple 4-20 mA current “loop” uses direct current to represent a process measurement in percentage of span, such as in this example:



The transmitter senses an applied fluid pressure from the process being measured, regulates electric current in the series circuit according to its calibration (4 mA = no pressure ; 20 mA = full pressure), and the indicator (ammeter) registers this measurement on a scale calibrated to read in pressure units. If the calibrated range of the pressure transmitter is 0 to 250 PSI, then the indicator’s scale will be labeled to read from 0 to 250 PSI as well. No human operator reading that scale need worry about how the measurement gets from the process to the indicator – the 4-20 mA signal medium is transparent to the end-user as it should be.

Air pressure may be used as an alternative signaling medium to electricity. Imagine a pressure transmitter designed to output a *variable air pressure* according to its calibration rather than a *variable electric current*. Such a transmitter would have to be supplied with a source of constant-pressure compressed air instead of an electric voltage, and the resulting output signal would be conveyed to the indicator via tubing instead of wires:




The indicator in this case would be a special pressure gauge, calibrated to read in units of process pressure although actuated by the pressure of clean compressed air from the transmitter instead of directly by process fluid. The most common range of air pressure for industrial pneumatic instruments is 3 to 15 PSI. An output pressure of 3 PSI represents the low end of the process measurement scale and an output pressure of 15 PSI represents the high end of the measurement scale. Applied to the previous example of a transmitter calibrated to a range of 0 to 250 PSI, a lack of process pressure would result in the transmitter outputting a 3 PSI air signal and full process pressure would result in an air signal of 15 PSI. The face of this special “receiver” gauge would be labeled from 0 to 250 PSI, while the actual mechanism would operate on the 3 to 15 PSI range output by the transmitter. As with the 4-20 mA loop, the end-user need not know how the information gets transmitted from the process to the indicator. The 3-15 PSI signal medium is once again transparent to the operator.


Typically, a 3 PSI pressure value represents 0% of scale, a 15 PSI pressure value represents 100% of scale, and any pressure value in between 3 and 15 PSI represents a commensurate percentage in between 0% and 100%. The following table shows the corresponding current and percentage values for each 25% increment between 0% and 100%. Every instrument technician tasked with maintaining 3-15 PSI pneumatic instruments commits these values to memory, because they are referenced so often:

Pressure value	% of scale
3 PSI	0%
6 PSI	25%
9 PSI	50%
12 PSI	75%
15 PSI	100%


Pneumatic temperature, flow, and level control systems have all been manufactured to use the same principle of 3-15 PSI air pressure signaling. In each case, the transmitter and controller are both supplied clean compressed air at some modest pressure (20 to 25 PSI, usually) and the instrument signals travel via tubing. The following illustrations show what some of these applications look like:

Biodiesel "wash column" temperature control





Two-element boiler steam drum level control



Instruments operating on compressed air, and process measurement signals transmitted as air pressures through long runs of metal tubing, was the norm for industrial instrumentation prior to the advent of reliable electronic instruments. In honor of this paradigm, instrument technicians were often referred to as *instrument mechanics*, for these air-powered devices were mechanically complex and in frequent need of adjustment to maintain high accuracy.

Back in the days of control room panels populated by rows and rows of pneumatic indicators, recorders, and controllers, clean and organized routing of all the instrument signal tubes was a significant concern. By contrast, electrical wires are relatively easy to organize through the use of marshaling panels and terminal blocks – bundles of tubes (especially metal tubes!) are not. A photograph taken of the upper rear portion of an old control room panel shows a portion of a marshaling board where dozens of bulkhead-style 1/4 inch instrument tube fittings are organized in neat rows¹, where a multitude of pneumatic instrument signal lines once attached:



Each bulkhead fitting bears a numbered tag², for easy identification and documentation of tube connections. Loop diagrams of pneumatic control systems documented each bulkhead fitting where an instrument signal passed, in the same way that modern loop diagrams document each terminal block where an electrical signal connection is made.


¹Note the staggered layout of the tube fittings, intended to improve access to each one. Remember that the technician used a 9/16 inch wrench to loosen and tighten the tube fitting nuts, so it was important to have working room between fittings in which to maneuver a wrench.

²The numbers are difficult to see here, because the entire panel has been painted in a thick coat of grey paint. This particular panel was stripped of all pneumatic instruments and outfitted with electronic instruments, so the rows of bulkhead fittings no longer serve a purpose, but to remind us of legacy technology. I must wonder if some day in the future I will include a photograph of an empty terminal strip in another chapter of this book, as I explain how wired “legacy” instruments have all but been replaced by wireless (radio) instruments! Let the ghosts of the past speak to you, dear reader, testifying to the endless march of technological evolution.

Pneumatic instruments still find wide application in industry, although it is increasingly rare to encounter completely pneumatic control loops. One of the most common applications for pneumatic control system components is control valve actuation, where pneumatic technology still dominates. Not only is compressed air used to create the actuation force in many control valve mechanisms, it is still often the signal medium employed to command the valve's position. Quite often this pneumatic signal originates from a device called an *I/P transducer*, or *current-to-pressure converter*, taking a 4-20 mA control signal from the output of an electronic controller and translating that information as a pneumatic 3-15 PSI signal to the control valve's positioner or actuator.

14.1 Pneumatic sensing elements

Most pneumatic instruments use a simple but highly sensitive mechanism for converting mechanical motion into variable air pressure: the *baffle-and-nozzle* assembly (sometimes referred to as a *flapper-and-nozzle* assembly). A baffle is nothing more than a flat object obstructing the flow of air out of a small nozzle by close proximity:




The physical distance between the baffle and the nozzle alters the resistance of air flow through the nozzle. This in turn affects the air pressure built up inside the nozzle (called the nozzle *backpressure*). Like a voltage divider circuit formed by one fixed resistor and one variable resistor, the baffle/nozzle mechanism “divides” the pneumatic source pressure to a lower value based on the ratio of restrictiveness between the nozzle and the fixed orifice.


This crude assemblage is surprisingly sensitive, as shown by the graph. With a small enough orifice, just a few thousandths of an inch of motion is enough to drive the pneumatic output between its saturation limits. Pneumatic transmitters typically employ a small sheet-metal lever as the baffle. The slightest motion imparted to this baffle by changes in the process variable (pressure,

temperature, flow, level, etc.) detected by some sensing element will cause the air pressure to change in response.

The principle behind the operation of a baffle/nozzle mechanism is often used directly in quality-control work, checking for proper dimensioning of machined metal parts. Take for instance this shaft diameter checker, using air to determine whether or not a machined shaft inserted by a human operator is of the proper diameter after being manufactured on an assembly line:



If the shaft diameter is too small, there will be excessive clearance between the shaft and the inside diameter of the test jig, causing less air pressure to register on the gauge. Conversely, if the shaft diameter is too large, the clearance will be less and the gauge will register a greater air pressure because the flow of air will be obstructed by the reduced clearance. The exact pressure is of no particular consequence to the quality-control operator reading the gauge. What does matter is that the pressure falls within an acceptable range, reflecting proper manufacturing tolerances for the shaft. In fact, just like the 3-15 PSI “receiver gauges” used as pneumatic instrument indicators, the face of this pressure gauge might very well lack pressure units (such as kPa or PSI), but rather be labeled with a colored band showing acceptable limits of mechanical fit:




This is another example of the *analog* nature of pneumatic pressure signals: the pressure registered by this gauge *represents* a completely different variable, in this case the mechanical fit of the shaft to the test jig.

Although it is possible to construct a pneumatic instrument consisting *only* of a baffle/nozzle mechanism, this is rarely done. Usually the baffle/nozzle mechanism is just one of several components comprising a “balancing” mechanism in a pneumatic instrument. It is this concept of self-balancing that we will study next.

14.2 Self-balancing pneumatic instrument principles

A great many precision instruments use the principle of *balance* to measure some quantity. Perhaps the simplest example of a balance-based instrument is the common balance-beam scale used to measure mass in a laboratory:




A specimen of unknown mass is placed in one pan of the scale, and precise weights are placed in the other pan until the scale achieves a condition of balance. When balance is achieved, the mass of the specimen is known to be equal to the sum total of mass in the other pan. An interesting detail to note about the scale itself is that it has no need of routine calibration. There is nothing to “drift” out of spec which would cause the scale to read inaccurately. In fact, the scale itself doesn’t even have a gauge to register the mass of the specimen: all it has is a single mark indicating a condition of balance. To express this more precisely, the balance beam scale is actually a *differential mass* comparison device, and it only needs to be accurate at a single point: zero. In other words, it only has to be correct when it tells you there is zero difference in mass between the specimen and the standard masses piled on the other pan.

The elegance of this mechanism allows it to be quite accurate. The only real limitation to accuracy is the certainty to which we know the masses of the balancing weights.


Imagine being tasked with the challenge of automating this laboratory scale. Suppose we grew weary of employing a lab technician to place standard weights on the scale to balance it for every new measurement, and we decided to find a way for the scale to balance itself. Where would we start? First we would need some sort of mechanism to sense when the scale was out of balance, and another mechanism to apply more or less downward force on the other pan whenever an out-of-balance condition was detected.

The baffle/nozzle mechanism previously discussed would suffice quite well as a detection mechanism: simply attach a baffle to the end of the pointer on the scale, and attach a nozzle adjacent to the baffle at the “zero” position (where the pointer should come to a rest at balance). Such a mechanism might look like this:



Now we have a highly sensitive means of indicating when the scale is balanced, but we still have not yet achieved full automation. The scale cannot balance itself, at least not yet.

Suppose now instead of using precise, machined, brass weights placed on the other pan to counter the mass of the specimen, we employ a pneumatically-actuated force generator operated by the backpressure of the nozzle. An example of such a “force generator” is a *bellows*: a device made of thin sheet metal with circular corrugations in it, such that it resembles the bellows on an accordion. Pneumatic pressure applied to the interior of the bellows causes it to elongate, the amount of force applied to the bellows’ end being the product of air pressure and the end surface area:

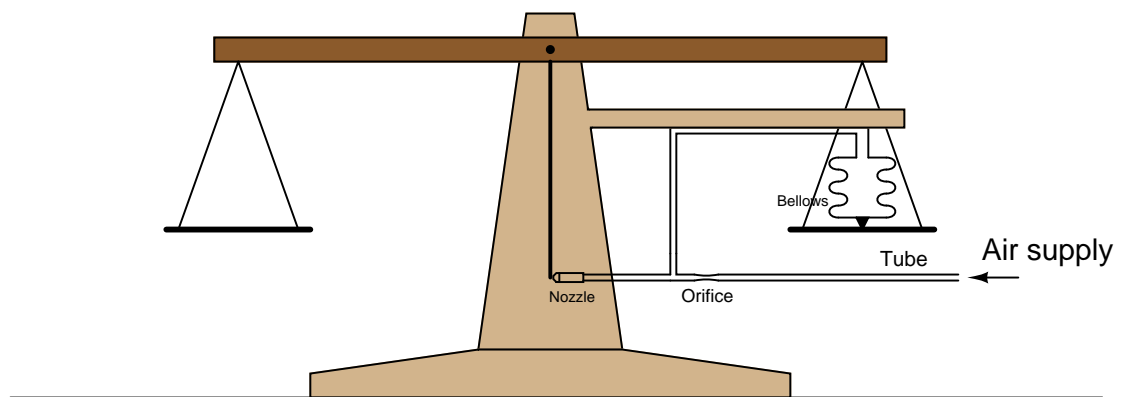


A photograph of a brass bellows unit appears here, taken from a Foxboro model 130 pneumatic controller:



If the bellows' expansion is externally restrained so it does not stretch appreciably – and therefore the metal never stretches enough to act as a restraining spring – the force exerted by the bellows on that restraining object will *exactly* equal the pneumatic pressure multiplied by the cross-sectional area of the bellows' end.

Applying this to the problem of the self-balancing laboratory scale, imagine fixing a bellows to the frame of the scale so it presses downward on the pan where the brass weights normally go, then connecting the bellows to the nozzle backpressure:




Now the scale *will* self-balance. When mass is added to the left-hand pan, the pointer (baffle) will move ever so slightly toward the nozzle until enough backpressure builds up behind the nozzle to make the bellows exert an equal and opposing force to re-balance the mechanism. This balancing action is entirely automatic: the nozzle backpressure adjusts to whatever value is necessary to

maintain the pointer in the balanced position, applying or venting pressure to the bellows as needed to keep the system in a condition of equilibrium. What we have created is a *negative feedback system*, where the output of the system (nozzle backpressure) continuously adjusts to match and balance the input (the applied weight).

This is all well and good, but how does this help us determine the weight of the specimen in the left-hand pan? What good is this self-balancing scale if we cannot *read* the balancing force? All we have achieved so far is to make the scale self-balancing. The next step is making the balancing force readable to a human operator.


Before we add the final piece to this automated scale, it is worthwhile to reflect on what has been done so far. By adding the baffle/nozzle and bellows mechanisms to the scale, we have abolished the need for brass weights and instead have substituted air pressure. In effect, the scale translates applied weight into a proportional air pressure: air pressure has now become an *analogue* for specimen weight. What we really need is a way to now translate that air pressure into a human-readable indication of weight.

To make this air pressure readable to a human operator, all we must add to the system is a pressure gauge. The gauge operates on air pressure, but now the air pressure is proportionately equivalent to specimen weight. In honor of this proportionality, we may label the face of the pressure gauge in units of ounces or grams instead of PSI or kPa:



It is very important to note how the pressure gauge performs an entirely different function with the feedback bellows in place. With just a baffle-nozzle mechanism at work, the pressure gauge was hyper-sensitive to any motion of the scale's balance beam – it served only as a highly sensitive indicator of balance. Now, with the addition of the feedback bellows, the pressure gauge actually indicates how much weight is in the specimen pan, not merely whether the scale is balanced or not. As we add more mass to the specimen pan, the gauge's indication proportionately increases. As we take away mass from the specimen pan, the gauge's indication proportionately decreases.

Although it may seem as though we are done with the task of fully automating the laboratory scale, we can go a step further. Building this pneumatic negative-feedback balancing system provides us with a capability the old manually-operated scale never had: *remote indication*. There is no reason why the indicating gauge must be located near the scale. Nothing prevents us from locating the receiver gauge some distance from the scale, and using long lengths of tubing to connect the two:



By equipping the scale with a pneumatic self-balancing apparatus, we have turned it into a *pneumatic weight transmitter*³, capable of relaying the weight measurement in analog pneumatic

³In ISA parlance, this would be a "WT" instrument, "W" signifying weight and "T" signifying transmitter.

form to an indicating gauge far away. This is the basic *force-balance* principle used in most pneumatic industrial transmitters to convert some process measurement into a 3-15 PSI pneumatic signal.

14.3 Pilot valves and pneumatic amplifying relays

Self-balancing mechanisms consisting solely of a baffle/nozzle detector coupled to a feedback bellows, while functional, are not always practical as field instruments. Nozzles and orifices must be made rather small in diameter in order to minimize compressed air usage⁴, but this means the mechanism will require significant time to alter its output pressure (i.e. to fill and empty the bellows and associated air tubing). Such a pneumatic mechanism will be impractically slow if connected to a remote indicating instrument through a long run of tubing, owing to the relatively large volume of that tube.


A practical solution to the compromise between air consumption and responsiveness inherent to simple baffle/nozzle/bellows mechanisms is to boost the nozzle backpressure (and volume) using a pneumatic “amplifier” device. With a pneumatic amplifier in place, the detector (baffle/nozzle) need not leak great quantities of compressed air, since the amplifier will provide the volume boost necessary to quickly fill (and vent) the bellows and signal tubing.

The design challenge for us, then, is how to construct such a pneumatic amplifier: a mechanism to amplify small pneumatic signal changes into larger pneumatic signal changes. In essence, we need to find a pneumatic equivalent of the electronic *transistor*: a device that lets a small signal control a larger signal.


⁴Compressed air is a valuable commodity because much energy is required to compress and distribute high-pressure air. Every pneumatic instrument’s nozzle is essentially a “leak” in the compressed air system, and the combined effect of many operating pneumatic instruments is that the air compressor(s) must continually run to meet demand.

First, let us analyze the following pneumatic mechanism and its electrical analogue (as shown on the right):

Pneumatic mechanism



Equivalent electrical circuit




As the control rod is moved up and down by an outside force, the distance between the plug and the seat changes. This changes the amount of resistance experienced by the escaping air, thus causing the pressure gauge to register varying amounts of pressure. Moving the control rod up opens the variable restriction formed by the plug and seat, venting air more easily and decreasing the output pressure. Moving the control rod down closes off the vent, causing output pressure to rise. These up-and-down rod motions are analogous to the variable resistor decreasing and increasing resistance, respectively, causing the output voltage to change in direct relation to the variable resistance.


There is little functional difference between this mechanism and a baffle/nozzle mechanism. Both work on the principle of one variable restriction and one fixed restriction (the orifice) “dividing” the pressure of the compressed air source to some lesser value.

The sensitivity of this pneumatic mechanism may be improved by extending the control rod and adding a second plug/seat assembly. The resulting mechanism, with dual plugs and seats, is known as a pneumatic *pilot valve*. An illustration of a pilot valve is shown here, along with its electrical analogue (on the right):

Pneumatic pilot valve




Equivalent electrical circuit




As the control rod is moved up and down, *both* variable restrictions change in complementary fashion. As the control rod moves up, the upper restriction closes off (restricting supply air) while the lower restriction opens up (venting more), causing the output pressure signal to decrease. As the rod moves down, the upper restriction opens (allowing more supply air in) and the lower restriction closes off (venting less), causing the output pressure to rise. The combination of *two* restrictions changing in opposite direction results in a much more aggressive change in output pressure for a given amount of rod motion than the previous mechanism with its fixed restriction and single variable restriction.

A similar design of pilot valve reverses the directions of the two plugs and seats. The only operational difference between this pilot valve and the previous design is an inverse relationship between control rod motion and pressure:

Pneumatic pilot valve



Equivalent electrical circuit




Now, moving the control rod up increases pressure while moving it down decreases pressure: precisely opposite the action of the previous pilot valve.


At this point, all we've managed to accomplish is build a better baffle/nozzle mechanism. We still do not yet have a pneumatic equivalent of an electronic transistor. To accomplish that, we must have some way of allowing an air pressure signal to control the motion of a pilot valve's control rod.

If we add a *diaphragm* to the pilot mechanism, we will create a proper *pneumatic relay*. The following relay and its electronic analogue are shown here:

Pneumatic relay



Equivalent electronic circuit




The diaphragm is nothing more than a thin disk of sheet metal, upon which an incoming air pressure signal presses. Force on the diaphragm is a simple function of signal pressure (P) and diaphragm area (A), as described by the standard force-pressure-area equation:


$$F = PA$$

If the diaphragm is taut, the elasticity of the metal allows it to also function as a spring. This allows the force to translate into displacement (motion), forming a definite relationship between applied air pressure and control rod position. Thus, the applied air pressure input will exert control over the output pressure.

It is easy to see how the input air signal exerts control over the output air signal in these two illustrations:



Since there is a direct relationship between input pressure and output pressure in this pneumatic relay, we classify it as a *direct-acting relay*. If we were to add an actuating diaphragm to the first pilot valve design, we would have a *reverse-acting relay* as shown here:




The *gain* (A) of any pneumatic relay is defined just the same as the gain of any electronic amplifier circuit, the ratio of output change to input change:

$$A = \frac{\Delta\text{Output}}{\Delta\text{Input}}$$

For example, if an input pressure change of $\Delta 2$ PSI results in an output pressure change of $\Delta 12.9$ PSI, the gain of the pneumatic relay is 6.45.

Whether or not a pneumatic relay provides a pressure gain, it is guaranteed to provide a *volume gain* which is necessary to make pneumatic field instruments practical. Note how the diaphragm chamber where the input pressure goes is sealed off: this means there will be no continual draw (or leakage) of input signal air volume. Any pneumatic sensing element sending a pressure signal to the input of a pneumatic relay will *not* be “loaded” by the relay. The relay, on the other hand, is able to supply or vent a continual flow of air at its output port as needed. Just as a transistor amplifier circuit presents a light load to the input signal and a comparatively “heavy” source/sink capacity to any load connecting to its output terminals, pneumatic relays similarly boost the volume capacity of a pneumatic signal. Recall that this was precisely our goal for increasing the responsiveness of a baffle/nozzle mechanism: to have a pneumatic amplifier capable of boosting the nozzle’s backpressure signal.

Adding a pneumatic pressure-amplifying relay to a force-balance system such as our hypothetical laboratory scale improves the performance of that pneumatic system in multiple ways:




The pressure gain of the pneumatic amplifying relay makes the force-balancing bellows respond more aggressively to changes in baffle position than it could on its own. This makes the scale more sensitive, better able to sense small changes in applied weight.

The volume gain of the pneumatic amplifying relay results in greatly decreased response time to changes in applied weight. Without a relay in the system, the rate at which the force-balance bellows fills and empties with compressed air is a direct function of the orifice's and nozzle's restrictiveness, respectively. Nozzles and orifices designed for high restriction (small diameters) work well to conserve air usage over time, but they also limit the rate of air flow in or out of the feedback bellows. With an amplifying relay in place, however, we get the best of both worlds: the nozzle and orifice bores may be minimized for minimum air consumption, while the relay's valves may be made large enough to ensure high flow capacity to and from the bellows for quick response.

It should be noted that the principles of self-balancing mechanisms, baffles and nozzles, amplifying relays, and the like are not limited to pneumatic systems. It is also possible to build self-balancing *hydraulic* systems using all the same principles, the only difference being the use of liquid (oil) as the working fluid instead of gas (air). An example of a force-balance hydraulic system is the ASCO model NH90 "Hydramotor" linear actuator, which uses a self-contained hydraulic pump and reservoir to provide pressurized oil for the mechanism, a baffle/nozzle mechanism to detect out-of-balance conditions, and a hydraulic amplifying relay to boost the nozzle backpressure signal to perform useful work through a hydraulic actuating cylinder.

The Foxboro corporation designed a great many of their pneumatic instruments using just one style of (highly sensitive) amplifying relay:




The motion of the diaphragm actuates a pair of valves: one with a cone-shaped plug and the other with a metal ball for a plug. The ball-plug admits supply air to the output port, while the cone-shaped “stem valve” plug vents excess air pressure to the vent port.

The following photograph shows one of these relay units:




Two slotted-drive screws attach the relay to the rest of the controller mechanism, while two smaller (Phillips-drive) screws hold the relay assembly together.

The Fisher corporation used a different style of amplifying relay in some of their legacy pneumatic instruments:



The following photograph shows one of these relay units (colored black) attached to the back of a model 546 I/P transducer (colored grey):




The pressure gain of this Fisher relay is much less than that of the Foxboro relay, since output pressure in the Fisher relay acts against input pressure by exerting force on a sizable diaphragm. The movable vent seat in the Fisher relay makes this design a “non-bleeding” type, meaning it possesses the ability to close both supply and vent valves at the same time, allowing it to hold an output air pressure between saturation limits without bleeding a substantial amount of compressed air to atmosphere through the vent. The Foxboro relay design, by contrast, is a “bleeding type,” whose ball and stem valves cannot close simultaneously, and thus always bleeds some compressed air to atmosphere so long as the output pressure remains somewhere between saturation limits.

14.4 Analogy to opamp circuits

Self-balancing pneumatic instrument mechanisms are very similar to negative-feedback operational amplifier circuits, in that negative feedback is used to generate an output signal in precise proportion to an input signal. This section compares simple operational amplifier (“opamp”) circuits with analogous pneumatic mechanisms for the purpose of illustrating how negative feedback works, and learning how to generally analyze pneumatic mechanisms.

In the following illustration, we see an opamp with no feedback (open loop), next to a baffle/nozzle mechanism with no feedback (open loop):



For each system there is an input and an output. For the opamp, input and output are both electrical (voltage) signals: V_{in} is the differential voltage between the two input terminals, and V_{out} is the single-ended voltage measured between the output terminal and ground. For the baffle/nozzle, the input is the physical gap between the baffle and nozzle (x_{in}) while the output is the backpressure indicated by the pressure gauge (P_{out}).

Both systems have very large gains. Operational amplifier open-loop gains typically exceed 200000 (over 100 dB), and we have already seen how just a few thousandths of an inch of baffle motion is enough to drive the backpressure of a nozzle nearly to its limits (supply pressure and atmospheric pressure, respectively).


Gain (A) is always defined as the ratio between output and input for a system. Mathematically, it is the quotient of output *change* and input *change*, with “change” represented by the triangular Greek capital-letter delta (Δ)⁵:

$$\text{Gain} = A = \frac{\Delta \text{Output}}{\Delta \text{Input}}$$


Normally, gain is a unitless ratio. We can easily see this for the opamp circuit, since both output and input are voltages, any unit of measurement for voltage would cancel in the quotient, leaving a unitless quantity. This is not so evident in the baffle/nozzle system, with the output represented in units of pressure and the input represented in units of distance.

⁵A more precise way to express gain as a ratio of changes is to use the “derivative” notation of calculus: $\frac{d\text{Output}}{d\text{Input}}$

If we were to add a bellows to the baffle/nozzle mechanism, we would have a system that inputs and outputs fluid pressure, allowing us to more formally define the gain of the system as a unitless ratio of $\frac{\Delta P_{out}}{\Delta P_{in}}$:



We may modify this mechanism slightly to make it an even more realistic analogue of an operational amplifier circuit by adding a second input bellows in direct opposition to the first:




Now our mechanism is a *differential-input pneumatic relay*. Pressure applied to the “noninverting” input bellows presses the baffle toward the nozzle, causing the output pressure to rise dramatically. Pressure applied to the “inverting” input bellows presses the baffle away from the nozzle, causing the output pressure to fall dramatically. Exactly equal pressures simultaneously applied to both bellows causes no baffle motion at all, resulting in no change in output pressure. This is analogous to an electronic operational amplifier: positive voltage applied to the noninverting (+) input strongly drives the output voltage more positive, while positive voltage applied to the inverting (–) input strongly drives the output voltage more negative.

Given the extreme sensitivity of a baffle/nozzle mechanism, the pneumatic gain of this device will be quite large. Like its electronic counterpart – the opamp circuit – miniscule input signal levels are sufficient to fully saturate the output.

Electronic operational amplifiers and pneumatic relays find their greatest application where we use *feedback*, sending all or part of the output signal back to one of the inputs of the device. The most common form of feedback is *negative feedback*, where the output signal works in such a way as to negate (compensate) itself. The general effect of negative feedback is to decrease the gain of a system, and also make that system's response more linear over the operating range. This is not an easy concept to grasp, however, and so we will explore the effect of adding negative feedback in detail for both systems.

The simplest implementation of negative feedback is a condition where the entire strength of the output signal gets “fed back” to the amplifier system in degenerative fashion. For an opamp, this simply means connecting the output terminal directly to the inverting input terminal, to form a circuit known as a *voltage follower*:



We call this “negative” or “degenerative” feedback because its effect is counteractive in nature. If the output voltage rises too high, the effect of feeding this signal to the inverting input will be to bring the output voltage back down again. Likewise, if the output voltage is too low, the inverting input will sense this and act to bring it back up again. *Self-correction* is the hallmark of any negative-feedback system.


Having connected the inverting input directly to the output of the opamp leaves us with the noninverting terminal as the sole input. Thus, our input voltage signal is a ground-referenced voltage just like the output. Electronics students learn that the voltage gain of this circuit is unity (1), meaning that the output will assume whatever voltage level is present at the input, within the limits of the opamp's power supply. This is why the circuit is called a “voltage follower”: the output voltage mimics, or follows, the input voltage. If we were to send a voltage signal of 5 volts to the noninverting terminal of this opamp circuit, it would output 5 volts, provided that the power supply exceeds 5 volts in potential from ground. What is not always taught to electronics students is *why* this is true.

Let's mathematically analyze why the gain of a "voltage follower" opamp circuit is unity. First, we will start with the equation representing the open-loop output of an opamp, as a function of its differential input voltage:

$$V_{out} = A_{OL}(V_{in(+)} - V_{in(-)})$$

As stated before, the open-loop voltage gain of an opamp is typically very large ($A_{OL} = 200000$ or more!) which means only a tiny voltage difference between the noninverting and inverting inputs ($V_{in(+)} - V_{in(-)}$) is necessary to drive the opamp's output voltage to saturation.

Connecting the opamp's output terminal to its own inverting input terminal simplifies the scenario because it makes those two terminals equipotential to each other (i.e. the output and inverting input terminals now have the same potential at all times):



With these two terminals directly connected, V_{out} and $V_{in(-)}$ are now one and the same. This means we may substitute V_{out} for $V_{in(-)}$ in the equation, while $V_{in(+)}$ simply becomes V_{in} since it is now the only remaining input. Reducing the equation to the two variables of V_{out} and V_{in} and a constant (A_{OL}) allows us to solve for overall voltage gain ($\frac{V_{out}}{V_{in}}$) as a function of the opamp's internal voltage gain (A_{OL}). The following sequence of algebraic manipulations shows how this is done:

$$V_{out} = A_{OL}(V_{in} - V_{out})$$

$$V_{out} = A_{OL}V_{in} - A_{OL}V_{out}$$

$$A_{OL}V_{out} + V_{out} = A_{OL}V_{in}$$

$$V_{out}(A_{OL} + 1) = A_{OL}V_{in}$$

$$\text{Overall gain} = \frac{V_{out}}{V_{in}} = \frac{A_{OL}}{A_{OL} + 1}$$


If we assume an internal opamp gain of 200000, the overall gain will be very nearly equal to unity (0.999995). Moreover, this near-unity gain will remain quite stable despite large changes in the opamp's internal (open-loop) gain. The following table shows the effect of major A_{OL} changes on overall voltage gain (A_V):

A_{OL} Internal gain	A_V Overall gain
100000	0.99999
200000	0.999995
300000	0.999997
500000	0.999998
1000000	0.999999

Note how an order of magnitude change⁶ in A_{OL} (from 100000 to 1000000) results in a miniscule change in overall voltage gain (from 0.99999 to 0.999999). Negative feedback clearly has a stabilizing effect on the closed-loop gain: the internal gain of the operational amplifier may drift considerably over time with negligible effect on the voltage follower circuit's overall gain. It was this principle that led Harold Black in the late 1920's to apply negative feedback to the design of very stable telephone amplifier circuits. His discovery led to the development of electronic amplifiers exhibiting very stable gains despite internal changes such as vacuum tube aging, power supply drift, etc.

⁶An "order of magnitude" is nothing more than a ten-fold change. Do you want to sound like you're really smart and impress those around you? Just start comparing ordinary differences in terms of orders of magnitude. "Hey dude, that last snowboarder's jump was an *order of magnitude* higher than the one before!" "Whoa, that's some big air . . ." Just don't make the mistake of using decibels in the same way ("Whoa dude, that last jump was at least 10 dB higher than the one before!") – you don't want people to think you're a nerd.

If we subject our voltage follower circuit to a constant input voltage of exactly 5 volts, we may expand the table to show the effect of changing open-loop gain on the output voltage, and also the differential voltage appearing between the opamp's two input terminals:




A_{OL} Internal gain	A_V Overall gain	V_{out} Output voltage	$V_{in(+)} - V_{in(-)}$ Differential voltage
100000	0.99999	4.99995	0.00005
200000	0.999995	4.999975	0.000025
300000	0.999997	4.99998	0.00002
500000	0.999998	4.99999	0.00001
1000000	0.999999	4.999995	0.000005

With such extremely high open-loop voltage gains, it hardly requires any difference in voltage between the two input terminals to generate the necessary output voltage to match the input. Negative feedback has “tamed” the opamp’s extremely large gain to a value that is nearly 1. Thus, $V_{out} = V_{in}$ for all practical purposes, and the opamp’s differential voltage input is zero for all practical purposes.

One of the “simplifying assumptions” electronics technicians and engineers make when analyzing opamp circuits is that the differential input voltage in *any* negative feedback circuit is zero. As we see in the above table, this assumption is very nearly true⁷. Following this assumption to its logical consequence allows us to predict the output voltage of any negative feedback opamp circuit quite simply. For example:


*Applying the "simplifying assumption"
of zero differential input voltage*



Building on this assumption, we may conclude that the opamp will output whatever voltage it *must* to maintain zero differential voltage between its inputs. In other words, we assume zero volts between the two input terminals of the opamp, then calculate what the output voltage must be in order for that condition to remain true. If we assume there will be zero differential voltage between the two input terminals of the opamp, we see that the output voltage will exactly equal the input voltage, for that is what *must* happen here in order for the two opamp input terminals to see equal potentials. We don't even need to evaluate a mathematical formula to tell what the voltage follower will do with a 5 volt input – the simplifying assumption lets us directly conclude $V_{out} = V_{in}$.

⁷In order for negative feedback to hold the input differential at zero volts, we must *also* assume the opamp has enough power supply voltage and output current capability to achieve this balance. No amplifier can output more voltage than its power supply gives it, nor can it output more current than its active components can conduct.

Now let us apply negative feedback to our differential-input pneumatic relay and analyze it similarly. The simplest and most direct form of negative feedback is to connect the output pressure line to the inverting bellows. This leaves only one input pressure port remaining:



It should be clear that the inverting bellows, which now experiences the same pressure (P_{out}) as the pressure gauge, introduces negative feedback into the system. If the output pressure happens to rise too high, the baffle will be pushed away from the nozzle by the force of the feedback bellows, causing backpressure to decrease and stabilize. Likewise, if the output pressure happens to go too low, the baffle will move closer to the nozzle and cause the backpressure to rise again.


As we have seen already, the baffle/nozzle is exceptionally sensitive: only a few thousandths of an inch of motion being sufficient to saturate the nozzle backpressure to either extreme (supply air pressure or zero, depending on which direction the baffle moves). This is analogous to the high gain of an operational amplifier, requiring only a few microvolts of potential difference between the input terminals to saturate the amplifier's output to full "rail" voltage. This being the case, we may conclude the nozzle backpressure can assume any value needed with negligible pressure difference between the two opposed bellows. From this we may conclude the system naturally seeks a condition where the pressure inside the feedback bellows matches the pressure inside the "input" bellows. In other words, P_{out} will (very nearly) equal P_{in} with negative feedback in effect.

Introducing negative feedback to the opamp led to a condition where the differential input voltage was held to (nearly) zero. In fact, this potential was so small as to safely consider it a constant zero microvolts for the purpose of more easily analyzing the output response of the system. *We may make the exact same "simplifying assumption" for the pneumatic mechanism:* we will assume zero baffle/nozzle gap motion (i.e. the gap remains constant) so long as negative feedback is at work.

Building on this assumption, we may conclude the nozzle backpressure will rise or fall to *whatever value it must* in order to maintain a constant baffle/nozzle gap at all times. In other words, we may analyze the operation of a pneumatic feedback mechanism for any given input condition by assuming a constant baffle/nozzle gap, then calculate what the output pressure must be in order for that assumption to remain true.


If we simply assume the baffle/nozzle gap will be held constant through the action of negative feedback, we may conclude in this case that the output pressure is exactly equal to the input pressure, since that is what *must* happen in order for the two pressures to exactly oppose each other through two identical bellows to hold the baffle at a constant gap from the nozzle.

The analytical technique of assuming perfect balance in a negative feedback system works just as well for more complicated systems. Consider the following opamp circuit:




Here, negative feedback occurs through a voltage divider from the output terminal to the inverting input terminal, such that only one-half of the output voltage gets “fed back” degeneratively. If we follow our simplifying assumption that perfect balance (zero difference of voltage) will be achieved between the two opamp input terminals due to the balancing action of negative feedback, we conclude V_{out} must be exactly *twice* the magnitude of V_{in} . In other words, the output voltage must increase to twice the value of the input voltage in order for the divided feedback signal to exactly match the input signal. Thus, feeding back half the output voltage yields an overall voltage gain of two.

If we make the same (analogous) change to the pneumatic system, we see the same effect:



Here, the feedback bellows has half the surface area of the input bellows. This results in half the amount of force applied to the beam for the same amount of pressure. If we follow our simplifying assumption that perfect balance (zero baffle motion) will be achieved due to the balancing action of negative feedback, we conclude P_{out} must be exactly *twice* the magnitude of P_{in} . In other words, the output pressure must increase to twice the value of the input pressure in order for the divided feedback force to exactly match the input force and prevent the baffle from moving. Thus, our pneumatic mechanism has a pressure gain of two, just like the opamp circuit with divided feedback had a voltage gain of two.

We could have achieved the same effect by moving the feedback bellows to a lower position on the force beam instead of changing its surface area:




This arrangement effectively reduces the feedback force by placing the feedback bellows at a mechanical disadvantage to the input bellows. If the distance between the feedback bellows tip and the force beam pivot is exactly half the distance between the input bellows tip and the force beam pivot, the effective force ratio will be one-half. The result of this “divided” feedback force is that the output pressure must rise to *twice* the value of the input pressure, since the output pressure is at a mechanical disadvantage to the input. Once again, we see a balancing mechanism with a gain of two.


Pneumatic instruments built such that bellows’ forces directly oppose one another in the same line of action to constrain the motion are known as “force balance” systems. Instruments built such that bellows’ forces oppose one another through different lever lengths pivoting around the same fulcrum point (such as in the last system) are technically known as “*moment*⁸ balance” systems. Instead of two equal forces balancing, we have two equal “moments” or torques balancing. However, one will often find that “moment balance” instruments are commonly referred to as “force balance” because the two principles are so similar. In either case, the result of the balance is that actual motion is constrained to an absolute minimum, like a tug-of-war where the two sides are perfectly matched in pulling force.

⁸In physics, the word *moment* refers to the product of force times lever length (the “moment arm”). This is alternatively known as *torque*. Thus, we could classify this pneumatic mechanism as a *torque-balance* system, since the two bellows’ forces are converted into torques (about the pivot point) which then cancel even though the forces themselves are unequal.

An entirely different classification of pneumatic instrument is known as *motion balance*. The same “simplifying assumption” of zero baffle/nozzle gap motion holds true for the analysis of these mechanisms as well:



In this particular mechanism there is no fixed pivot for the beam. Instead, the beam hangs between the ends of two bellows units, affixed by pivoting links. As input pressure increases, the input bellows expands outward, attempting to push the beam closer to the nozzle. However, if we follow our assumption that negative feedback holds the nozzle gap constant, we see that the feedback bellows must expand the same amount, and thus (if it has the same area and spring characteristics⁹ as the input bellows) the output pressure must equal the input pressure:



We call this a *motion* balance system instead of a *force* balance system because we see two *motions* working in complementary fashion to maintain a constant baffle/nozzle gap instead of two *forces* working against each other to maintain all components in their original positions.


The distinction between motion-balance and force-balance is a source of confusion for many people, but it is an important concept to master in order to analyze the responses of each and to understand how to calibrate each type of instrument. Perhaps the easiest way to distinguish one from the other is to apply the “simplifying assumption” of a constant baffle/nozzle gap and ask the question, “Must the components *move* as the signal pressures explore their full ranges?” An analogy to help distinguish force-balance from motion-balance is that of tug-of-war contestants versus ballroom dancers. Two people competing in a tug-of-war maintain a steady gap between themselves and the line only by precisely countering each other’s force: if their forces are perfectly balanced (equal and opposite), they will not move relative to the line. By contrast, two ballroom dancers maintain a steady gap between themselves only by moving the same distance: if their motions are perfectly balanced, the gap between them will not change. Again, the same question applies: do

⁹An important feature of motion-balance mechanisms is that the bellows function as calibrated spring elements in addition to being force generators. Force-balance systems move so slightly that the spring characteristics of the bellows is irrelevant – not so with motion-balance mechanisms! In fact, some motion-balance mechanisms actually place coil springs inside of brass bellows to more precisely fix the elastic properties of the assembly.

the people actually *move* in the act of maintaining equilibrium? If so, theirs is a motion-balance system; if not, theirs is a force-balance system.

In a force-balance pneumatic mechanism, the components (ideally) do not move. In fact, if there were such a thing as an infinitely sensitive baffle/nozzle mechanism (where the baffle/nozzle gap held perfectly constant at all times), *there would be absolutely no motion at all occurring in a force-balance mechanism!* However, even with an infinitely sensitive baffle/nozzle detector, a motion-balance mechanism would still visibly move in its effort to keep the baffle-nozzle gap constant as the pneumatic pressures rose and fell over their respective ranges. This is the most reliable “test” I have found to distinguish one type of mechanism from the other¹⁰.


The gain of a motion-balance pneumatic instrument may be changed by altering the bellows-to-nozzle distance such that one of the two bellows has more effect than the other. For instance, this system has a gain of 2, since the feedback bellows must move twice as far as the input bellows in order to maintain a constant nozzle gap:



Force-balance (and moment-balance) instruments are generally considered more accurate than motion-balance instruments because motion-balance instruments rely on the pressure elements (bellows, diaphragms, or bourdon tubes) possessing predictable spring characteristics. Since pressure must accurately translate to motion in a motion-balance system, there must be a predictable relationship between pressure and motion in order for the instrument to maintain accuracy. If anything happens to affect this pressure/motion relationship such as metal fatigue or temperature change, the instrument’s calibration will drift. Since there is negligible motion in a force-balance system, pressure element spring characteristics are irrelevant to the operation of these devices, and their calibrations remain more stable over time.

¹⁰In my teaching experience, students try hard to find simplistic ways to distinguish force-balance from motion-balance systems. For example, many will try to associate fulcrums with force-balance, assuming all motion-balance systems lack pivot points (which is not true!). Another example is to associate pivoting links with motion-balance mechanisms, which is likewise untrue. The problem with these efforts is that they are usually based on analysis of just a few different pneumatic mechanisms, making it easy to over-generalize. The truth of the matter is that a wide variety of pneumatic designs exist, defying easy categorization. My advice to you is the same as my advice to my students: you are going to have to *think* your way through the analysis of these mechanisms rather than *memorize* simple rules. Perform “thought experiments” whereby you imagine the effects of an increasing or decreasing input signal and then “see” for yourself whether the mechanism balances force with force or motion with motion, keeping in mind the simplifying assumption of an absolutely constant baffle/nozzle gap.

Both force- and motion-balance pneumatic instruments are usually equipped with an *amplifying relay* between the nozzle backpressure chamber and the feedback bellows. The purpose of an amplifying relay in a self-balancing pneumatic system is twofold: first, it boosts the open-loop gain of the mechanism so that its overall gain may be more predictable and stable; second, it provides additional flow capacity to fill and empty long pneumatic signal tubes necessary to convey the output air pressure signal to remote locations. The following illustration shows how a pneumatic amplifying relay may be used to improve the performance of our demonstration force-balance mechanism:




Adding a relay to a self-balancing pneumatic system is analogous to increasing the open-loop voltage gain of an opamp (A_{OL}) by several-fold: it makes the overall gain *closer to ideal*. The overall gain of the system, though, is dictated by the ratio of bellows leverage on the force beam, just like the overall gain of a negative-feedback opamp circuit is dictated by the feedback network and *not* by the opamp's internal (open-loop) voltage gain.

14.5 Analysis of practical pneumatic instruments


To better understand the design and operation of self-balancing pneumatic mechanisms, it is helpful to examine the workings of some actual instruments. In this section, we will explore three different pneumatic instruments: the Foxboro model 13A differential pressure transmitter, the Foxboro model E69 I/P (electro-pneumatic) transducer, the Fisher model 546 I/P (electro-pneumatic) transducer, and the Fisher-Rosemount model 846 I/P (electro-pneumatic) transducer.

14.5.1 Foxboro model 13A differential pressure transmitter

Perhaps one of the most popular pneumatic industrial instruments ever manufactured is the Foxboro model 13 differential pressure transmitter. A photograph of one with the cover removed is shown here:



A functional illustration of this instrument identifies its major components:



Part of the reason for this instrument's popularity is the extreme utility of differential pressure transmitters in general. A "DP cell" may be used to measure pressure, vacuum, pressure differential,

liquid level, liquid or gas flow, and even liquid density. A reason for this *particular* differential transmitter's popularity is its excellent design: the Foxboro model 13 transmitter is rugged, easy to calibrate, and quite accurate.

Like so many pneumatic instruments, the model 13 transmitter uses the *force-balance* (more precisely, the *moment-balance*) principle whereby any shift in position is sensed by a detector (the baffle/nozzle assembly) and immediately corrected through negative feedback to restore equilibrium. As a result, the output air pressure signal becomes an analogue of the differential process fluid pressure sensed by the diaphragm capsule. In the following photograph you can see my index finger pointing to the baffle/nozzle mechanism at the top of the transmitter:



Let's analyze the behavior of this transmitter step-by-step as it senses an increasing pressure on the "High pressure" input port. As the pressure here increases, the large diaphragm capsule is forced to the right. The same effect would occur if the pressure on the "Low pressure" input port were to decrease. This is a *differential* pressure transmitter, meaning it responds to fluid pressure *differences* sensed between the two input ports.


This resultant motion of the capsule tugs on the thin flexure connecting it to the force bar. The force bar pivots at the fulcrum (where the small diaphragm seal is located) in a counter-clockwise rotation, tugging the flexure at the top of the force bar. This motion causes the range bar to also pivot at its fulcrum (the sharp-edged "range wheel"), moving the baffle closer to the nozzle.

As the baffle approaches the nozzle, air flow through the nozzle becomes more restricted, accumulating backpressure in the nozzle. This backpressure increase is greatly amplified in the relay, sending an increasing air pressure signal both to the output line and to the bellows at the bottom of the range bar. Increasing pneumatic pressure in the bellows causes it to push harder on the bottom of the range bar, negating the initial motion¹¹ and returning the range bar (and force


¹¹This negating action is a hallmark of force-balance systems. When the system has reached a point of equilibrium, the components will have returned to (very nearly) their original positions. With motion-balance systems, this is not the case: one component moves, and then another component moves in response to keep the baffle/nozzle detector at a near-constant gap, but the components definitely do *not* return to their original positions or orientations.

bar) to their near-original positions.

Calibration of this instrument is accomplished through two adjustments: the zero screw and the range wheel. The zero screw simply adds tension to the bottom of the range bar, pulling it in such a direction as to further oppose the bellows' force as the zero screw is turned clockwise. This action attempts to push the baffle closer to the nozzle and therefore increases air pressure to the bellows to achieve equilibrium. Turning the range wheel alters the lever ratio of the range bar, changing the ratio of capsule force to bellows force and thereby adjusting the transmitter's span. The following photograph shows the range bar and range wheel of the instrument:




As in all instruments, the zero adjustment works by *adding or subtracting* a quantity, while the span adjustment works by *multiplying or dividing* a quantity. In the Foxboro model 13 pneumatic transmitter, the quantity in question is force, since this is a force-balance mechanism. The zero screw adds or subtracts force to the mechanical system by tensioning a spring, while the range wheel multiplies or divides force in the system by changing the mechanical advantage (force ratio) of a lever.



14.5.2 Foxboro model E69 “I/P” electro-pneumatic transducer


The purpose of any “I/P” transducer is to convert an electrical signal into a corresponding pneumatic signal. In most cases, this means an input of 4-20 mA DC and an output of 3-15 PSI, but alternative ranges do exist.

An example of an I/P transducer manufactured by Foxboro is the model E69, shown here:



Two pressure gauges indicate supply and output pressure, respectively. Wires convey the 4-20 mA electrical signal into the coil unit inside the transducer.

A view with the cover removed shows the balancing mechanism used to generate a pneumatic pressure signal from the electric current input. The baffle/nozzle may be seen at the left of the mechanism, the nozzle located at the end of a bent tube, facing the flat baffle on the surface of the circular coil unit:



As electric current passes through the coil, it produces a magnetic field which reacts against a permanent magnet's field to generate a torque. This torque causes the coil to rotate counter-clockwise (as viewed in the picture), with the baffle connected to the rotating assembly. Thus, the baffle moves like the needle of an analog electric meter movement in response to current: the more current through the coil, the more the coil assembly moves (and the baffle moves with it).


The nozzle faces this baffle, so when the baffle begins to move toward the nozzle, backpressure within the nozzle rises. This rising pressure is amplified by the relay, with the output pressure applied to a bellows. As the bellows expands, it draws the nozzle away from the advancing baffle, achieving balance by matching one motion (the baffle's) with another motion (the nozzle's). In other words, the nozzle "backs away" as the baffle "advances toward:" the *motion* of one is matched by the *motion* of the other, making this a *motion-balance* instrument.

A closer view shows the baffle and nozzle in detail:




Increased current through the wire coil causes the baffle to move toward the right (as pictured) toward the nozzle. The nozzle in response backs away (also to the right) to hold the baffle/nozzle gap constant.

Interestingly, the model E69 transducer employs the same pneumatic amplifying relay used in virtually every Foxboro pneumatic instrument:




This amplifying relay makes the system more responsive than it would be otherwise, increasing sensitivity and precision. The relay also serves as an air volume amplifier, either sourcing (supplying) or sinking (venting) air to and from a control valve actuator much more rapidly than the nozzle and orifice could do alone.

As in all instruments, the zero adjustment works by *adding or subtracting* a quantity, while the span adjustment works by *multiplying or dividing* a quantity. In the Foxboro model E69 transducer, the quantity in question is motion, since this is a motion-balance mechanism. The zero adjustment adds or subtracts motion by offsetting the position of the nozzle closer to or farther away from the baffle. A close-up photograph of the zero adjustment screw shows it pressing against a tab to rotate the mounting baseplate upon which the coil unit is fixed. Rotating this baseplate add or subtracts angular displacement to/from the baffle's motion:



The span adjustment consists of changing the position of the nozzle relative to the baffle's center of rotation (axis), so that a given amount of rotation equates to a different amount of balancing motion required of the nozzle. If the nozzle is moved farther away from the baffle's axis, the same rotation (angle) will result in greater nozzle motion (more output pressure) because the nozzle "sees" greater baffle movement. If the nozzle is moved closer toward the baffle's axis, the same rotation (angle) will result in less nozzle motion (less output pressure) because the nozzle "sees" less baffle movement¹². The effect is not unlike the difference between a baseball striking the tip of a swung bat versus striking in the middle of a swung bat: the baseball struck by the tip of the bat "sees" a faster-moving bat than the baseball struck by the middle of the bat.

This span adjustment in the E69 mechanism consists of a pair of nuts locking the base of the bellows unit at a fixed distance from the baffle's axis. Changing this distance alters the effective radius of the baffle as it swings around its center, therefore altering the gain (or span) of the motion-balance system:




¹²A good problem-solving technique to apply here is *limiting cases*, where we imagine the effects of extreme changes. In this case, we may imagine what would happen if the nozzle were moved all the way to the baffle's axis, as a limiting case of moving closer to this axis. With the nozzle in this position, no amount of baffle rotation would cause the nozzle to move away, because there is no lateral motion at the axis. Only at some radius away from the axis will there be any tangential motion for the nozzle to detect and back away from, which is why the gain of the mechanism may be altered by changing the nozzle's location with respect to the baffle's axis.

14.5.3 Fisher model 546 “I/P” electro-pneumatic transducer

The Fisher model 546 I/P transducer performs the same signal-conversion function (mA into PSI) as the Foxboro model E69, but it does so quite differently. The following photograph shows the internal mechanism of the model 546 transducer with its cover removed:



This particular instrument's construction tends to obscure its function, so I will use an illustrative diagram to describe its operation:




The heart of this mechanism is a ferrous¹³ beam, located between the poles of a permanent magnet assembly, and centered within an electromagnet coil (solenoid). Current passing through the electromagnet coil imparts magnetic poles to the ends of the beam. Following the arrow head/tail convention shown in the coil windings (the dots versus X marks) representing conventional flow vectors pointing out of the page (top) and going into the page (bottom) for the coil wrapped around the beam, the right-hand rule tells us that the beam will magnetize with the right-hand side being “North” and the left-hand side being “South.” This electro-magnetic polarity interacts with the permanent-magnetic poles to torque the beam clockwise around its pivot point (fulcrum), pushing the right-hand side down toward the nozzle.

¹³“Ferrous” simply means any substance containing the element *iron*.

Any advance of the beam toward the nozzle will increase nozzle backpressure, which is then fed to the balancing bellows at the other end of the beam. That bellows provides a restoring force to the beam to return it (nearly) to its original position. The phenomenon of an input force being counter-acted by a balancing force to ensure negligible motion is the defining characteristic of a *force-balance* system. This is the same basic principle applied in the Foxboro model 13 differential pressure transmitter: an input force countered by an output force.

If you examine the diagram carefully, you will notice that this instrument's amplifying relay is not located within the force-balance feedback loop. The nozzle's backpressure is *directly* fed back to the balancing bellows with no amplification at all. A relay does exist, but its purpose is to provide a modest (approximately 2:1) pressure gain to raise the nozzle backpressure to standard levels (3-15 PSI, or 6-30 PSI).


The next photograph shows the solenoid coil, force beam, and nozzle. If you look closely, you can see the copper-colored windings of the coil buried within the mechanism. The zero-adjustment spring is located above the beam, centered with the nozzle (below the beam):



Fisher manufactured these I/P transducers with two different pneumatic ranges: 3-15 PSI and 6-30 PSI. The mechanical difference between the two models was the size of feedback bellows used in each. In order to achieve the greater pressure range (6-30 PSI), a *smaller* feedback bellows was used. This may seem backward at first, but it makes perfect sense if you mentally follow the operation of the force-balance mechanism. In order to generate a greater air pressure for a given electric current through the coil, we must place the air pressure at a mechanical *disadvantage* to force it to rise higher than it ordinarily would in achieving balance. One way to do this is to decrease the effective area of the bellows¹⁴, so that it takes a greater air pressure to generate the same amount of balancing force on the beam.

¹⁴Recall the mathematical relationship between force, pressure, and area: $F = PA$. If we desire a greater pressure (P) to generate the same force (F) as before, we must decrease the area (A) upon which that pressure acts.

A 3-15 PSI bellows (left) is contrasted against a 6-30 PSI bellows (right) in this pair of photographs:



The span adjustment for this I/P transducer functions by varying the permanent-magnetic field strength acting against the beam's electro-magnetic field. Adjustment occurs through the use of a magnetic *shunt*: a ferrous plate moved closer to or farther away from the permanent magnets, providing an alternate (shunt, or bypass) path for magnetic flux away from the force beam. Moving the shunt farther away from the magnets strengthens the magnetic field "seen" by the beam, resulting in a multiplication of force on the beam and therefore a multiplication of output pressure. Moving the shunt closer to the magnets bypasses more of the magnetic flux, weakening the magnetic field "seen" by the beam and thereby diminishing the reaction force and also the output pressure.

A view of the mechanism's other side reveals the magnetic shunt plate, complete with an instructional arrow showing the correct direction to turn the adjustment screw to increase output span:



14.5.4 Fisher-Rosemount model 846 “I/P” electro-pneumatic transducer

The Fisher-Rosemount model 846 is a more modern I/P transducer than either the Foxboro model E69 or the Fisher model 546. It employs neither the force-balance nor the motion-balance principle in its operation, which makes it unique to analyze. This I/P unit is also unique in that it features a modular design allowing very convenient replacement of internal components when in service.

This next photograph shows three model 846 I/P transducers attached to a metal panel, below a set of five Rosemount model 1151 pressure transmitters:





A closer photograph reveals the unit in more detail:



When one of the end-covers is unscrewed, the internal workings of the I/P may be removed as a single module. Both the removed module and the housing are shown in this photograph:



Shown separately, you can see where the module's current input terminals connect with matching pins in the housing. Even the zero and span adjustment potentiometers on the module circuit board are equipped with Velcro (hook and loop) pads, matching with pads attached to calibration screws on the housing. This simple yet effective mechanical coupling allows screws located on the exterior housing to adjust resistances on the module's circuit board for zero and span calibration, yet without exposing those delicate potentiometers to ambient weather conditions:




Pneumatic (air) connections are made to the housing through standard 1/4 inch female NPT pipe threads. Compressed air passes to the module (and from the module back out to the housing) through ports, sealed from each other by O-rings¹⁵ located on the module.

The primary benefit of this modular design is ease of maintenance in the field. If a module fails for any reason, it may be very quickly removed and replaced, with no disconnection and re-connection of signal wires or pneumatic tubes necessary.

¹⁵It is quite easy to dislodge these small-section, large-diameter O-rings from their respective grooves during re-assembly of the unit. Be very careful when inserting the module back into the housing!

As mentioned before, the feedback mechanism for this particular I/P transducer employs neither the force-balance nor the motion-balance principle. Rather, the negative feedback and balancing of this unit is done electronically rather than mechanically. The following diagram shows how this works:



An electronic pressure sensor continuously monitors the output pressure, with its signal being electronically compared to the input (4-20 mA) signal by the control circuit to check for equivalence. If the output does not match the input, the control circuit drives the deflector motor with more or less current as needed, to deflect the air jet more or less as it exits one nozzle and is intercepted by the other to stimulate the pneumatic amplifying relay. Thus, we see the “balancing” internal to this I/P is done electronically rather than mechanically as it was in the other I/P relays (Foxboro model E69, Fisher model 546) explored in this section.

Electronic components are less likely to drift in their calibration, and are less susceptible to the effects of mechanical vibration and mounting orientation, than mechanical balancing components.

14.6 Proper care and feeding of pneumatic instruments

Perhaps the most important rule to obey when using pneumatic instruments is to *maintain clean and dry instrument air*. Compressed air containing dirt, rust, oil, water, or other contaminants will cause operational problems for pneumatic instruments. First and foremost is the concern that tiny orifices and nozzles inside the pneumatic mechanisms will clog over time. Clogged orifices tend to result in decreased output pressure, while clogged nozzles tend to result in increased output pressure. In either case, the “first aid” repair is to pass a welding torch tip cleaner through the plugged hole to break loose the residue or debris plugging it.

Moisture in compressed air tends to corrode metal parts inside pneumatic mechanisms. This corrosion may break loose to form debris that plugs orifices and nozzles, or it may simply eat through thin diaphragms and bellows until air leaks develop. Grossly excessive moisture will cause erratic operation as “plugs” of liquid travel through thin tubes, orifices, and nozzles designed only for air passage.

A common mistake made when installing pneumatic instruments is to connect them to a general-service (“utility”) compressed air supply instead of a dedicated instrument-service compressed air system. Utility air systems are designed to supply air tools and large air-powered actuators with pneumatic power. These high-flow compressed air systems are often seeded with antifreeze and/or lubricating chemicals to prolong the operating life of the piping and air-consuming devices, but the same liquids will wreak havoc on sensitive instrumentation. Instrument air supplies should be sourced by their own dedicated air compressor(s), complete with automatic air-dryer equipment, and distributed through stainless steel, copper, or plastic tubing (never black iron or galvanized iron pipe!).

The worst example of moisture in an instrument air system I have ever witnessed is an event that happened at an oil refinery where I worked as an instrument technician. Someone on the operations staff decided they would use 100 PSI instrument air to purge a process pipe filled with acid. Unfortunately, the acid pressure in the process pipe exceeded 100 PSI, and as a result acid flushed backward into the instrument air system. Within days most of the pneumatic instruments in that section of the refinery failed due to accelerated corrosion of metal components within the instruments. The total failure of multiple instruments over such a short time could have easily resulted in a disaster, but fortunately the crisis was minimal. Once the first couple of faulty instruments were disassembled after removal, the cause of failure became evident and the technicians took action to flush the lines of acid before too many more instruments suffered the same fate.

Pneumatic instruments must be fed compressed air of the proper pressure as well. Just as electronic circuits require power supply voltages within specified limits, pneumatic instruments do not operate well if their air supply pressure is too low or too high. If the supply pressure is too low, the instrument cannot generate a full-scale output signal. If the supply pressure is too high, internal failure may result from ruptured diaphragms, seals, or bellows. Many pneumatic instruments are equipped with their own local pressure regulators directly attached to ensure each instrument receives the correct pressure despite pressure fluctuations in the supply line.

Another “killer” of pneumatic instruments is mechanical vibration. These are precision mechanical devices, so they do not generally respond well to repeated shaking. At the very least, calibration adjustments may loosen and shift, causing the instrument’s accuracy to suffer. At worst, actual failure may result from component breakage¹⁶.

¹⁶Having said this, pneumatic instruments can be remarkably rugged devices. I once worked on a field-mounted

14.7 Advantages and disadvantages of pneumatic instruments

The disadvantages of pneumatic instruments are painfully evident to anyone familiar with both pneumatic and electronic instruments. Sensitivity to vibration, changes in temperature, mounting position, and the like affect calibration accuracy to a far greater degree for pneumatic instruments than electronic instruments. Compressed air is an expensive utility – much more expensive per equivalent watt-hour than electricity – making the operational cost of pneumatic instruments far greater than electronic. The installed cost of pneumatic instruments can be quite high as well, given the need for special (stainless steel, copper, or tough plastic) tubes to carry supply air and pneumatic signals to distant locations. The volume of air tubes used to convey pneumatic signals over distances acts as a low-pass filter, naturally damping the instrument’s response and thereby reducing its ability to respond quickly to changing process conditions. Pneumatic instruments cannot be made “smart” like electronic instruments, either. With all these disadvantages, one might wonder why pneumatic instruments are still used at all in modern industry.

Part of the answer is legacy. For an industrial facility built decades ago, it makes little sense to replace instruments that still work just fine. The cost of labor to remove old tubing, install new conduit and wires, and configure new (expensive) electronic instruments often is not worth the benefits.

However, pneumatic instruments actually enjoy some definite technical advantages which secure their continued use in certain applications even in the 21st century. One decided advantage is the *intrinsic safety* of pneumatic field instruments. Instruments that do not run on electricity cannot generate electrical sparks. This is of utmost importance in “classified” industrial environments where explosive gases, liquids, dusts, and powders exist. Pneumatic instruments are also self-purging. Their continual bleeding of compressed air from vent ports in pneumatic relays and nozzles acts as a natural clean-air purge for the inside of the instrument, preventing the intrusion of dust and vapor from the outside with a slight positive pressure inside the instrument case. It is not uncommon to find a field-mounted pneumatic instrument encrusted with corrosion and filth on the outside, but factory-clean on the inside due to this continual purge of clean air. Pneumatic instruments mounted inside larger enclosures with other devices tend to protect them all by providing a positive-pressure air purge for the entire enclosure.

Some pneumatic instruments can also function in high-temperature and high-radiation environments that would damage electronic instruments. Although it is often possible to “harden” electronic field instruments to such harsh conditions, pneumatic instruments are practically immune by nature.

An interesting feature of pneumatic instruments is that they may operate on compressed gases other than air. This is an advantage in remote natural gas installations, where the natural gas itself is sometimes used as a source of pneumatic “power” for instruments. So long as there is compressed natural gas in the pipeline to measure and to control, the instruments will operate. No air compressor or electrical power source is needed in these installations. What *is* needed, however,

pneumatic controller attached to the same support structure as a severely cavitating control valve. The vibrations of the control valve transferred to the controller through the support, causing the baffle to hammer repeatedly against the nozzle until *the nozzle’s tip had been worn down to a flattened shape*. Remarkably, the only indication of this problem was the fact the controller was having some difficulty maintaining setpoint. Other than that, it seemed to operate adequately! I doubt any electronic device would have fared as well, unless completely “potted” in epoxy.

is good filtering equipment to prevent contaminants in the natural gas (dirt, debris, liquids) from causing problems within the sensitive instrument mechanisms.

14.8 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Linear equations:** any function represented by a straight line on a graph may be represented symbolically by the slope-intercept formula $y = mx + b$. Relevant to instrument input/output scaling.
- **Pascal's principle:** changes in fluid pressure are transmitted evenly throughout an enclosed fluid volume. Relevant to pneumatic signaling, where air pressure is evenly distributed throughout a signal tube so that the pressure at one end will be equal to the pressure at the other.
- **Amplification:** the control of a relatively large signal by a relatively small signal. Relevant to the role of pneumatic relays, controlling relatively large amounts of air pressure and air flow based on the command of a much smaller air pressure signal generated by a baffle/nozzle assembly.
- **Negative feedback:** when the output of a system is degeneratively fed back to the input of that same system, the result is decreased (overall) gain and greater stability. Relevant to the internal construction of pneumatic instruments, where negative feedback is used in the form of either force-balance or motion-balance to achieve a highly linear response between input and output.
- **Self-balancing pneumatic mechanisms:** all self-balancing pneumatic instruments work on the principle of negative feedback maintaining a nearly constant baffle-nozzle gap. Force-balance mechanisms maintain this constant gap by balancing force against force with negligible motion, like a tug-of-war. Motion-balance mechanisms maintain this constant gap by balancing one motion with another motion, like two dancers moving in unison.
- **Self-balancing opamp circuits:** all self-balancing operational amplifier circuits work on the principle of negative feedback maintaining a nearly zero differential input voltage to the opamp. Making the "simplifying assumption" that the opamp's differential input voltage is exactly zero assists in circuit analysis, as does the assumption that the input terminals draw negligible current.

References

“13A d/p Cell Transmitter”, instruction manual (document MI 022-310), The Foxboro Company, Foxboro, MA, 1999.

“ASCO Nuclear Catalog”, ASCO Valve, Inc.

Black, Harold S., “Wave Translation System”, US Patent 2,102,671, filed 22 April 1932, issued to Bell Labs on 21 December 1937.

“E69 Current-to-Pneumatic Signal Converter”, instruction manual (document MI 018-430), The Foxboro Company, Foxboro, MA, 1995.

Patrick, Dale R. and Patrick, Steven R., *Pneumatic Instrumentation*, Delmar Publishers, Inc., Albany, NY, 1993.

“Type 546, 546S, and 546NS Electro-Pneumatic Transducers”, instruction manual form 1783, Fisher Controls International, Marshalltown, IA, 1997.

Chapter 15

Digital data acquisition and networks

The advent of digital electronic circuitry has brought a steady stream of technological progress to industrial instrumentation. From early applications of digital computing in the 1960's to the first distributed control systems (DCS) in the 1970's to the "smart" transmitter revolution of the 1980's, digital technology has expanded on the capabilities and information-sharing ability of measuring and control instruments. It is the purpose of this chapter to give a general overview of digital technology as it applies to data acquisition (measuring and recording process data) and digital communication, highlighting some of the more common standards used in industry. Subsequent chapters will be devoted to more in-depth discussions of specific digital instrumentation standards.

One of the greatest advantages of digital technology over analog is the ability to communicate vast amounts of data over a limited number of data channels. In the world of 4-20 mA signaling (or 3-15 PSI signaling, for that matter!) each pair of electrical wires (or pneumatic tube) could communicate only *one* variable. In the world of digital networks, one pair of wires can communicate a nearly limitless number of variables, the only limit being the *speed* of that data communication¹.

This one-signal-per-channel limit of 4-20 mA analog signals represents a technological "bottleneck" restricting data transfer between instruments and control systems. While it certainly is possible to devote a dedicated wire pair to each and every variable in an instrument system, this is very expensive to do. It is particularly cumbersome for instruments generating multiple variables of measurement, such as Coriolis flowmeters which simultaneously measure process fluid mass flow rate, fluid density, and fluid temperature; or "smart" valve positioners which continuously measure the stem position, actuator air pressure(s), air supply pressure, and temperature of a control valve. The data-rich capabilities of digital field instruments demands a digital form of communication to overcome the "bottleneck" of analog 4-20 mA signals.

Rosemount's HART standard was an early attempt to provide the "best of both worlds" in industrial instrumentation. With HART digital signals superimposed on 4-20 mA analog signals, one could retain the simplicity, fast response, and reliability of analog signaling while enjoying the multi-

¹The technical term for the "speed limit" of any data communications channel is *bandwidth*, usually expressed as a frequency (in Hertz).

variable communication benefits offered by digital signaling. However, wired-HART communication is rather slow by any standard, restricting its use to maintenance (range changes, diagnostic data polling) and process control for slow processes² only.


There exist many different digital communication standards (generally called “fieldbuses”) designed to interconnect industrial instruments. An incomplete list is shown here:

- HART
- Modbus
- FOUNDATION Fieldbus
- Profibus PA
- Profibus DP
- Profibus FMS
- EtherNet/IP
- AS-I
- CANbus
- ControlNET
- DeviceNet
- LonWorks
- BACnet

The utility of digital “fieldbus” instruments becomes apparent through the host system these instruments are connected to (typically a *Distributed Control System*, or *DCS*). Fieldbus-aware host systems usually have means to provide instrument information (including diagnostics) in very easy-to-navigate formats.

²HART communications occur at a rate of 1200 bits per second, and it is this slow by design: this slow data rate avoids signal reflection problems that would occur in unterminated cables at higher speeds. For more insight into how and why this works, refer to the “transmission lines” section 5.10 beginning on page 475. An example of a “slow” process variable suitable for HART digital monitoring or control is the *temperature* of a large building or machine, where the sheer mass of the object makes temperature changes slow by nature, and therefore does not require a fast digital data channel to report that temperature.

For example, the following computer screenshot shows the field instrument devices connected to a small-scale DCS used in an educational lab. Each instrument appears as an icon, which may be explored further simply by pointing-and-clicking with the mouse³:



³The host system in this case is an Emerson DeltaV DCS, and the device manager software is Emerson AMS.

Another application of digital communication technology to industrial measurement and control is what is generally referred to as a *SCADA* (“Supervisory Control And Data Acquisition”) system. A *SCADA* system might be thought of as a distributed control system (DCS) spread over a geographically large area, such as across the span of a city or even across national borders. Typical applications of *SCADA* technology include:

- Electric power generation and distribution (power line, substation) systems
- Water and wastewater treatment and distribution (water line, pumping stations) systems
- Gas and oil exploration and distribution (pipeline) systems
- Large-scale agricultural (irrigation, harvesting) systems
- Storage tank monitoring systems

Process data in a *SCADA* system is sensed by various measurement devices (transmitters), converted to digital form by a device called an *RTU* (“Remote Terminal Unit”), and communicated to one or more *MTUs* (“Master Terminal Units”) at a central location where human operators may monitor the data and make command decisions.

If the flow of information is one-way (simplex, from measurement devices to human operators), the system is more properly referred to as a *telemetry* system rather than a *SCADA* system. “*SCADA*” implies two-way (duplex) information flow, where human operators not only monitor process data but also issue commands back to the remote terminal units to effect change.

The saying “necessity is the mother of invention” certainly holds true for the development of remote telemetry and *SCADA* systems. The need for remote monitoring and control of electric power distribution systems led to the development of “power line carrier” analog telemetry systems as far back in time as the 1940’s. These systems superimposed⁴ high-frequency (50 kHz to 150 kHz) “carrier” signals on low-frequency (50 Hz and 60 Hz) power line conductors to communicate such basic information as human voice (like a telephone network, only dedicated for power system operators), power flow (wattmeter, MVAR meter) monitoring, and protective relay (automatic trip) controls. These *SCADA* systems were among the first to enjoy the benefits of digital technology in the 1960’s. Large-scale electric power systems simply cannot be operated safely and effectively without remote data monitoring and control, and this operational necessity pushed technological development of telemetry and *SCADA* systems beyond their small-scale (industrial manufacturing) counterparts.

Whether it is a “smart” temperature transmitter, a panel-mounted process controller with Ethernet communication capability, a variable-speed electric motor drive with Modbus signaling, a large-scale DCS controlling an oil refinery, or a *SCADA* system monitoring an international power distribution system, digital measurement and communication is an essential part of modern measurement and control systems. This chapter focuses on some of the basic principles of digital data formatting and communication, referencing practical applications wherever possible.

⁴This concept is not unlike HART, where audio-tone AC signals are superimposed on DC signal cables, so that digital data may be communicated along with DC signal and power.

15.1 Digital representation of numerical data

Process measurements are often of an analog nature: the temperature of a furnace, the rate of fluid flow through a pipe, the pressure of a fluid, etc. These data are all *analog* quantities: infinitely variable, not discrete. Discrete process measurements such as the number of units passed by on a conveyor belt are relatively easy to sense and to log in digital form, but analog measurements require special circuitry and methods of numerical representation in order for control systems to manage them.

In order for any digital device to successfully interface with an analog signal, that signal must be *digitized* by means of an *analog-to-digital converter* or *ADC*. In general, ADC circuits receive analog voltage signals from sensors and output binary numbers (*count* values) proportional to those signals. A typical scale for an ADC is a count of zero (all binary bits = 0) with 0 millivolts input, and a full-scale count (all binary bits = 1) when the voltage input value is at its maximum (full-scale). The “full scale” value of an ADC circuit is often a modest DC voltage value of 5 or 10 volts. Details of ADC performance are discussed in a subsequent section of this chapter.

Many of the concerns discussed in this section are relevant to circuits converting digital values into analog signals as well. These *digital-to-analog converters*, or *DACs*, are generally used to produce the analog drive signals required of final control elements (e.g. the output of a digital PID controller driving a 4-20 mA analog signal to a control valve positioner).

15.1.1 Integer number formats

An *integer* is any whole number or its negative counterpart (e.g. 2, 7, 241, 0, -35, -1022). All modern digital systems represent integer quantities using *binary* numeration, where integer numbers consists of strings of “bits,” each bit having two possible values: 0 or 1. Unlike the *decimal* numeration system most people are familiar with where the place-weight values are powers of ten, the place-weight values of binary numeration are powers of two. The following example shows how the integer number six hundred and three is represented in both decimal and binary formats:

Six hundred and three

$$\frac{6}{10^2} \frac{0}{10^1} \frac{3}{10^0} = (6 \times 10^2) + (3 \times 10^0) \quad \text{Decimal representation}$$

$$\frac{1}{2^9} \frac{0}{2^8} \frac{0}{2^7} \frac{1}{2^6} \frac{0}{2^5} \frac{1}{2^4} \frac{1}{2^3} \frac{0}{2^2} \frac{1}{2^1} \frac{1}{2^0} = 2^9 + 2^6 + 2^4 + 2^3 + 2^1 + 2^0 \quad \text{Binary representation}$$

The largest integer value representable in any positive place-weighted format is the base raised to the power of the number of places, minus one. So, for a three-digit decimal number, the largest possible value is $10^3 - 1$ ($1000 - 1 = 999$). For a ten-bit binary number, the largest possible value is $2^{10} - 1$ ($1024 - 1 = 1023$).

The beauty of binary representation is that the individual “bits” are easy to encode in physical form. “1” and “0” values may take the form of voltage levels (“high” or “low”) in an electronic circuit, magnetization states on a magnetic disk or magnetic tape, AC signal frequencies (“high” or “low”) transmitted on a two-conductor cable, pulses of light transmitted through a fiber optic cable, or any other medium capable of representing two different states. This makes binary a natural form of numerical representation for computers.

When representing negative integers, the convention for written numerals is to prepend a “negative” symbol to the left-hand side of the number. In digital machines, this may be emulated by including an extra “sign bit” having a value of either 0 or 1 indicating whether the number is positive or negative. A more common method for negative value representation in digital machines, however, is called *two’s complement* notation.

Two’s complement notation works by assigning the most significant place-weight a negative weight value. To use a ten-bit binary example showing the number negative three hundred and ninety six:

Negative three hundred and ninety six
(*Two’s complement binary notation*)

$$\begin{array}{cccccccccc} 1 & 0 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 0 \\ \hline -2^9 & 2^8 & 2^7 & 2^6 & 2^5 & 2^4 & 2^3 & 2^2 & 2^1 & 2^0 \end{array} = -2^9 + 2^6 + 2^5 + 2^4 + 2^2$$


Using this system to represent integer values, the greatest positive number possible is when all bits are “1” with the exception of the most significant bit (MSB). In the case of this ten-bit field, this number would be 011111111, equaling $2^9 - 1$ or 511. The greatest negative number possible is when the only “1” bit is the MSB. In the case of this ten-bit integer, this number would be 100000000, equaling -2^9 or -512 . Thus, a ten-bit two’s complement integer may represent values ranging from -512 to $+511$, inclusive.

Whether binary numbers in a computer use the two’s complement system of notation or not is often the choice of the human being programming the computer. For this reason, it is important to define whether a binary integer used within a computer system to represent data is *signed* or *unsigned*. If the integer in question is “signed,” it means the binary number is capable of representing both positive and negative integer values (usually via two’s complement notation). If the integer in question is “unsigned,” it means the binary number is capable only of representing positive values.

For example, an “unsigned” 16-bit integer has a possible numerical value range of 0 to 65535. However, the same 16-bit field – if designated as “signed” using two’s complement place-weighting – has a range of -32768 to $+32767$.

Digital computers usually organize binary numbers as clusters of bits with a set number of bits. A cluster of eight bits is called a *byte*. The standard number of bits for most modern computers exceeds one byte, with 16, 32, or even 64 bits’ length being common. Whatever the standard bit-length of a computer’s binary number field may be, it is defined for that computer as a *word*. Thus, if a digital computer is built to process binary data in 32-bit clusters, one “word” of data within that computer will be 32 bits in length. A “double-word” value for that computer is one where the binary number consists of 64 bits.

Given the typically large bit-lengths of numbers in modern computers, it is conventional to express such quantities in “shorthand” form using *hexadecimal* numeration. The basic idea of hexadecimal is that every four bits (half of a byte, or one *nybble*) of a binary number are represented by a single hex character ranging in value from 0 through F (0,1,2,3,4,5,6,7,8,9,A,B,C,D,E,F). For example, the 16-bit binary number 1001001111011110 is equivalent to the four-digit hexadecimal number 93DE:



Quite obviously, the hexadecimal number 93DE is a lot easier for human beings to manage than the binary number 1001001111011110. Computer systems providing users with live “views” of numerical values often format those values using hexadecimal numeration, just to simplify the task of interpretation.

15.1.2 Fixed-point number formats

The main limitation of integer numbers is their inability to represent quantities between whole numbers such as fractions. For example, a 4-bit integer number may have one of sixteen possible values (zero through fifteen), but it can never represent *two-and-a-half*.

One solution to this limitation is to re-assign place-weight values to the bits of a binary number, giving some of those bits fractional weights. For example, here is a ten-bit binary number field with a “binary point” separating whole-weighted bits from fractional-weighted bits, representing the quantity 70.375:

Fixed-point binary notation

$$\begin{array}{cccccccccc} 1 & 0 & 0 & 0 & 1 & 1 & 0 & \bullet & 0 & 1 & 1 \\ \frac{2^6}{2^6} & \frac{2^5}{2^5} & \frac{2^4}{2^4} & \frac{2^3}{2^3} & \frac{2^2}{2^2} & \frac{2^1}{2^1} & \frac{2^0}{2^0} & & \frac{2^{-1}}{2^{-1}} & \frac{2^{-2}}{2^{-2}} & \frac{2^{-3}}{2^{-3}} \end{array} = 2^6 + 2^2 + 2^1 + 2^{-2} + 2^{-3}$$

$$= \mathbf{70.375} \text{ (decimal notation)}$$

This form of binary notation is referred to as *fixed-point*, because the “binary point” is fixed in its location. Like any integer variable, this 10-bit binary number consisting of seven whole-weighted bits and three fractional-weighted bits has a definite range of possible values: in this particular case, from 0 to 127.875.

A variation on this theme is to use an integer number whose decimal value is some power-of-ten-times larger than it needs to be, but regard it as though it had a decimal point between two of its digits. I will illustrate this concept by way of anecdote: I once worked on a digital travel-measurement system which displayed the distance a hydraulic ram’s piston traveled using an LED numerical display. The sensor was a *pulse encoder* providing ten pulses per inch of travel, and this pulse signal directly fed a digital counter circuit which drove the LED display. When the hydraulic ram moved 19.3 inches, the LED display incremented by 193 counts. In order to make this LED display easy to read for human operators, I placed a dab of bright-red fingernail polish on the LED display to make it look as though a decimal point were permanently fixed between the least-significant digit and the next-significant digit. With the fingernail polish dot in place, 193 counts would read as 19.3. A text label saying “inches” was affixed to the right of the LED display, completing the indicator.

This solution is common to many industrial devices. Variable-frequency motor drives (VFDs), for example, often accept speed command values in integer format even though it is often necessary to specify drive frequency down to fractions of a Hertz. In such cases, the VFD manufacturer will program the VFD to imply a decimal point in the received integer number. For instance, if the fixed-point format is specified as **XXX.X** it means an integer value of 2057 sent to the VFD will be interpreted as 205.7 Hz; a value of 439 will be interpreted as 43.9 Hz; etc.

While fixing a “binary point” in the bit field escapes the limitation of whole-number resolution, a more sophisticated approach known as *floating-point* representation is a far more flexible way of representing fractional quantities in binary form. This is how most digital computers represent real-world analog quantities in digital form, and is the topic of the next subsection.

15.1.3 Floating-point number formats

Integer numbers, while very commonly used in computers for representing data, are insufficient for tasks requiring the representation of fractional (non-whole) values. Most industrial data points fall into this category, whether it be a temperature measurement (e.g. 428.21 °C), a gas concentration (e.g. 5.09 ppm of carbon monoxide), a control valve position (e.g. 27.58% open), or a motor speed (e.g. 1592.1 RPM). What we need is a form of binary representation capable of expressing fractional quantities.

One way to do this is to use *floating-point* number format, where we represent binary quantities using a form resembling scientific notation. Some control systems refer to floating-point numbers as *real* numbers. Recall how we use scientific notation to represent decimal numbers:

$$1002349 = 1.002349 \times 10^6$$

$$678.331 = 6.78331 \times 10^2$$

$$0.00004532 = 4.523 \times 10^{-5}$$

The power-of-ten multiplier has the effect of “floating” the decimal point to any position we desire, so that the significant digits of the number may be expressed in “normalized” form (between 1 and 10). This is where the term “floating-point” comes from.

In binary floating-point notation, we use a designated set of bits to represent the significant bits of the number (m), and a second designated set of bits to represent the power-of-two exponent (E):

$$\pm 1.m \times 2^E$$

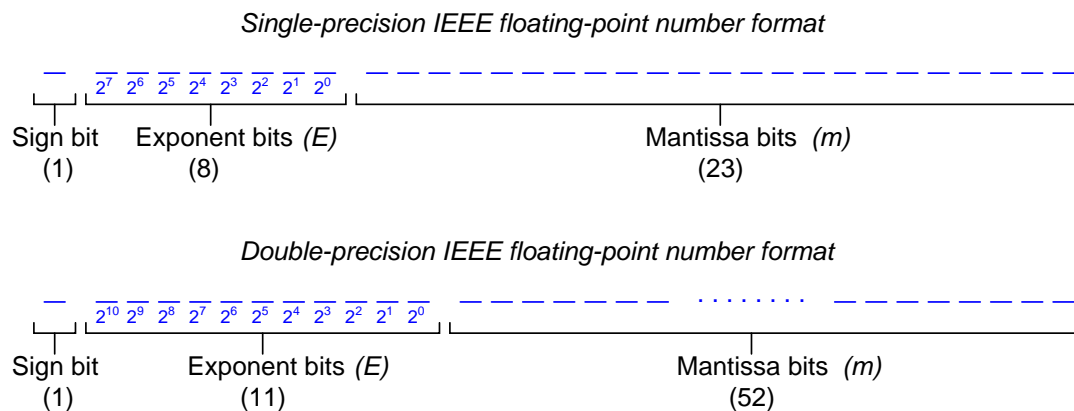
Alert readers will note that the format shown here ($\pm 1.m \times 2^E$) provides no way to represent the number zero, since 1.0×2^0 ($m = 0$ and $E = 0$) is actually equal to one! Here, floating-point notation must provide *special cases* of representation. In the IEEE standard, the format shifts to $\pm 0.m \times 2^E$ if ever all exponent bits are 0 ($E = 0$). In order to still be able to represent 1 (1.0×2^0), the IEEE standard assumes the exponent value is *biased* with a negative number, so that an exponent bit field of 0 does not mean 2^0 , but rather 2^{-bias} . This makes it possible to have an exponent value of 0 using non-zero E bits. Similarly, an exponent field consisting entirely of 1 bits is used as special representation for *infinity* or for an error code called *Not a Number* (NaN), depending on the values of the m bits. These special representations are important for handling the results of calculations such as division-by-zero and the square-roots or logarithms of negative numbers.

Given these special cases of representation required in floating-point notation, the task of doing calculations with floating-point numbers requires special processor circuitry designed to implement these rules. Inside a digital computer, this task is managed by a *floating-point processor* unit, usually a special section⁵ of the microprocessor. Some simpler microprocessors cannot support floating-point

⁵In the early days of personal computers, many microprocessor chips lacked floating-point processing capability. As a result, floating-point calculations had to be implemented in software, with programmed algorithms instructing the microprocessor how to do floating-point arithmetic. Later, floating-point processor chips were added alongside the regular microprocessor to implement these algorithms in hardware rather than emulated them through software, resulting in increased processing speed. After that, these floating-point circuits were simply added to the internal architecture of microprocessor chips as a standard feature. Even now, however, computer programmers understand that floating-point math requires more processor cycles than integer math, and should be avoided in applications

arithmetic, and thus some control system hardware (e.g. low-cost PLCs) must do all tasks using integer numbers (or fixed-point notation, if fractional quantities must be represented).

The ANSI/IEEE standard 754-1985 specifies multiple floating-point number formats, including one that is 32 bits in length (“single-precision”) and another that is 64 bits in length (“double-precision”). In the IEEE standard, one bit is reserved for the sign of the number (0 for positive and 1 for negative), a certain number of bits for the power-of-two exponent⁶, and the rest of the bits for the *mantissa* (the fractional portion of the normalized value). Both formats are shown here:



A third floating-point IEEE standard called *extended* uses 80 total bits: 1 for the sign, 15 for the exponent, and 64 for the mantissa.

Floating-point number representation greatly simplifies the task of digitally calculating real-world values. Integer numbers, by contrast, are rather clumsy when used to represent most real-world measurements or statistics. For this reason, floating-point numbers are sometimes referred to as *real* values in digital computer systems.

Specific details on IEEE floating-point representation and arithmetic rules extend beyond the scope of this book. For a good primer on the subject, I recommend Michael L. Overton’s excellent paper “Floating Point Representation”.

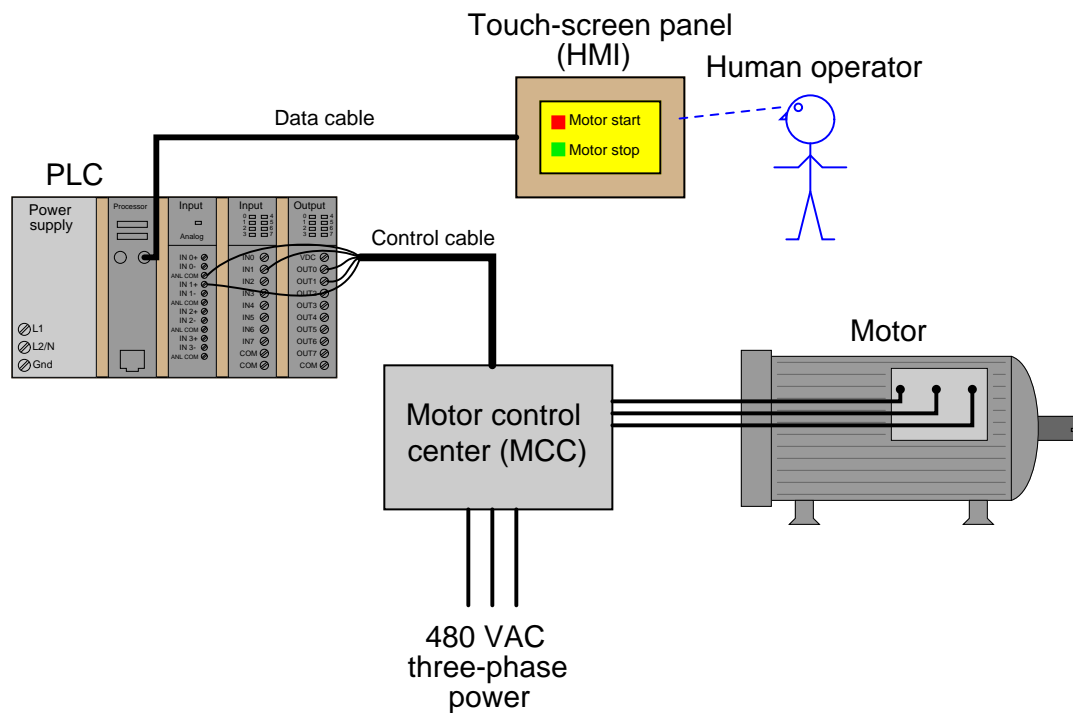
where speed is essential and floating-point representation is not. In applications demanding a small microprocessor chip and optimum speed (e.g. embedded systems), fixed-point notation is best for representing numbers containing fractional quantities.

⁶Note how the place-weights shown for the exponent field do not seem to allow for negative values. There is no negative place-weight in the most significant position as one might expect, to allow negative exponents to be represented. Instead the IEEE standard implies a *bias value* of -127 . For example, in a single-precision IEEE floating-point number, an exponent value of 11001101 represents a power of 78 (since $11001101 = 205$, the exponent’s actual value is $205 - 127 = 78$).

15.1.4 Example of industrial number formats

To illustrate how different digital data formats are used in industrial systems, we shall explore an example of a human-machine interface (HMI) panel providing graphical view and control of data within a programmable logic controller (PLC). HMI panels are nothing more than special-purpose computers equipped with graphic display screens for displaying data to human operators and pushbuttons (sometimes touch-sensitive graphic screens) for receiving input from human operators.

The following illustration shows how an HMI may be used with a PLC to control electric power to a motor:



When the operator presses a button on the HMI display to tell the motor to start, the HMI panel sends digital data to the PLC instructing the program within the PLC to activate the motor. The PLC, in turn, receives this data from the HMI and determines whether or not⁷ to start the motor. If the motor starts, the PLC sends data back to the HMI confirming the motor’s “run” status, which the HMI indicates to the operator in graphic form.

⁷This motor may be “interlocked” to prevent start-up if certain conditions are not met, thereby automatically prohibiting the operator’s instruction to start.

The configuration software used to set up communication between an HMI panel and the control device typically requires the human programmer (engineer or technician) to specify which data points inside the control device will be accessed by the HMI panel. This is specified in a list called the *tag name database*. A screen-shot of the tag name database for a particular HMI panel configured to communicate with a PLC using Modbus protocol appears here:

No.	Tag Name	Data Type	PLC Address	Device Name	Attribute	
1	START_PUSHBUTTON	Discrete	10024	PLC_01	R	
2	STOP_PUSHBUTTON	Discrete	10031	PLC_01	R	
3	MOTOR_RUN	Discrete	00005	PLC_01	R/W	
4	MOTOR_RUN_TIMER	Unsigned int 32	40010	PLC_01	R/W	
5	START_COUNTER	Unsigned int 16	40242	PLC_01	R/W	
6	MOTOR_TEMPERATURE	Floating PT 32	30008	PLC_01	R	
7	ERROR_MESSAGE	Ascii String	40560	PLC_01	R/W	
8	MOTOR_SPEED	Floating PT 32	30017	PLC_01	R	

This database table shows the “tag name” assigned to each data point (for easy programming reference), the type of digital data format used, the memory address within the control device (the PLC), the network name of the PLC, and whether the data point is read-only (R), write-only (W), or read/write (R/W) from the perspective of the HMI panel.

Note how the first three tag names are all classified as *discrete* data. The word “discrete” here is synonymous with *Boolean*, because it represents single-bit data having only two possible states: on or off, 1 or 0. As the tag names imply, these particular data points are associated with simple on/off states: a pushbutton for starting a motor, another pushbutton for stopping a motor, and a bit representing the motor’s “run” status. Once these tag names, data types, and device address values have been established, the programmer may then link the tags to graphic objects on the screen of the HMI panel. Perhaps a square box with the words “Motor start” may serve to activate the START_PUSHBUTTON bit and cause the PLC to energize the electric motor. An animated graphic image of a motor may be programmed to show rotation if the MOTOR_RUN bit is set, and show no motion if that same bit is cleared.

Four other data types appear in this tag name database, each one representing different forms of information:

- Unsigned 16-bit integer
- Unsigned 32-bit integer
- Floating-point 32-bit (also known as “Real” number format)
- ASCII string

Two “unsigned integer” data points represent motor run time and a “start counter” for keeping count of how many times the motor has been started. The timer is a 32-bit integer value, while the counter is a 16-bit integer value. Being unsigned integers (no two’s complement notation), the numerical range of the timer is 0 to 4,294,967,295 ($2^{32} - 1$) and the numerical range of the counter is 0 to 65535 ($2^{16} - 1$). The large range afforded by the timer’s 32-bit word length allows up to 136.2 years’ worth of recorded continuous motor run time, assuming a timer resolution of one second per count.

Two of these data points use 32-bit floating-point format to represent analog motor data: motor temperature and motor speed. Given the units of measurement common to these types of variables (degrees Celsius or Fahrenheit, revolutions per minute), a resolution of ± 1 as given by an integer number would probably be insufficient for the task. Instead, the operator should be able to read *fractions* of a degree in temperature, and *fractions* of an RPM motor speed. This necessity to display numerical values with resolution less than 1 unit dictates the use of the floating-point data type⁸.

The last data type represented in this tag name database is an “ASCII string,” which means a series of bytes (8-bit groupings) representing text data, such as a message. In this particular case, the tag name itself betrays the purpose of the data: `ERROR.MESSAGE`. The PLC obviously has the ability to generate text messages and send them to the HMI panel to communicate specific errors to the operator.

⁸It is also possible to “simulate” fractional resolution using an integer number, by having the HMI insert a decimal point in the numerical display. For instance, one could use a 16-bit signed integer having a numerical range of -32768 to $+32767$ to represent motor temperature by programming the HMI to insert a decimal point between the hundreds’ and the tens’ place. This would give the motor temperature tag a (displayed) numerical range of -327.68 degrees to $+327.67$ degrees, and a (displayed) resolution of ± 0.01 degree. This is basically the concept of a *fixed-point* number, where a fixed decimal point demarcates whole digits (or bits) from fractional digits (or bits).

15.2 Digital representation of text

Binary patterns are not just able to represent numerical values. Given a standardized *code*, they may represent other types of data as well, such as alphabetical characters. The ability to encode non-numerical data in digital form is what allows computers to manipulate and communicate text. In this subsection, we will explore some of the ways language characters are digitally encoded.

The representation of text by discrete (on/off) signals is rooted in ancient practice. The Greek historian Polybius described one such system in his *Histories* written in the second century BCE, used to communicate information about military maneuvers. In this passage Polybius describes the problem posed by primitive fire signals, and presents an improved method:

“It is evident to all that in every matter, and especially in warfare, the power of acting at the right time contributes very much to the success of enterprises, and fire signals are the most efficient of all the devices that aid us to do this. For they show what has recently occurred and what is still in the course of being done, and by means of them anyone who cares to do so even if he is at a distance of three, four, or even more days’ journey can be informed. So that it is always surprising how help can be brought by means of fire messages when the situation requires it. Now in former times, as fire signals were simple beacons, they were for the most part of little use to those who used them. It was possible for those who had agreed to convey a signal that a fleet had arrived in Oreus, Peparethus, or Chalcis, but when it came to some of the citizens having changed sides or having been guilty of treachery or a massacre having taken place in town, or anything of the kind, things that often happen, but cannot all be foreseen – and it is chiefly unexpected occurrences which require instant consideration and help – all such matters defied communication by fire signal. It was quite impossible to have a preconceived code for things which there was no means of foretelling.

This is the vital matter; for how can anyone consider how to render assistance if he does not know how many of the enemy have arrived, or where? And how can anyone be of good cheer or the reverse, or in fact think of anything at all, if he does not understand how many ships or how much corn has arrived from the allies?

The most recent method, devised by Cleoxenus and Democleitus and perfected by myself, is quite definite and capable of dispatching with accuracy every kind of urgent messages, but in practice it requires care and exact attention. It is as follows: We take the alphabet and divide it into five parts, each consisting of five letters. Each of the two parties who signal to each other must get ready five tablets and write one division of the alphabet on each tablet. The dispatcher of the message will raise the first set of torches on the left side indicating which tablet is to be consulted; i.e., one torch if it is the first, two if it is the second, and so on. Next he will raise the second set on the right on the same principle to indicate what letter of the tablet the receiver should write down.”

We no longer use burning torches to convey information over long distances, but we do face a similar challenge: *how might we convey messages of arbitrary length and complexity using a limited range of signals?* Like Polybius, many modern solutions are based on encodings of an alphabet, which then may be used to communicate any message reducible to written language.

15.2.1 Morse and Baudot codes

In the early days of communication, Morse code was used to represent letters of the alphabet, numerals (0 through 9), and some other characters in the form of “dot” and “dash” signals. Each “dot” or “dash” signal is one *bit* of data, equivalent to “1” and “0”. In the International Morse Code, no character requires more than five bits of data, and some (such as the common letters E and T) require only one bit.

A table showing how International Morse code represents all characters of the English alphabet and the numerals 0 through 9 appears here:

International Morse Code (English letters and Arabic numerals only)

A	•—	J	•— — —	S	•••	0	— — — — —
B	— •••	K	— • —	T	—	1	• — — — —
C	— • — •	L	• — ••	U	•• —	2	•• — — —
D	— ••	M	— —	V	••• —	3	••• — —
E	•	N	— •	W	• — —	4	•••• —
F	•• — •	O	— — — —	X	— •• —	5	•••••
G	— — •	P	• — — •	Y	— • — —	6	— ••••
H	••••	Q	— — • —	Z	— — ••	7	— — — ••
I	••	R	• — •			8	— — — — ••
						9	— — — — — •

The variable bit-length of Morse code, though very efficient⁹ in terms of the total number of “dots” and “dashes” required to communicate text messages, was difficult to automate in the form of teletype machines. In answer to this technological problem, Emile Baudot invented a different code where each and every character was five bits in length. Although this gave only 32 characters, which is not enough to represent the 26-letter English alphabet, plus all ten numerals and punctuation symbols, Baudot successfully addressed this problem by designating two of the characters as “shift” characters: one called “letters” and the other called “figures.” The other 30 characters had dual (overloaded) meanings, depending on the last “shift” character transmitted¹⁰.

⁹Morse code is an example of a *self-compressing* code, already optimized in terms of minimum bit count. Fixed-field codes such as Baudot and the more modern ASCII tend to waste bandwidth, and may be “compressed” by removing redundant bits.

¹⁰For example, the Baudot code 11101 meant either “Q” or “1” depending on whether the last shift character was “letters” or “figures,” respectively. The code 01010 meant either “R” or “4”. The code 00001 meant either “T” or a “5”. This overloading of codes is equivalent to using the “shift” key on a computer keyboard to switch between numbers and symbols (e.g. “5” versus “%”, or “8” versus “*”). The use of a “shift” key on a keyboard allows single keys on the keyboard to represent multiple characters.

15.2.2 EBCDIC and ASCII

A much more modern attempt at encoding characters useful for text representation was *EBCDIC*, the “Extended Binary Coded Decimal Interchange Code” invented by IBM in 1962 for use with their line of large (“mainframe”) computers. In EBCDIC, each character was represented by a one-byte (eight bit) code, giving this code set 256 (2^8) unique characters. Not only did this provide enough unique characters to represent all the letters of the English alphabet (lower-case *and* capital letters separately!) and numerals 0 through 9, but it also provided a rich set of *control characters* such as “null,” “delete,” “carriage return,” “linefeed,” and others useful for controlling the action of electronic printers and other peripheral devices.

A number of EBCDIC codes were unused (or seldom used), though, which made it somewhat inefficient for large data transfers. An attempt to improve this state of affairs was *ASCII*, the “American Standard Code for Information Interchange” first developed in 1963 and then later revised in 1967, both by the American National Standards Institute (ANSI). ASCII is a seven-bit code, one bit shorter per character than EBCDIC, having only 128 unique combinations as opposed to EBCDIC’s 256 unique combinations. The compromise made with ASCII versus EBCDIC was a smaller set of control characters.

IBM later created their own “extended” version of ASCII, which was eight bits per character. In this extended code set were included some non-English characters plus special graphic characters, many of which may be placed adjacently on a paper printout or on a computer console display to form larger graphic objects such as lines and boxes.

ASCII is by far the most popular digital code for representing English characters, even today. Nearly every digital transmission of English text in existence employs ASCII as the character encoding¹¹. Nearly every text-based computer program’s source code is also stored on media using ASCII encoding, where 7-bit codes represent alphanumeric characters comprising the program instructions.

¹¹Including the digital source code for this textbook!

The basic seven-bit ASCII code is shown in this table, with the three most significant bits in different columns and the four least significant bits in different rows. For example, the ASCII representation of the upper-case letter “F” is 1000110, the ASCII representation of the equal sign (=) is 0111101, and the ASCII representation of the lower-case letter “q” is 1110001.

ASCII code set

↓ LSB / MSB →	000	001	010	011	100	101	110	111
0000	NUL	DLE	SP	0	@	P	‘	p
0001	SOH	DC1	!	1	A	Q	a	q
0010	STX	DC2	”	2	B	R	b	r
0011	ETX	DC3	#	3	C	S	c	s
0100	EOT	DC4	\$	4	D	T	d	t
0101	ENQ	NAK	%	5	E	U	e	u
0110	ACK	SYN	&	6	F	V	f	v
0111	BEL	ETB	’	7	G	W	g	w
1000	BS	CAN	(8	H	X	h	x
1001	HT	EM)	9	I	Y	i	y
1010	LF	SUB	*	:	J	Z	j	z
1011	VT	ESC	+	;	K	[k	{
1100	FF	FS	,	<	L	\	l	
1101	CR	GS	-	=	M]	m	}
1110	SO	RS	.	>	N	^	n	~
1111	SI	US	/	?	O	_	o	DEL

The aforementioned “control characters” occupy the “000” and “001” columns of the table. These characters, while not associated with a printed character on the page, nevertheless play roles the other “visible” characters cannot. The “LF” character (“line feed”), for example, tells the printer to go to the next line on the paper. The “CR” character (“carriage return”) tells the printing head to return to the left-hand side of the paper. The “ESC” character (“escape”) does not result in a printed character on paper, nor does it usually command a printing machine to do anything specific, but it is often used within software as a signal to jump up one level in a menu structure, or to switch modes.

15.2.3 Unicode

There exist *many* written languages whose characters cannot and are not represented by either EBCDIC or ASCII. In an attempt to remedy this state of affairs, a new standardized code set is being developed called *Unicode*, with sixteen bits per character. This large bit field gives 65536 possible combinations, which should be enough to represent every unique character in every written language in the entire world. In deference to existing standards, Unicode encapsulates both ASCII and EBCDIC as sub-sets within its defined character set¹².

And no, I am not going to include a table showing all the Unicode characters!


15.3 Analog-digital conversion

In order to provide an interface between the internal (digital) world of a computer and the external (analog) world of process measurement and control, there must be some form of *conversion* taking place between these two types of data. Devices that generate digital representations of analog measurements are called *analog-to-digital converters*, or *ADCs*. You will find ADC circuits in process transmitters, where some process variable measurement (e.g. pressure, temperature, flow, etc.) must be converted into digital form for a computer to interpret. Devices that generate an analog representation of digital quantities are called *digital-to-analog converters*, or *DACs*. These are found in the output channels of digital control systems, where the controller must drive some field device (e.g. I/P transducer) with a 4-20 mA analog electronic signal based on a digital calculation (e.g. PID).

¹²To illustrate, the first 128 Unicode characters (0000 through 007F hexadecimal) are identical to ASCII's 128 characters (00 through 7F hexadecimal)

15.3.1 Converter resolution

Both ADCs and DACs relate analog voltage signals to multi-bit binary (digital) data. Perhaps the most obvious measure of ADC or DAC performance, then, is how many bits of output represent how many volts of analog signal. A simplified diagram of a 12-bit ADC is shown here for illustration:



Since the digital data “width” of this ADC is 12 bits, its digital output ranges from 000000000000 to 111111111111, constituting a 12-bit binary integer with a range extending from 000 hexadecimal to FFF hexadecimal, or 0 decimal to 4095 decimal. This integer number is called the *count* value of the ADC circuit. Although the ADC shown outputs its digital data in *parallel* form (with separate terminals for the 12 individual bits), many modern ADC chips are designed for *serial* data output, where a single terminal generates a sequential series of bits timed to the pulse of a clock signal.

Supposing this 12-bit ADC has an analog input voltage range of 0 to 10 volts, how do we relate any given digital number value to a voltage value, or vice-versa? The key here is to understand that the 12-bit *resolution* of this ADC means it has 2^{12} , or 4096 possible count values. The 10 volt DC input range is therefore divided up into $2^{12} - 1$, or 4095, discrete increments:

$$\text{Analog resolution} = \frac{\text{Analog span}}{2^n - 1}$$

Where,

n = Number of binary bits in the output “word”

For our hypothetical 0-10 VDC, 12-bit converter, the analog resolution is 2.442 millivolts. Thus, for any analog signal between 0 mV and 2.442 mV, the ADC’s output should be zero (binary 000000000000); for any analog signal between 2.442 mV and 4.884 mV, the ADC’s output should be one (binary 000000000001); and so on.

As previously mentioned, the digital value output by an ADC is commonly referred to as a *count*¹³. The word “count” is used in this context as a unit of measurement. For instance, if we subjected our 12-bit ADC to a full-scale input signal of 10 VDC, we would expect to see a full-scale digital output (binary 111111111111) of 4095 “counts.” Since most ADC circuits are designed to be linear, the mathematical relationship between input voltage and digital output “counts” is a simple proportionality:

$$\frac{V_{in}}{V_{fullscale}} = \frac{\text{Counts}}{2^n - 1}$$

We may use this formula to generate a partial table of input and output values for our 0-10 VDC, 12-bit ADC:

V_{in}	Counts (decimal)	Counts (hex)
0 V	0	000
2.46 mV	1	001
3.85 V	1576	628
4.59 V	1879	757
6.11 V	2502	9C6
9.998 V	4094	FFE
10 V	4095	FFF

In order to calculate a digital count value from a given input voltage, simply divide that voltage value by the full-scale voltage, then multiply by the full-scale count value and round down¹⁴ to the nearest whole number. For any given voltage value input to the ADC, there is (ideally) one corresponding output “count” value. The converse cannot be said, however: for any given output “count” value, there is actually a *range* of possible input voltages (the span of that range being the analog resolution of the ADC, in this case 2.442 mV).


To illustrate, let us take one of the table entries as an example: an analog input of 6.11 volts should yield a digital output of (precisely) 2502 counts. However, a digital output of 2502 counts could represent any analog input voltage ranging between 6.10989 volts and 6.11233 volts. This uncertainty is inherent to the process of “digitizing” an analog signal: by using a discrete quantity to represent something infinitely variable, some detail is inevitably lost. This uncertainty is referred to as *quantization error*: the (potential) error resulting from “quantizing” (digitizing) an inherently analog quantity into a discrete representation.

¹³The origin of this word has to do with the way many ADC circuits are designed, using binary *counters*. In the *tracking* design of ADC, for instance, an up-down binary counter “tracks” the varying analog input voltage signal. The binary output of this counter is fed to a DAC (digital-to-analog converter) sending an analog voltage to a comparator circuit, comparing the digital counter’s equivalent value to the value of the measured analog input. If one is greater than the other, the up-down counter is instructed to either count up or count down as necessary to equalize the two values. Thus, the up-down counter repeatedly steps up or down as needed to keep pace with the value of that analog voltage, its digital output literally “counting” steps along a fixed scale representing the full analog measurement range of the ADC circuit.

¹⁴Whether or not the actual ADC will round *down* depends on how it is designed. Some ADCs round down, others “bobble” equally between the two nearest digital values, and others yet “bobble” proportionately between the two nearest values. No matter how you round in your calculation of count value, you will never be more than 1 count off from the real ADC’s value.

Quantization error may be reduced (but never eliminated) by using an ADC with greater resolution. A 14-bit ADC operating over the same 0-10 VDC analog input range would have approximately one-quarter the uncertainty of the 12-bit ADC (0.610 mV instead of 2.442 mV). A 16-bit ADC's uncertainty would only be (approximately) one-sixteenth that of the 12-bit ADC. The number of bits chosen for any particular ADC application is therefore a function of how precise the digitization must be.

In some data acquisition systems the digital “count” scale maps to a live-zero analog range. For example, the Siemens model 353 process controller represents process variable, setpoint, and output (“valve”) percentages on a displayed scale of -3.3% to 103.3% with a 12-bit ADC count. For this controller, a digital count of zero (000 hexadecimal) represents an analog signal displayed on the controller's digital readout of -3.3% , and a full-scale digital count value (FFF hexadecimal) represents a displayed value of 103.3% . We may show the relationship between these two scales in graphical form, like a number line:



At first, the task of translating any analog value into its corresponding digital count value for this scale may seem daunting, but it is really no different than translating an input value into its corresponding output value for *any* linear instrument. This is something every student of instrumentation learns to do when working with analog-signal instrumentation, where process measurements are represented by 4-20 mA current signals or 3-15 PSI pneumatic pressure signals. If you were given a temperature transmitter with an input range of -100 to 450 degrees Celsius and an output range of 4 to 20 milliamps, and then asked to calculate the appropriate current value corresponding to 300 degrees in that range, the task would be simple. First, you could translate 300 degrees into a per-unit expression of range by subtracting the lower-range value (LRV) of -100 and then dividing by the span (URV $-$ LRV):

$$\frac{300 - (-100)}{450 - (-100)} = 0.7273 \text{ per unit}$$

After that, you would translate that per-unit figure into a current value:

$$(0.7273)(16) + 4 = 15.636 \text{ mA}$$

This general procedure of converting an input value into a per-unit representation of range, then converting that per-unit quantity into an output value, is described in detail in section 13.2.8 beginning on page 883. Other methods of performing the same translation between linear input and output ranges are explained in sections 13.2.7 and 13.2 (beginning on pages 881 and 865, respectively).


Converting a digital count value to its respective analog value (or vice-versa) follows the exact same procedure. All we need to do is identify the input and output ranges of the converter circuit, then apply the same formula. One of these ranges will be an analog voltage or current range, while the other will be a count range. If needed, the count value may be translated to or from hexadecimal or binary, because your manual calculations of course would need to be done in decimal format.

For example, suppose one of these Siemens controllers displayed a value of 26.7% on its front panel, and we wished to know what digital (count) value corresponded to this. First, we could translate the panel indication of 26.7% into a per-unit expression of range based on the scale being -3.3% to 103.3% :

$$\frac{26.7 - (-3.3)}{103.3 - (-3.3)} = \frac{30}{106.6} = 0.2814 \text{ per unit}$$

Then, we may take this per-unit expression of 0.2814 and multiply it by the span of the known count range (4095 for a 12-bit converter) to arrive at 1152, or 480 hexadecimal (480h).

Similarly, if we knew the range of this 12-bit ADC in actual process engineering units, we could translate between ADC counts and the process value by the same method. Suppose we used one of these same controllers to display the temperature of a furnace, where the lower- and upper-range values were 900 deg F and 1850 deg F, respectively. It is important to realize that these lower- and upper-range values represent 0% and 100% of scale, and that the ADC is actually configured to cover a range from -3.3% to 103.3% . In other words, the ADC range “spills over” the specified range of 0 to 100% (900 to 1850 °F) by 3.3% on both ends, equating to 31.35 °F above and below the specified range limits. Thus, the 12-bit ADC’s full-count range of 0 to 4095 (000 to FFF hex) actually covers a temperature range of 868.65 °F to 1881.35 °F:



Suppose we happened to know the ADC count inside this Siemens controller for a certain furnace temperature was A59 hexadecimal (A59h), equal to 2649 in decimal form:

To convert this count value into a temperature, first we convert the count value of 2649 into a per-unit value:

$$\frac{2649}{4095} = 0.6469 \text{ per unit}$$

Next, we translate this per-unit value into a temperature based on the known LRV and URV range points:

$$(0.6469)(1881.35 - 868.65) + 868.65 = 1523.75 \text{ degrees F}$$

15.3.2 Converter sampling rate and aliasing

The next major performance metric for analog signal digitization is how often the analog signal gets converted into digital form. Each time an ADC circuit “samples” its analog input signal, the resulting digital number is fixed until the next sample. This is analogous to monitoring a continuously moving object by taking a series of still-photographs. Any changes happening to the analog signal between sampling events are not detected by the converter, and therefore are not represented in the digital data coming from the converter.

It stands to reason, then, that the sampling rate of any ADC must be at least as often as significant changes are expected to take place in the analog measurement. According to the *Nyquist Sampling Theorem*, the absolute minimum sample rate necessary to capture an analog waveform is twice the waveform’s fundamental frequency. More realistic is to have the ADC sample the waveform *ten times* or more per cycle.

In general electronics work, for example with the design of electronic test equipment such as digital multimeters (DMMs) and digital storage oscilloscopes (DSOs), sampling rates must be rather fast. Modern digital oscilloscopes may have sampling rates in the *billions* of samples per second, to allow for the successful digitization of radio-frequency analog signals.


Industrial process measurements are far more forgiving than measurements commonly performed on an electronic technician’s workbench, thankfully. The temperature of a large furnace may be adequately sampled at a rate of only once per minute, if need be. Even “fast” feedback processes such as liquid flow and pressure control may be controlled with reasonable stability by digital systems sampling just a few times per second.

A sampling rate that is too slow (infrequent) may detrimentally affect a measurement or control system in more than one way. First, the time between samples is *dead time* to the system: time during which the digital system will be completely unresponsive to any changes in process measurement. Excessive dead time in an alarm system means an unnecessary time delay between the alarm event and the alarm signal. Excessive dead time in a feedback control loop leads to oscillation and instability. Another detrimental effect of low sampling rate is something called *aliasing*¹⁵: a condition where the digital system “thinks” the frequency of an analog signal is far lower than it actually is.

¹⁵A less-commonly-used synonym for aliasing is *folding*.


An example of signal aliasing is shown in the following illustration, where a sinusoidal signal (colored blue) is sampled at periods slightly slower than once per cycle (samples marked by red dots). The result (the red, dashed curve) is what appears to be a much lower-frequency signal as seen by the digital system, which only “sees” the values represented by the red dots¹⁶:

Sample interval slightly longer than one wave cycle



Aliasing can even occur when the sampling rate is slightly *faster* than the sinusoidal signal’s period, as shown in this illustration:


Sample interval slightly shorter than one wave cycle



As you can see from these illustrative examples, the troubling nature of aliasing is that it causes the ADC to report a *completely incorrect*, yet *completely plausible* signal. The only way to avoid aliasing is to ensure that the sampling rate is *much* faster than the signal period: a good rule here is to sample at least 10 times per period, for the highest signal frequency of interest.

¹⁶A mechanical demonstration of aliasing may be seen by using a *stroboscope* to “freeze” the motion of a rotating object. If the frequency of a flashing strobe light is set to exactly match the rotational speed of the object (e.g. 30 Hz flashing = 1800 RPM rotation), the object will appear to stand still because your eyes only see the object when it is at the exact same position every flash. This is equivalent to sampling a sinusoidal signal exactly once per cycle: the signal appears to be constant (DC) because the sine wave gets sampled at identical points along its amplitude each time. If the strobe light’s frequency is set slightly slower than the object’s rotational speed, the object will appear to slowly rotate in the forward direction because each successive flash reveals the object to be in a slightly further angle of rotation than it was before. This is equivalent to sampling a sinusoidal signal at a rate slightly slower than the signal’s frequency: the result appears to be a sinusoidal wave, but at a much slower frequency.

If we cannot set the sampling rate of the ADC to be significantly faster than the highest signal frequency we might encounter, we may avoid aliasing by preventing those high signal frequencies from ever reaching the ADC. This may be done by placing an analog low-pass filter circuit before the ADC's input. Such a "front-end" circuit is called an *anti-aliasing filter*:



Aliasing may still occur within digital systems, though, if one portion of a system "samples" the digital output of another portion at too slow of a rate. An example of this might be the rate at which a digital control system (such as a DCS) polls a process variable value collected by a digital sensor network (such as a network of radio-linked process transmitters, or digital fieldbus transmitters). If the DCS polling rate is sufficiently slow compared to the frequency of the signal reported by the digital transmitters, aliasing may result. The best guard against such potential troubles is to synchronize the sampling rates throughout the system, or (alternatively) ensure data sources always output values at a significantly slower rate than any functions reading them. Remember that a practical (minimum) sample rate to signal period ratio is 10:1 to ensure good analog-to-digital conversion integrity.

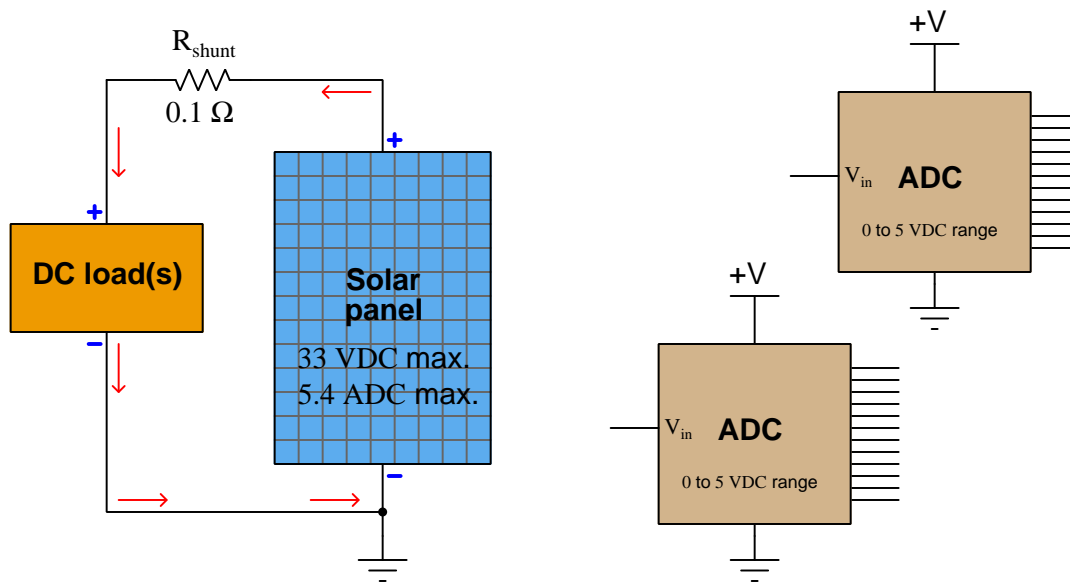
15.4 Analog signal conditioning and referencing

Modern analog-to-digital converters are phenomenally accurate, dependable, repeatable, and surprisingly inexpensive as integrated circuits considering their capabilities. However, even the best ADC is useless in a real application unless the analog voltage signal input to it is properly *conditioned*, or “made ready” for the ADC to receive. We have already explored one form of signal conditioning for ADC circuits, and that is the *anti-alias filter*, designed to block any signals from reaching the ADC with frequencies higher than the ADC can faithfully sample. An even more fundamental signal-conditioning requirement, though, is to ensure the analog input signal voltage is a good match for the voltage range of the ADC.

Most analog signals take the form of a varying *voltage*, but not all voltages are equally referenced. Recall the fundamental principle in electricity that “voltage” or “potential” is a relative quantity: *there is no such thing as a voltage existing at a single point*. Voltage is something that exists *between* two points. In electronic circuits, most voltage signals are referenced to a common point called “ground.” However, in many industrial measurement applications, the voltage signal of interest may not have one of its poles connected to ground. A voltage source may be “floating,” as in the case of an ungrounded thermocouple junction. A voltage source may also be “elevated,” which means both of its connection points are at some substantial amount of voltage with reference to ground. Whether or not a voltage signal source is “referenced” to ground poses a challenge for accurate and safe signal measurement, and it is a subject fraught with much confusion. This section of the book will explore some of these principles and applications, showing how to properly connect data acquisition (DAQ) hardware to voltage signal sources so as to overcome these problems.

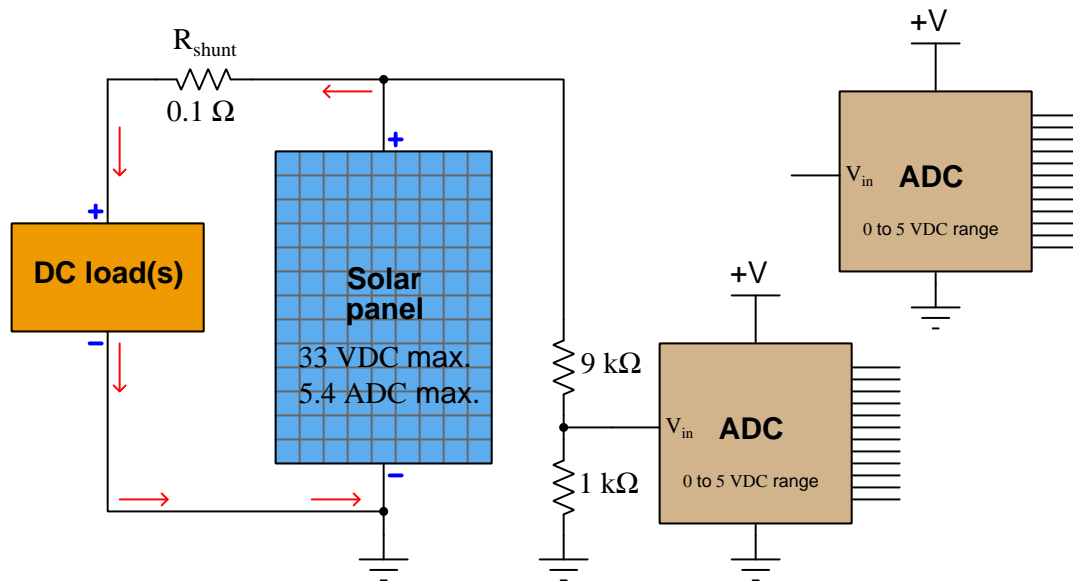
15.4.1 Instrumentation amplifiers

To illustrate the necessity of signal conditioning, and to introduce a critically important conditioning circuit called an “instrumentation amplifier,” let us examine a data acquisition application applied to a photovoltaic power system (solar electric panel) generating electricity from sunshine. Here, our goal is to use a pair of analog-to-digital converter circuits to monitor the solar panel’s voltage as well as its output current to the DC load. Since ADC circuits typically require *voltage* signals rather than current, a precision *shunt resistor* is placed in series with the solar panel to produce a measurable voltage drop directly proportional to load current:



The question we face now is, “how do we connect each ADC to the photovoltaic power system?” Each ADC is designed to digitize a DC voltage signal with reference to its ground connection, 0 to 5 volts DC producing a full range of “count” values. Right away we see that the maximum output voltage of the photovoltaic panel (33 volts DC) significantly exceeds the maximum input range of each ADC (5 volts DC), while the voltage produced by the shunt resistor for measuring load current will be very small (0.54 volts DC) compared to the input range of each ADC. Excessive signal voltage is obviously a problem, while a small voltage range will not effectively utilize the available measurement span or signal resolution.

The first problem – how to measure panel voltage when it greatly exceeds the ADC’s 5-volt maximum – may be easily solved by connecting one ADC to the panel through a precision voltage divider. In this particular case, a 10:1 divider circuit will do nicely:

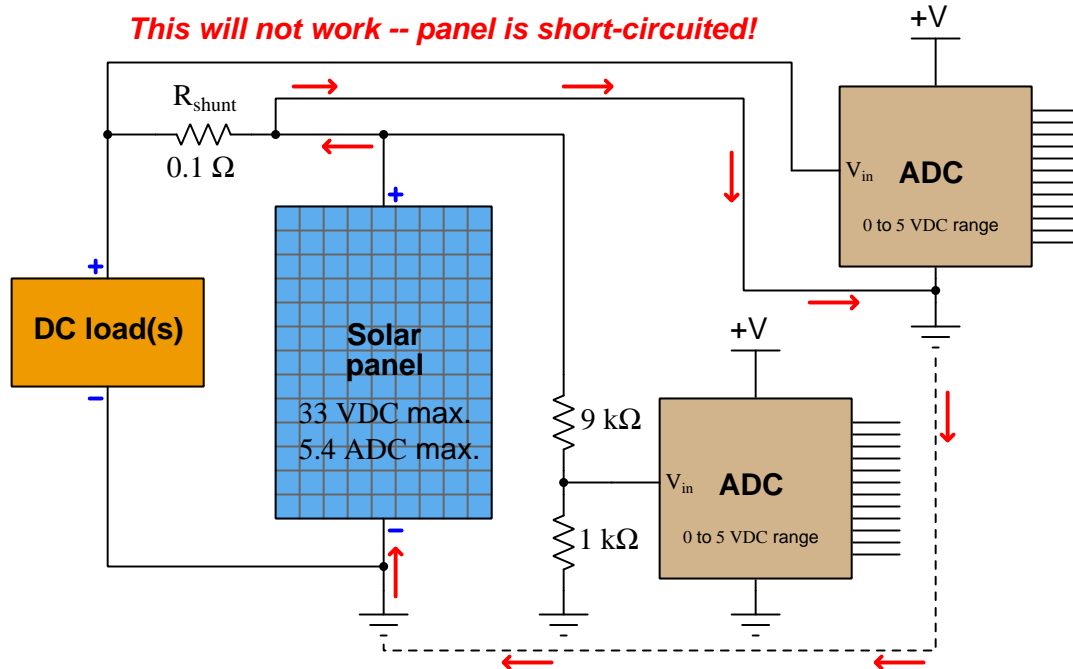


With this 10:1 voltage divider circuit in place, the panel’s 33 VDC maximum output will be seen as a 3.3 VDC maximum signal value at the ADC, which is both within its measurement range and yet spans a majority of the available range for good measurement resolution¹⁷. This simple voltage divider network thus *conditions* the solar panel’s 33 volt (maximum) output to a range acceptable to the ADC. Without such a divider in place, the ADC would be over-ranged at the very least – but most likely destroyed – by the solar panel’s relatively high voltage.

Please note how the ADC is really nothing but a voltmeter: sampling whatever voltage it senses between its V_{in} terminal and its Ground terminal. If you wish, you may visualize the ADC as being a voltmeter with red and black test leads, the red test lead being V_{in} and the black test lead being Ground.


¹⁷Remember that an ADC has a finite number of “counts” to divide its received analog signal into. A 12-bit ADC, for example, has a count range of 0 to 4095. Used to digitize an analog signal spanning the full range of 0 to 5 VDC, this means each count will be “worth” 1.22 millivolts. This is the minimum amount of signal voltage that a 12-bit, 0-5 VDC converter is able to resolve: the smallest increment of signal it is able to uniquely respond to. 1.22 mV represents 0.037% of 3.3 volts, which means this ADC may “resolve” down to the very respectable fraction 0.037% of the solar panel’s 33 volt range. If we were to use the same ADC range to directly measure the shunt resistor’s voltage drop (0 to 0.54 VDC), however, it would only be able to resolve down to 0.226% of the 0 to 5.4 amp range, which is much poorer resolution.

Connecting the second ADC to the shunt resistor presents an even greater challenge, as we see in the following schematic. Treating the second ADC as a voltmeter (its “red” test lead being the V_{in} terminal and its “black” test lead being the Ground terminal), it might seem appropriate to connect those two terminals directly across the shunt resistor. However, doing so will immediately result in the panel’s full current output flowing through the Ground conductors:



Attempting to connect the ADC in parallel with the shunt resistor in order to measure its voltage drop results in unintentionally *short-circuiting* the solar panel through the ADC’s ground connection, as shown by the “fault current” path depicted in the schematic! Not only will this circuit configuration fail to function properly, but it may even result in overheated conductors. A failure to recognize the measurement problems inherent to “elevated” voltage signals is no academic matter: a mistake like this could very well end in disaster, especially if the power source in question is much larger than a single solar panel!


One way to try eliminating the fault current path is to avoid connecting the ADC to the same signal ground point shared by the first ADC. We could power the second ADC using a battery, and simply let it “float” at an elevated potential (up to 33 volts) from ground:



While this “floating ADC” solution does avoid short-circuiting the solar panel, it does not completely eliminate the fundamental problem. When we connect the ADCs’ digital output lines to a microprocessor so as to actually do something useful with the digitized signals, we face the problem of having the first ADC’s digital lines referenced to ground, while the second ADC’s digital lines are at an elevated potential from ground (up to 33 volts!). To a microprocessor expecting 5.0 volt TTL logic signals (0 volts = “low”, 5 volts = “high”) from each ADC, this makes the second ADC’s digital output unreadable (33 volts = ???, 38 volts = ???). The microprocessor must share the same ground connection as each ADC, or else the ADCs’ digital output will not be readable.

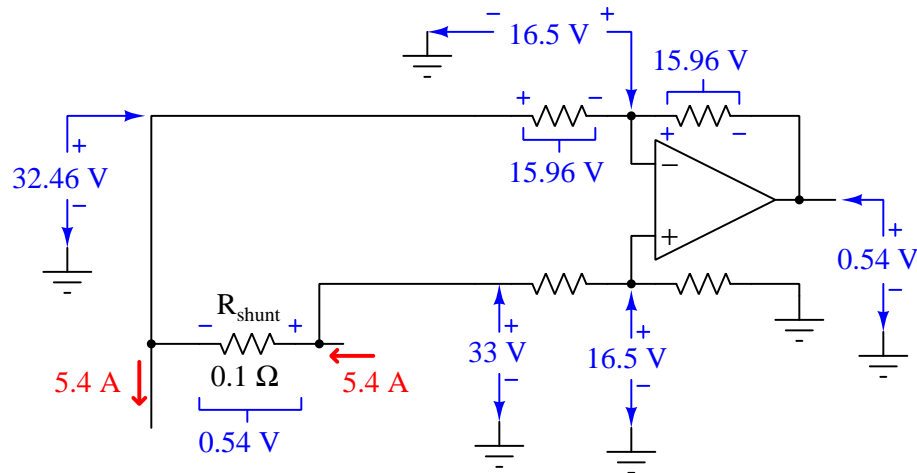
We refer to the added 33 volts as a *common-mode voltage* because that amount of voltage is common to both poles of the signal source (the shunt resistor terminals), and now is common to the digital output lines of the ADC as well. Most sensitive electronic circuits – microprocessors included – cannot effectively interpret signals having significant common-mode voltages. Somehow, we must find a way to eliminate this common-mode potential so that a microprocessor may sample both ADCs’ digital outputs.

An elegant solution to this problem involves the use of a *differential amplifier* to sample the voltage drop of the shunt resistor, then translate that voltage drop into a ground-referenced voltage ready for input to the second ADC, sharing the same ground as the first ADC. So long as this differential amplifier can tolerate the 33 VDC “common mode” voltage presented by the shunt resistor’s location on the ungrounded side of the solar panel, the shunt resistor’s signal will be properly conditioned for the ADC:



The task of the differential amplifier is to take the *difference* in potential between its two input lines and repeat that voltage at its output terminal, with reference to ground: effectively “shifting” the common-mode voltage from 33 volts to 0 volts. Thus, the differential amplifier takes a “floating” voltage signal and converts it into a ground-referenced voltage signal.


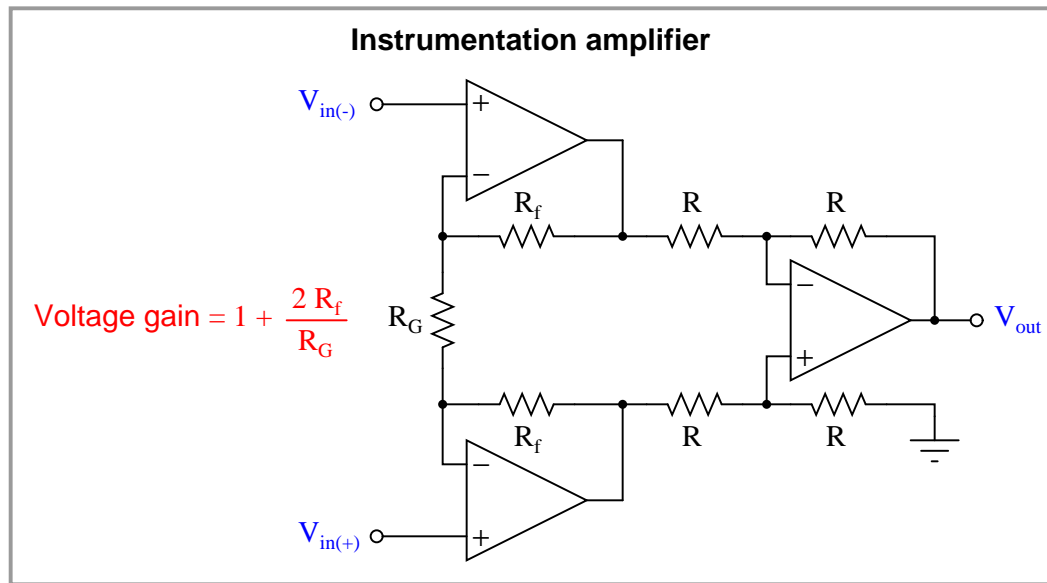
The following schematic shows how the differential amplifier does this, assuming a condition of maximum solar panel voltage and current (33 volts at 5.4 amps), and equal-value resistors in the differential amplifier circuit:



All voltages in the above schematic may be derived from the signal source (shunt resistor) and the general rule that an operational amplifier does its best to maintain zero differential voltage between its input terminals, through the action of negative feedback. The lower voltage-divider network presents half of the 33 volt solar panel potential (with reference to ground) to the noninverting opamp terminal. The opamp does its best to match this potential at its inverting input terminal (i.e. trying to keep the voltage difference between those two inputs at zero). This in turn drops 15.96 volts across the upper-left resistor (the difference between the “downstream” shunt resistor terminal voltage of 32.46 volts and the 16.5 volts matched by the opamp, both with respect to ground). That 15.96 volt drop results in a current through both upper resistors, dropping the same amount of voltage across the upper-right resistor, resulting in an opamp output voltage that is equal to 0.54 volts with respect to ground: the same potential that exists across the shunt resistor terminals, just lacking the common-mode voltage.

Not only does a differential amplifier translate an “elevated” voltage signal into a ground-referenced signal the ADC can digitize, but it also has the ability to overcome another problem we haven’t even discussed yet: amplifying the rather weak 0 to 0.54 volt shunt resistor potential into something larger to better match the ADC’s 0 to 5 volt input range. Most of the ADC’s 0 to 5 volt input range would be wasted digitizing a signal that never exceeds 0.54 volts, so amplification of this signal by some fixed gain would improve the resolution of this data channel.


Fortunately, it is a rather simple matter to equip our differential amplifier circuit with variable gain capability by adding two more operational amplifiers and three more resistors. The resulting configuration is called an *instrumentation amplifier*, so named because of its versatility in a wide variety of measurement and data acquisition applications:



A very convenient feature of the instrumentation amplifier is that its gain may be set by changing the value of a single resistor, R_G . All other resistors in an instrumentation amplifier IC are laser-trimmed components on the same semiconductor substrate as the opamps, giving them extremely high accuracy and temperature stability. R_G is typically an external resistor connected to the instrumentation amplifier IC chip by a pair of terminals.


As the formula shows us, the basic gain of an instrumentation amplifier may be adjusted from 1 (R_G open) to infinite (R_G shorted), inclusive. The input voltage range is limited only by the opamp power supplies. Thus, the instrumentation amplifier is a versatile signal-conditioning circuit for translating virtually *any* voltage signal into a ground-referenced, buffered, and amplified signal suitable for an analog-to-digital converter.

A typical DAQ (Data AcQuisition) module will have one instrumentation amplifier for every analog-to-digital converter circuit, allowing independent signal conditioning for each measurement “channel”:



The “MUX” module shown inside this data acquisition unit is a *digital multiplexer*, sequentially sampling the count values output by each ADC (one at a time) and transmitting those digital count values out to the network or “bus” cable to be read by some other digital device(s).

A final solution showing this DAQ module measuring solar panel voltage as well as current appears here:




15.4.2 Analog input references and connections

Most analog signals in industry come in the form of *voltages*. Even the ubiquitous 4 to 20 milliamp DC analog signal standard is typically converted into a 1 to 5 volt DC voltage signal before entering an electronic recorder, indicator, or controller. Therefore, the most common form of data acquisition device for analog measurement is one that accepts a DC voltage input signal. However, voltage signals cannot all be treated the same, especially with regard to a very important concept called *ground reference*. This portion of the book is devoted to an exploration of that concept, and how we may measure different kinds of voltage signals with real data acquisition devices.

First, let's examine some examples of analog voltage signal sources. For consistency, we will focus on different circuits that all sense temperature and output proportional DC voltages. Our first example is a simple voltage-divider circuit using an RTD (Resistance Temperature Detector) as the primary sensing element. An RTD is a variable resistance with a positive temperature coefficient: increasing resistance as temperature increases. Connected as shown, it will generate a signal voltage roughly¹⁸ proportional to sensed temperature:

Ground-referenced voltage signal




The power source for this circuit is commonly referred to as an “excitation” source, hence the label V_{excite} . The voltage signal measurable across the RTD is what we might refer to as a *ground-referenced* voltage signal, because one pole of it is directly connected (common) to ground.

¹⁸The relationship of temperature to V_{signal} in this sensor circuit will not be precisely linear, especially if R_{fixed} is not tremendously larger than R_{RTD} .

Our next example is a device called a *thermocouple* – a pair of dissimilar-metal wires joined together to form a junction. Thermocouple junctions produce small amounts of voltage directly proportional to temperature. As such, they are self-powered devices, needing no “excitation” power sources to operate:


Floating voltage signal



If the thermocouple junction is insulated from earth ground, we refer to it as a *floating* voltage signal. The word “floating” in this context refers to a complete lack of electrical connection to earth ground.

Returning to the use of RTDs for measuring temperature, another circuit design is the so-called *bridge* configuration, where an RTD comprises one or more “active” legs of a dual voltage divider. The excitation voltage source connects across two opposite ends of the bridge (powering both voltage dividers), while the signal voltage is measured across the other two opposite ends of the bridge (from one divider mid-point to the other):

Elevated voltage signal (signal + common-mode voltage)




The purpose of a bridge circuit is to subtract the “live zero” voltage otherwise dropped by the RTD, which cannot produce a zero-ohm resistance at any temperature. This makes it possible to have a signal voltage range beginning at 0 volts, even though the RTD’s resistance will always be non-zero. The price we pay for this elimination of signal offset is the *elevation* of the signal from ground potential.

If the fixed-value resistors on the left-hand side of this bridge circuit each have the same resistance, the “common-mode” voltage will be one-half the excitation voltage. This presents an interesting situation from the perspective of measuring V_{signal} , as the common-mode voltage may greatly exceed the signal voltage. We are not particularly interested in measuring the common-mode voltage because it tells us nothing about the sensed temperature, yet this relatively large voltage is “elevating” our signal voltage from ground potential whether we like it or not, and any data acquisition hardware we connect to the bridge circuit must deal effectively with this common-mode voltage (i.e. not let it corrupt or otherwise influence the accuracy of the desired signal measurement).


Yet another type of analog voltage signal is one where the signal is “centered” around ground potential, as is the case with a grounded-tip thermocouple:

Center-grounded voltage signal

Real circuit





Voltage-source model




If the centering is perfectly symmetrical, the signal voltage will be evenly “split” about ground potential. The two poles of a 30 millivolt thermocouple signal, for example, will measure +15 mV and -15 mV from ground. This is electrically equivalent to the elevated voltage signal model except with a negative common-mode voltage equal to half the signal voltage:

Center-grounded voltage signal




The type of analog voltage signal posed by our measurement application will become relevant when we connect it to a data acquisition device. Floating, ground-referenced, and elevated voltage signal sources each pose their own unique challenges to measurement, and any engineer or technician tasked with accurately measuring these signal types must understand these challenges. Data acquisition devices come in more than one type as well, and must be matched to the type of voltage signal in order to achieve good results.

Let's begin with our ground-referenced signal source: the simple RTD/resistor voltage divider circuit. If the divider is located close to the data acquisition (DAQ) analog input device, a single wire will suffice for connecting the two:



Each analog-to-digital converter (ADC) inside the DAQ unit is built to digitize a voltage with reference to ground, which is exactly the type of signal generated by our simple RTD voltage divider circuit. This type of DAQ analog input is called *single-ended*, and it is generally the default configuration on inexpensive DAQ units.

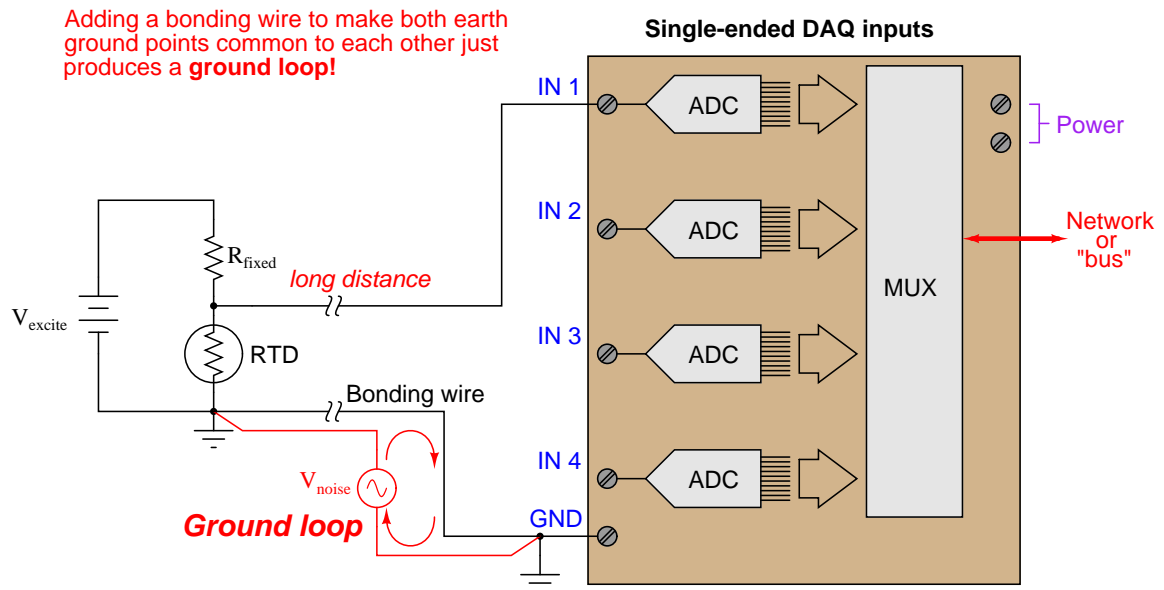
We cannot simply connect a single-ended DAQ input to a ground-referenced signal source using a single wire, however, if the two are located far apart from each other:



The problem here is that “all grounds are not created equal” over significant distances. If the ground path is literally through the earth (soil), there will be a myriad of noise sources adding spurious voltage to the measured signal: lightning strikes, ground leakage currents from AC power devices, and other sources of “noise” potential will become a part of the signal loop. Even continuous-metal ground paths over long distances can incur voltage drops significant enough to corrupt precision signal measurements. The grounding conductors used in AC power systems, for example, while continuous (no interruptions) between all points of use in the power system still drop enough millivoltage to significantly compromise instrumentation signals.

In essence, what appears to be a ground-referenced voltage signal source is actually an *elevated* voltage signal source, with the common-mode voltage being “noise” present between the two different ground locations.

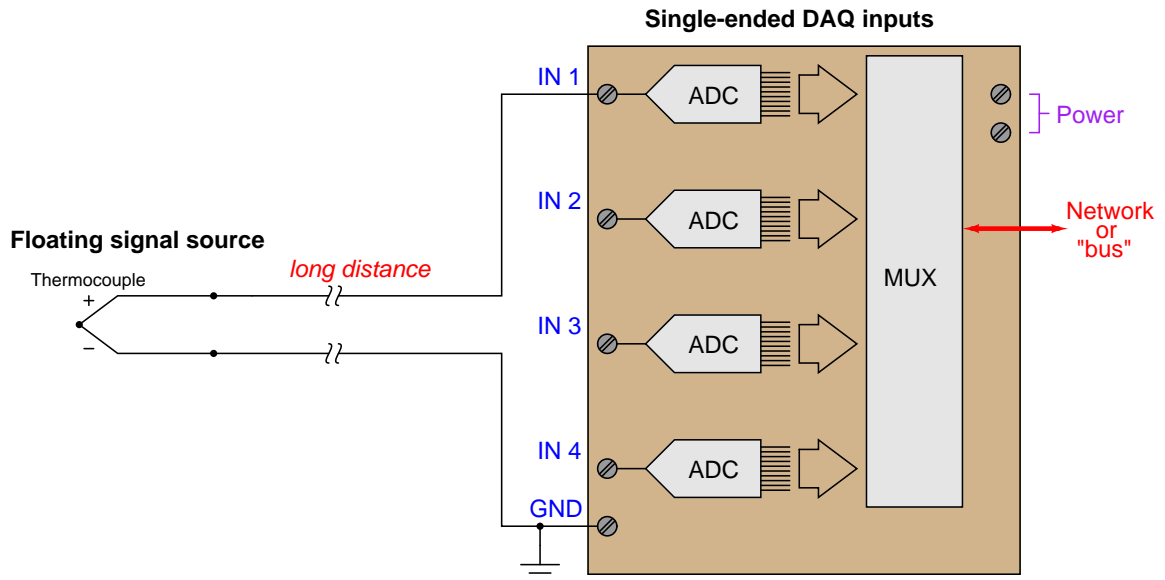
At first, it would seem that connecting the two grounds together with a dedicated length of wire would solve the differential ground problem. Unfortunately, it does not. The noise sources intercepted through the earth are often of significant power, which just means any conductor stretched between two different earth grounds may end up carrying substantial amounts of current, and subsequently dropping voltage along the way due to wire resistance ($V_{noise} = I_{ground}R_{wire}$).



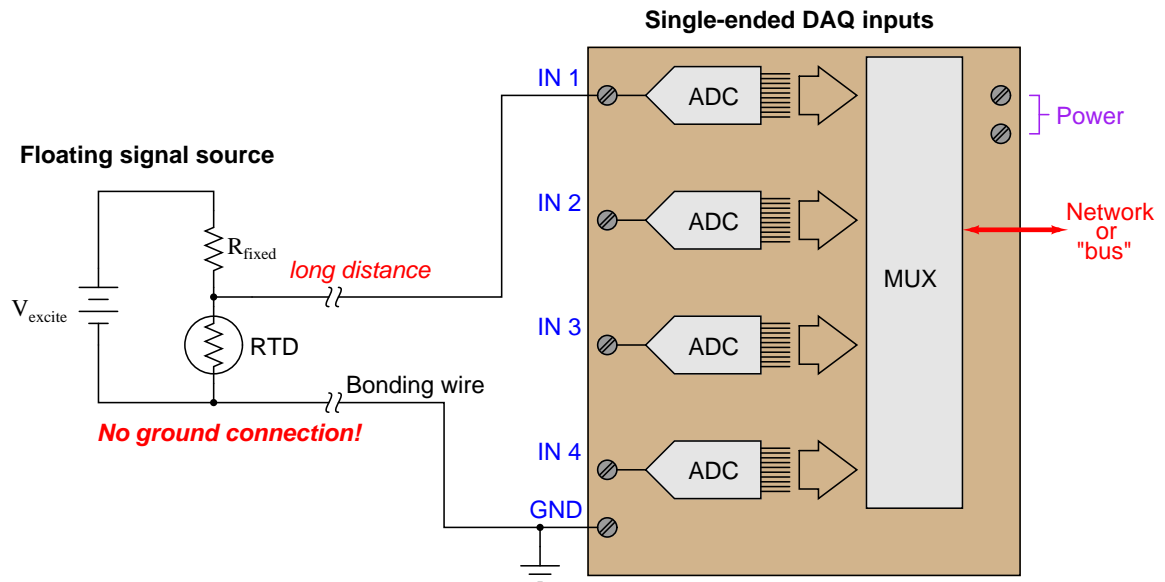
This is called a *ground loop*, and it should be avoided in signal circuits at all cost! Not only may the ground currents still produce significant noise voltage in the measurement circuit, but the ground currents may even become strong enough to damage the bonding wire! Ground loops are often unintentionally formed when the *shield* conductor of a long signal cable is earth-grounded at both ends.

A reasonable question to ask at this point is, “What constitutes a *long* distance when connecting ground-referenced signal sources to DAQ modules?” A simple rule to follow is that one cannot rely on ground points to be electrically common to each other (at least not common enough for precise signal-measurement purposes) if those points lie outside the same metal enclosure. If the signal source and DAQ analog input physically dwell inside the same metal enclosure, you can probably rely on the ground points being truly common to each other. If not, you should use some other means of measuring the signal. When in doubt, a sure test is to actually measure the potential difference between the source and DAQ ground points to see if any exists, being sure to check for AC noise voltage as well as DC.

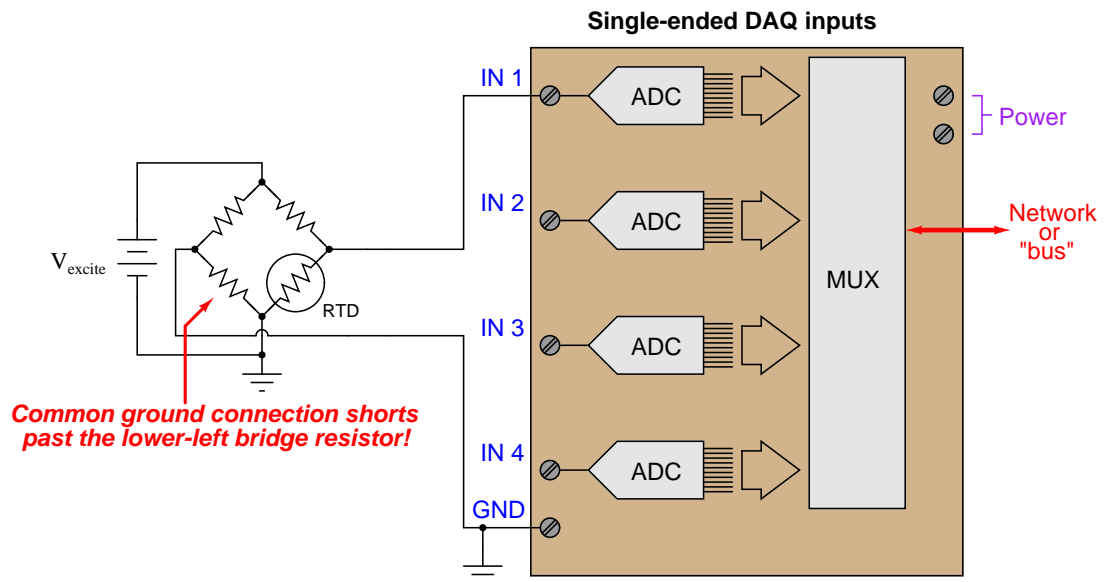
Here is an example of a single-ended DAQ module successfully measuring a voltage signal source located far away:



Since an ungrounded thermocouple has no connection whatsoever to any ground, there will be no ground loop when connected to a single-ended DAQ input. The same is true for battery-powered sensor circuits lacking connection to earth ground:




Single-ended analog DAQ inputs have trouble measuring elevated signal voltages regardless of distance. Here, we see an example where someone has tried to connect a single-ended DAQ input to a grounded-excitation RTD bridge, with common-mode voltage equal to one-half the excitation source voltage. The result is disastrous:



If you follow the bottom wire in this diagram carefully, you will see how it effectively jumpers past the lower-left resistor in conjunction with the two ground connections. Yet, eliminating the wire simply substitutes one problem for another: without the bottom wire in place, the voltage seen by input channel 1 on the DAQ will be $V_{signal} + V_{common-mode}$ rather than V_{signal} all by itself.

A clever way to solve the problem of measuring elevated signal sources is to use *two* analog DAQ channels: one to measure $V_{\text{signal}} + V_{\text{common-mode}}$ and the other to measure $V_{\text{common-mode}}$, and then digitally subtract one measurement from the other. Thus, two single-ended input channels may function as one *differential* input channel.


Here, we see an example of this measurement technique applied to the grounded-excitation RTD bridge:



The subtraction of the two channels' digitized values may take place in a controller separate from the DAQ module, or within the DAQ module if it is equipped with enough "intelligence" to perform the necessary calculations.


An added benefit of using this dual-channel method is that any noise voltage existing between the signal ground and the DAQ ground will be common to both channels, and thus should cancel when the two channels' signals are mathematically subtracted.

A more versatile hardware solution for measuring any form of voltage signal is a DAQ equipped with true *differential* input channels. Here, each ADC is equipped with its own instrumentation amplifier, measuring the difference in potential between two ungrounded input terminals:




In this example we see one DAQ module measuring four different voltage signal sources, with no interference from differential ground potentials or between the sources themselves. Each input channel of the ADC is electrically independent from the rest. The only limitation to this independence is a certain maximum common-mode voltage between the signal source and the DAQ's own power supply, determined by the design of the instrumentation amplifiers inside the DAQ module.

One important limitation of differential input channels, however, is that there must be some path for the instrumentation amplifiers' bias currents to the DAQ's power supply ground or else the channel will not function properly. This poses a problem where we intend to connect a differential analog input channel to a floating signal source like this:



One of the “simplifying assumptions” students learn about operational amplifier circuits is that the input terminals of an opamp draw negligible current. While this may be close enough to the truth when performing calculations on an opamp circuit, it is not absolutely true. All opamps exhibit some amount of *bias current* at their input terminals, small as these currents may be. Without a complete path to ground for these currents, the input transistor stage of the operational amplifier will not be properly biased, and the amplifier will fail to work as designed.

For this reason, we must connect high-value resistors (typically in the mega-ohm range so as not to load the signal voltage being measured) to each of the differential input terminals, and then to ground like this:



In summary, we may list some essential rules to follow when connecting analog DAQ inputs to voltage signal sources:

- Ground points in different locations may not actually be common (enough) to each other
- Never create ground loops through wiring (any conductor connected to ground at each end)
- Beware of common-mode (elevated) signal voltages
- Always ensure a path to power supply ground for amplifier bias currents

15.5 Digital data communication theory

One of the great benefits of digital technology is the ability to *communicate* vast amounts of information over networks. This very textbook you are reading was transmitted in digital form over the electronic network we call the *Internet*: a feat nearly impossible with any sort of analog electronic technology. The main benefit of digital data communication in industrial control is simple: no longer must we dedicate a single pair of wires to each and every variable we wish to measure and control in a facility as is necessary with analog (4-20 mA) signaling. With digital signaling, a single pair of wires or coaxial cable is able to convey a theoretically unlimited number of data points.

This benefit comes at a price, though: in order to communicate multiple variables (data points) over a single channel (wire pair), we must transmit and receive those signals one at a time. This means a digital communications system will necessarily exhibit some degree of *time delay* in acquiring, transmitting, receiving, and interpreting a signal. Analog systems, by contrast, are virtually instantaneous¹⁹. Thus, we see a contrast between analog and digital communication pitting channel capacity against speed:

Analog	Digital
Only one signal per channel	Many signals per channel possible
Instantaneous	Time-delayed

With modern electronic technology it is possible to build digital communication systems that are so fast, the time delays are negligible for most industrial processes, which renders the second comparison (instantaneous versus time-delayed) moot. If time is no longer an issue, the advantage that digital communication has over analog in terms of channel usage makes it the superior choice²⁰.

Another important advantage of digital data communication for industrial processes is increased *noise immunity*. Analog data is *continuous* by nature: a signal of 11.035 milliamps has a different meaning than a signal of 11.036 milliamps, because any measurable increment in signal represents a corresponding increment in the physical variable represented by that signal. A voltage value in a 0-5 volt digital signaling system of 0.03 volts, however, means *the exact same thing* as a voltage value of 0.04 volts: either one is still interpreted as a “0” or “low” state. *Any* amount of electrical noise imposed on an analog signal corrupts that signal to some degree. A digital signal, however, may tolerate a substantial amount of electrical noise with no corruption whatsoever.

¹⁹To be fair, there is such a thing as a time-multiplexed analog system for industrial data communication (I’ve actually worked on one such system, used to measure voltages on electrolytic “pots” in the aluminum industry, communicating the voltages across hundreds of individual pots to a central control computer).

²⁰There is, of course, the issue of *reliability*. Communicating thousands of process data points over a single cable may very well represent a dramatic cost savings in terms of wire, junction boxes, and electrical conduit. However, it also means you will lose all those thousands of data points if that one cable becomes severed! Even with digital technology, there may be reason to under-utilize the bandwidth of a signal cable.

Not surprisingly, though, the noise immunity enjoyed by digital signals comes with a price: a sacrifice in *resolution*. Analog signals are able to represent the smallest imaginable changes because they are continuously variable. Digital signals are limited in resolution by the number of bits in each data “word.” Thus, we see another contrast between analog and digital data representation:

Analog	Digital
Corrupted by any amount of noise	Immune to certain (limited) amounts of noise
Unlimited resolution	Limited resolution

With modern digital electronic technology, however, the “limited resolution” problem is almost nonexistent. 16-bit converter chipsets are commonly available today for input/output (I/O) modules on digital systems, providing a resolution of 2^{16} (65536) counts, or $\pm 0.00153\%$, which is good enough for the vast majority of industrial measurement and control applications.

This section will focus on *serial* data transmission, as opposed to *parallel*. In order to transmit digital data in parallel form, the number of wires scales directly with the number of bits in each data “word.” For example, if a 16-bit ADC chip were to communicate its data to some other digital device using a parallel network, it would require a cable with 16 wires (plus a common “ground” wire) *at minimum*²¹. Since this approach undercuts the “fewer wires” advantage that digital communications theoretically enjoys over analog communication, parallel data transmission is rarely seen in industry except for within the internal construction of a digital device (e.g. a parallel data bus inside a personal computer, or inside a PLC or DCS rack).

In serial communications systems, digital data is sent over a wire pair (or fiber optic cable, or radio channel) *one bit at a time*. A 16-bit digital “word” (two *bytes* in length) then will require a succession of 16 bits transmitted one after the other in time. How we represent each bit as an electrical signal, how we arrange those bits in time to group them into meaningful “words,” and how multiple devices share access to a common communications channel, is our next subject of exploration: the technical details of serial data communication.

²¹A common technique for high-speed parallel data communication over short distances (e.g. on a printed circuit board) is *differential signaling*, where each bit requires its own dedicated *pair* of conductors. A 16-bit parallel digital signal communicated this way would require 32 conductors between devices!

15.5.1 Serial communication principles

The task of encoding real-life data as a series of on-and-off electrical signals, and then sending those signals long distances over electrical cables (or optical fibers, or radio waves) requires mutually-agreed *standards* for the encoding, the “packaging” of those bits, the speed at which the bits are sent, methods for multiple devices to use a common channel, and a host of other concerns. This subsection will delineate the major points of compatibility necessary for digital devices to communicate serially. We begin with a brief exploration of some of the standards used in early *telegraph* systems.

An early form of digital communication was *Morse Code*, used to communicate alpha-numerical information as a series of “dots” and “dashes” over telegraph²² systems. Each letter in the alphabet, and each numerical digit (0 through 9) was represented in Morse Code by a specific series of “dot” and “dash” symbols, a “dot” being a short pulse and a “dash” being a longer pulse. A similar code system called the *Continental Code* was used for early radio (“radiotelegraph”) communications.

As primitive as these codes were, they encapsulated many of the basic principles we find in modern digital serial communication systems. First, a system of codes was necessary in order to represent English letters and numerals by electrical pulses. Next, there needed to be some way to delineate the beginning and end of each character.

For example, consider the Continental Code encoding for the word **NOWHERE**. By placing an extra space (a pause in time) between characters, it is easy to represent individual characters in the message:


"NOWHERE"

N
O
W
H
E
R
E

²²I do not expect any reader of this book to have firsthand knowledge of what a “telegraph” is, but I suspect some will have never heard of one until this point. Basically, a telegraph was a primitive electrical communication system stretching between cities using a keyswitch at the transmitting end to transmit on-and-off pulses and a “sounder” to make those pulses audible on the receiving end. Trained human operators worked these systems, one at the transmitting end (encoding English-written messages into a series of pulses) and one at the receiving end (translating those pulses into English letters).

If this space between characters were not present, it would be impossible to determine the message with certainty. By removing the spaces, we find multiple non-sensical interpretations are possible for the same string of “dots” and “dashes:”

*Same sequence of "dots" and "dashes,"
with multiple interpretations!*



For that matter, it is even possible to confuse the meaning of the text string “NOWHERE” when the individual characters are properly interpreted. Does the string of characters say “NOWHERE,” or does it say “NOW HERE”?

This simple example illustrates the need for *delimiting* in serial data communication. Some means must be employed to distinguish individual groups of bits (generally called *frames* or *packets*) from one another, lest their meanings be lost. In the days when human operators sent and interpreted Morse and Continental code messages, the standard delimiter was an extra time delay (pause) between characters, and between words. This is not much different from the use of whitespace to delineate words, sentences, and paragraphs typed on a page. Sentences would certainly be confusing to read if not for spaces!

In later years, when *teletype* machines were designed to replace skilled Morse operators, the concept of frame delineation had to be addressed more rigorously. These machines consisted of a typewriter-style keyboard which marked either paper strips or pages with dots corresponding to a 5-bit code called the *Baudot code*. The paper strip or sheets were then read electrically and converted into a serial stream of on-and-off pulses which were then transmitted along standard telegraph circuit lines. A matching teletype machine at the receiving end would then convert the signal stream into printed characters (a telegram). Not only could unskilled operators use teletype machines, but the data rate far exceeded what the best human Morse operators could achieve²³. However, these machines required special “start” and “stop” signals to synchronize the communication of each character, not being able to reliably interpret pauses like human operators could.

Interestingly, modern asynchronous²⁴ serial data communication relies on the same concept of

²³A test message sent in 1924 between two teletype machines achieved a speed of 1920 characters per minute (32 characters per second), sending the sentence fragments “THE WESTERN ELECTRIC COMPANY”, “FRESHEST EGGS AT BOTTOM MARKET PRICES”, and “SHE IS HIS SISTER”.

²⁴“Asynchronous” refers to the transmitting and receiving devices not having to be in perfect synchronization in order for data transfer to occur. Every industrial data communications standard I have ever seen is asynchronous rather than synchronous. In synchronous serial networks, a common “clock” signal maintains transmitting and receiving devices in a constant state of synchronization, so that data packets do not have to be preceded by “start” bits or followed by “stop” bits. Synchronous data communication networks are therefore more efficient (not having

“start” and “stop” bits to synchronize the transmission of data packets. Each new packet of serial data is preceded by some form of “start” signal, then the packet is sent, and followed up by some sort of “stop” signal. The receiving device(s) synchronize to the transmitter when the “start” signal is detected, and non-precision clocks keep the transmitting and receiving devices in step with each other over the short time duration of the data packet. So long as the transmitting and receiving clocks are close enough to the same frequency, and the data packet is short enough in its number of bits, the synchronization will be good enough for each and every bit of the message to be properly interpreted at the receiving end.


to include “extra” bits in the data stream) but also more complex. Most long-distance, heavy traffic digital networks (such as the “backbone” networks used for the Internet) are synchronous for this reason.

15.5.2 Physical encoding of bits

Telegraph systems were Boolean in nature: representing “dots” and “dashes” by one electrical state of the telegraph line, and pauses by another. When manually-actuated keyswitches were abandoned in favor of teletype machines, and Morse code abandoned in favor of the *Baudot* (5-bit) code for representing alphanumeric characters, the electrical nature of the telegraph (at least initially²⁵) remained the same. The telegraph line would either be energized or not, corresponding to *marks* or *spaces* made on the teletype paper.

Many modern digital communication standards represent binary “1” and “0” values in exactly this way: a “1” is represented by a “mark” state and a “0” is represented by a “space” state. “Marks” and “spaces” in turn correspond to different voltage levels between the conductors of the network circuit. For example, the very common EIA/TIA-232 serial communications standard (once the most popular way of connecting peripheral devices to personal computers, formerly called RS-232) defines a “mark” (1) state as -3 volts between the data wire and ground, and a “space” (0) state as $+3$ volts between the data wire and ground. This is referred to as *Non-Return-to-Zero*²⁶ or NRZ encoding:

Non-Return-to-Zero (NRZ) encoding




An easy way to remember the difference between a “mark” and a “space” in this scheme is to recall the operation of old telegraph printing units, specifically how they created marks and spaces on moving paper strip. When the printing unit was energized (i.e. the transmitting key was pressed, sending current through the solenoid coil of the printer, corresponding to a “1” state), the printer’s iron armature would be pulled *down* to draw a mark on the paper strip. When de-energized (transmitting key released, stopping current in the telegraph line, corresponding to a “0” state), the printer’s armature would spring-return *up* from the paper to leave a blank space.

²⁵Later versions of teletype systems employed audio tones instead of discrete electrical pulses so that many different channels of communication could be funneled along one telegraph line, each channel having its own unique audio tone frequency which could be filtered from other channels’ tones.

²⁶This simply refers to the fact that the signal never settles at 0 volts.

This is not the only way to represent binary bits, though. An alternative method is to use an oscillating (square-wave) signal, counting *up* and *down* transitions (pulse edges) at specific times to represent 1 and 0 states. This is called *Manchester encoding*, and it is used in the 10 Mbps (10 million bits per second) version of *Ethernet* and in both the *FOUNDATION Fieldbus* “H1” and *Profibus* “PA” instrumentation network standards:

Manchester encoding




Note how each binary bit (0 or 1) is represented by the *direction* of the voltage signal’s transition. A low-to-high transition represents a “1” state while a high-to-low transition represents a “0” state. Extra “reversal” transitions appear in the pulse stream only to set up the voltage level as needed for the next bit-representing transitions. The representation of bits by transitions rather than by static voltage levels guarantees the receiving device can naturally detect the clock frequency of the transmitted signal²⁷. Manchester data is therefore referred to as *self-clocking*.

Interpreting a Manchester waveform is easier than it first appears. The key is identifying which transitions represent “clocked” bits and which transitions represent “reversals” prior to bits. If we identify the widest periods of the waveform, we know the transitions in these periods must represent real bits because there are no reversals. Another way to say this is the greatest time period found between successive transitions in a Manchester waveform *is* the clock period. Once we identify this clock period, we may step along the waveform at that same period width to distinguish clocked bits from reversals.

²⁷This is most definitely *not* the case with NRZ encoding. To see the difference for yourself, imagine a continuous string of either “0” or “1” bits transmitted in NRZ encoding: it would be nothing but a straight-line DC signal. In Manchester encoding, it is *impossible* to have a straight-line DC signal for an indefinite length of time. Manchester signals *must* oscillate at a minimum frequency equal to the clock speed, thereby guaranteeing all receiving devices the ability to detect that clock speed and thereby synchronize themselves with it.

Yet another method for encoding binary 1 and 0 states is to use sine waves of different frequencies (“tone bursts”). This is referred to as *Frequency Shift Keying*, or *FSK*, and it is the method of encoding used in the HART “smart” instrument communications standard.

Frequency Shift Key (FSK) encoding



In the HART standard – based on the Bell 202 standard used in telephone data exchange – two complete cycles at 2200 Hz represents a “0” bit (space), while one complete cycle at 1200 Hz represents a “1” bit (mark). This standard, along with a series of previous and subsequent modulation standards, was invented as a way to exchange digital data over telephone networks, which were built to communicate audio-frequency²⁸ AC signals and thus could not reliably communicate the square-wave signals associated with direct digital data. By assigning digital values to different audio frequencies, serial data could be communicated over telephone channels as a series of sine-wave tones.

The same principle of FSK allows HART-compatible instruments to communicate digital data over cables simultaneously carrying DC current (4 to 20 mA) signals representing control system data. Since each bit of FSK-encoded data consists of complete AC cycles (one full positive swing for every full negative swing), the superposition of AC tones does not affect the time-averaged value of the DC milliamp signal.

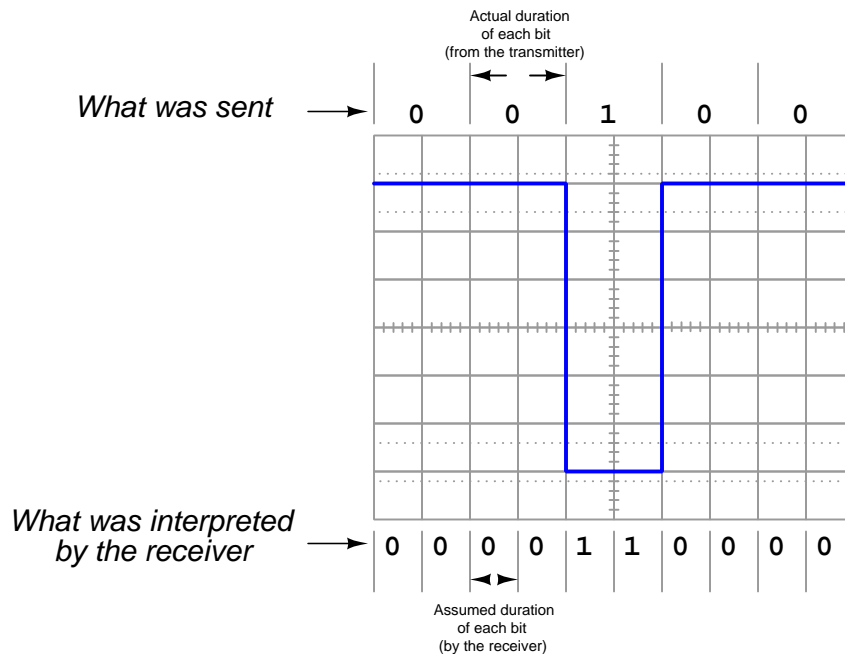
Other methods exist as well for encoding digital data along network cables, but these three are the most popular in industrial networks.

²⁸It is relatively easy to build an apparatus that makes HART tone signals audible: simply connect a small audio speaker to the low-impedance side of an audio transformer (8 ohms) and then connect the high-impedance side of that transformer (typically 1000 ohms) to the HART signal source through a coupling capacitor (a few microfarads is sufficient). When HART communications are taking place, you can hear the FSK tones reproduced by the speaker, which sound something like the noises made by a fax machine as it communicates over a telephone line.

15.5.3 Communication speed

In order to successfully communicate digital data along a network, there must not only be a standard agreed upon between transmitting and receiving devices for encoding bits (NRZ, Manchester, FSK, etc.), but there must also be a standard in place for the *speed* at which those bits will be sent. This is especially true for NRZ and FSK encoding, where the “clock” speed is not explicitly present in the signal²⁹.

For example, consider the confusion that could arise interpreting a NRZ signal if the transmitting device sends data at half the speed assumed by the receiving device:



Thus, one of the essential parameters in a serial data communication system is the *bit rate*, measured in *bits per second* (bps). Some communications standards have fixed bit rates, such as FOUNDATION Fieldbus H1 and Profibus PA, both standardized at exactly 31.25 kbps. Some, such as Ethernet, have a few pre-defined speeds (10 Mbps, 100 Mbps, 1 Gbps) defined by the specific transmitting and receiving hardware used. Others, such as EIA/TIA-232 may be arbitrarily set by the user at speeds ranging from 300 bps to over 115 kbps.

An older term sometimes used synonymously with bit rate is *baud rate*, however “bits per second” and “baud” are actually different things. “Baud” refers to the number of voltage (or current) alternations per second of time, whereas “bits per second” refers to the actual number of binary data bits communicated per second of time. Baud is useful when determining whether or not the bandwidth (the maximum frequency capacity) of a communications channel is sufficient for a certain communications purpose. For a string of alternating bits (e.g. 0101010101) using NRZ encoding,

²⁹This is one of the advantages of Manchester encoding: it is a “self-clocking” signal.

the baud rate is equivalent³⁰ to the bit rate: exactly one voltage transition for each bit. For a string of unchanging bits (e.g. 000000000000 or 111111111111) using NRZ encoding, the baud rate is far less than the bit rate. In systems using Manchester encoding, the worst-case³¹ baud rate will be exactly *twice* the bit rate, with two transitions (one up, one down) per bit. In some clever encoding schemes, it is possible to encode multiple bits per signal transition, such that the bit rate will actually be greater than the baud rate.


³⁰This is likely why “bit rate” and “baud rate” became intermingled in digital networking parlance: the earliest serial data networks requiring speed configuration were NRZ in nature, where “bps” and “baud” are one and the same.

³¹For Manchester encoding, “worst-case” is a sequence of identical bit states, such as 111111111111, where the signal must make an extra (down) transition in order to be “ready” for each meaningful (up) transition representing the next “1” state.

15.5.4 Data frames

As mentioned earlier in this section, serial data is usually communicated *asynchronously* in industrial networks. This means the transmitting and receiving hardware need not be in perfect synchronization to reliably send and receive digital data. In order for this to work, data must be sent in “frames” or “packets” of fixed (maximum) length, each frame preceded by a special “start” signal and concluded with a special “stop” signal. As soon as the transmitting device issues the “start” signal, the receiving device synchronizes to that start time, and runs at the pre-determined clock speed to gather the successive bits of the message until the “stop” signal is received. So long as the internal clock circuits of the transmitting and receiving devices are running at *approximately* the same speed, the devices will be synchronized closely enough to exchange a short message without any bits being lost or corrupted. There is such a thing as a *synchronous* digital network, where all transmitting and receiving devices are locked into a common clock signal so they cannot stray out of step with each other. The obvious advantage of synchronous communication is that no time need be wasted on “start” and “stop” bits, since data transfer may proceed continuously rather than in packets. However, synchronous communication systems tend to be more complex due to the need to keep all devices in perfect synchronization, and thus we see synchronous systems used for long-distance, high-traffic digital networks such as those use for Internet “backbones” and not for short-distance industrial networks.

Like bit rate, the particular scheme of start and stop bits must also be agreed upon in order for two serial devices to communicate with each other. In some networks, this scheme is fixed and cannot be altered by the user. Ethernet is an example of this, where a sequence of 64 bits (an alternating string of “1” and “0” bits ending with a “1, 1”; this is the “preamble” and “start frame delimiter” or “SFD” bit groups) is used to mark the start of a frame and another group of bits specifies the length of the frame (letting the receiver know ahead of time when the frame will end). A graphic description of the IEEE 802.3 standard for Ethernet data frames is shown here, illustrating the lengths and functions of the bits comprising an Ethernet frame:




Other serial networks offer choices for the user to select regarding these parameters. One such example is EIA/TIA-232, where the user may specify not only the bit rate, but also how many bits will be used to mark the end of the data frame. It is imperative in such systems that *all* transmitting and receiving devices within a given network be configured exactly the same, so that they will all “agree” on how to send and receive data. A screenshot from a UNIX-based serial communication terminal program (called *minicom*³²) shows these options:



In this particular screenshot, you can see the data rate options (extending from 300 bps all the way up to 230400 bps!), the number of data bits (from 5 to 8), and the number of stop bits (1 or 2), all configurable by the user. Of course, if this program were being used for communication of data between two personal computers, *both* of those computers would need these parameters set identically in order for the communication to take place. Otherwise, the two computers would not be in agreement on speed, number of data bits, and stop bits; their respective data frames simply would not match.

³²An equivalent program for Microsoft Windows is *Hyperterminal*. A legacy application, available for both Microsoft Windows and UNIX operating systems, is the serial communications program called *kermit*.

To give an example of an EIA/TIA-232 data frame might look like as a series of voltage states, consider this waveform communicating a string of eight bits (01011001), using NRZ encoding. Here, a single “start” marks the beginning of the data frame, while two successive “stop” bits end it. Also note how the bit sequence is transmitted “backwards,” with the least-significant bit (LSB) sent first and the most-significant bit (MSB) sent last³³:



Serial bitstream for the digital byte 01011001, where the least-significant bit (LSB) is sent first

Interestingly, the “mark” state (corresponding to a binary bit value of “1”) is the default state of the communications channel when no data is being passed. The “start” bit is actually a space (0). This is the standard encoding scheme for EIA/TIA-232, EIA/TIA-485, and some other NRZ serial communication standards.


One of the options you probably noticed in the “minicom” terminal program screenshot was something called *parity*. This is a simple form of error-checking used in many serial communication standards. The basic principle is quite simple: an extra bit is added at the end of the data frame (between the data and stop bits) to force the total number of “1” states to be either odd or even. For example, in the data stream just shown (10011010), there is an *even* number of “1” bits. If the serial device sending this eight-bit data group were configured for “odd” parity, it would append an additional “1” to the end of that frame to make the total number of “1” bits odd rather than even. If the next data group were 11001110 instead (already having an odd number of “1” bits), the transmitting device would have to attach a “0” parity bit on to the data frame in order to maintain an odd count of “1” bits.

Meanwhile, the receiving device is programmed to count up all the “1” bits in each data frame (including the parity bit), and check to see that the total number is still odd (if the receiving device is configured for odd parity just as the transmitting device, which the two should *always* be in agreement). Unlike the transmitting device which is tasked with *creating* the parity bit state, the receiving device is tasked with *reading* all the data bits plus the parity bit to check if the count is still as it should be. If any one bit somehow gets corrupted during transmission, the received frame will not have the correct parity, and the receiving device will “know” something has gone wrong. Parity

³³This is standard in EIA/TIA-232 communications.

does not suggest *which* bit got corrupted, but it will indicate if there was a single-bit³⁴ corruption of data, which is better than no form of error-checking at all.


The following example shows how parity-checking would work to detect a transmission error in a 7-bit data word. Suppose a digital device asynchronously transmits the character “T” using ASCII encoding (“T” = 1010100), with one start bit, one stop bit, and “odd” parity. Since the “start” bit is customarily a 0 state (space), the data transmitted in reverse order (LSB first, MSB last), the parity bit transmitted after the data’s MSB, and the “stop” bit represented by a 1 state (mark), the entire frame will be the following sequence of bits: 0001010101. Viewed on an oscilloscope display where a negative voltage represents a “mark” and a positive voltage represents a “space,” the transmitted data frame will look like this:



Note how the parity bit in this particular frame has been generated by the transmitting device as a 0 state, because the parity type is set for “odd,” and the transmitting device realizes that the 7-bit data word already has an odd number of 1 bits in it and doesn’t need another “1” for the parity bit. The pulse waveform you see above is how this data frame will be transmitted onto the network.

³⁴It should take only a moment or two of reflection to realize that such a parity check cannot detect an *even* number of corruptions, since flipping the states of any *two* or any *four* or any *six* (or even all eight!) bits will not alter the evenness/oddness of the bit count. So, parity is admittedly an imperfect error-detection scheme. However, it is certainly better than no error detection at all!

Now suppose this transmitted frame encounters a significant amount of electrical noise as it travels to the receiving device. If the frame reaches the receiver as shown in the next illustration, the receiving device will interpret the message incorrectly:



One of the bits has been corrupted by noise, such that the fifth transmitted data bit (which should be 1) is instead received as a 0. The receiving device, of course, has no knowledge of the noise present on the NRZ signal because all it “sees” is the “mark” or “space” states as interpreted by its input buffer circuitry. When the receiving device goes to count the number of 1 bits in the message (data plus parity bit, disregarding start and stop bits), however, it will count an even number of 1’s instead of an odd number of 1’s. Since the receiving device is also set for “odd” parity to match the transmitting device, it expects an odd number of 1’s in the received message. Thus, it “knows” there is a problem somewhere in this transmission, because the received parity is not odd as it should be.

Parity-checking does not tell us *which* bit is corrupted, but it does indicate that *something* has gone wrong in the transmission. If the receiving device is programmed to take action on receipt of a non-matching parity, it may reply with a request for the transmitting device to re-send the data as many times as necessary until the parity is correct.

If we look at the “minicom” terminal screenshot again to analyze the parity options, we see there are several to choose from:



The five options for parity in this program include *None*, *Even*, *Odd*, *Mark*, and *Space*. “No” parity is self-explanatory: the transmitting device does not attach an extra bit for parity at all, and the receiving device does not bother to check for it. Since the inclusion of a parity bit does add to the bulk of a data frame, it has the unfortunate effect of slowing down communications (more bit “traffic” occupying the channel than would otherwise need to be), thus the option to waive parity altogether for a more compact (faster) data frame. “Even” and “Odd” parity options work as previously described, with the transmitting device adding a parity bit to each frame to bring the total “1” bit count either to an even number or to an odd number (depending on the user’s configuration), and the receiving device checks for the same. “Mark” and “Space” are really of limited usefulness. In either of these two options, a parity bit is added, but the transmitting device does not bother to calculate the evenness or oddness of the data bits, rather simply making the parity bit always equal to 1 (“mark”) or 0 (“space”) as chosen by the user. The receiving device checks to see that the parity bit is always that value. These two options are of limited usefulness because the parity bit fails to reflect the status of the data being transmitted. The only corruption the receiving device can detect, therefore, is a corruption of the parity bit itself!

One will often find the communications parameters of a serial network such as this displayed in “shorthand” notation as seen at the top of the “minicom” terminal display: 38400 8N1. In this case, the terminal program is configured for a bit rate of 38400 bits per second, with a data field 8 bits long, no parity bit, and 1 stop bit. A serial device configured for a bit rate of 9600 bps, with a 7-bit data field, odd parity, and 2 stop bits would be represented as 9600 7O2.

Parity bits are not the only way to detect error, though. Some communication standards employ more sophisticated means. In the Ethernet (IEEE 802.3) standard, for example, each data frame is concluded with a *frame check sequence*, which is a collection of bits mathematically calculated by the transmitting device based on the content of the data. The algorithm is called a *cyclic*

redundancy check, or *CRC*, and is similar to the concept of “checksum” used by computers to check the integrity of data stored in hard disks and other “permanent” media. Like a parity algorithm, the CRC algorithm runs through a mathematical process whereby all the bits in the data field are counted, and a number is generated to reflect the statuses of those bits. The receiving device takes the received data field and performs the exact same mathematical algorithm, generating its own CRC value. If any of the data’s bits become corrupted during transmission, the two CRC values will not match, and the receiving device will know *something* has gone wrong.

Like parity, the CRC algorithm is not perfect. There exists a chance that just the right combination of errors may occur in transmission causing the CRC values at both ends to match even though the data is not identical, but this is highly unlikely (calculated to be one chance in 10^{14}). It is certainly better than having no error detection ability at all.


If the communications software in the receiving device is configured to take action on a detection of error, it may return a “request for re-transmission” to the transmitting device, so the corrupted message may be re-sent. This is analogous to a human being hearing a garbled transmission in a telephone conversation, and subsequently requesting the other person repeat what they just said.

Another option often found in serial data communications settings is something called *flow control*, not to be confused with the actual control of fluid through a pipe. In the context of digital communications, “flow control” refers to the ability of a receiving device to request a reduction in speed or even a complete cessation of data transmission if the speed of the transmitted data is too fast for the receiving device to keep pace. An example common to personal computers is that of a mechanical printer: while the computer may be able to transmit data to be printed at a very rapid pace, the printer is limited by the speed of its printing mechanism. In order to make the printing process go more smoothly, printers are equipped with *buffer memory* to store portions of the print job received from the transmitting computer that have not had time to print yet. However, these buffers are of finite size, and may become overwhelmed on large print jobs. So, if and when a printer detects its buffer near full capacity, it may issue a command to the computer to freeze serial data transmission until the printer’s buffer has had some time to empty. In other words, the printer can send a message to the computer saying “Stop!” when its buffer is full, then later send another message saying “Resume” when its buffer is empty enough to resume filling. Thus, the receiving device has control over the flow of data necessary to manage its buffer resources.

Flow control in serial networks may take place in either *hardware* mode or *software* mode. “Hardware” mode refers to the existence of additional connector pins and cable conductors specifically designated for such “halt” and “resume” signals. “Software” mode refers to data codes communicated over the regular network channel telling the transmitting device to halt and resume. Software flow control is sometimes referred to as XON/XOFF in honor of the names given to these codes³⁵. Hardware flow control is sometimes referred to as RTS/CTS in honor of the labels given to the serial port pins conveying these signals.

³⁵The “XOFF” code tells the transmitting device to halt its serial data stream to give the receiving device a chance to “catch up.” In data terminal applications, the XOFF command may be issued by pressing the key combination <Ctrl><S>. This will “freeze” the stream of text data sent to the terminal by the host computer. The key combination <Ctrl><Q> sends the “XON” code, enabling the host computer to resume data transmission to the terminal.

The following screen shot shows options for flow control in the “minicom” terminal program:



```
Welcome to minicom 2.1
OPTI
Comp A - Serial Device      : /dev/tty50
      B - Lockfile Location : /var/lock
Pres  C - Callin Program    :
      D - Callout Program   :
      E - Bps/Par/Bits     : 38400 8N1
      F - Hardware Flow Control : Yes
      G - Software Flow Control : No

Change which setting? █

Screen and keyboard
Save setup as df1
Save setup as..
Exit

CTRL-A Z for help | 38400 8N1 | NOR | Minicom 2.1 | VT102 | Offline
```

Here, you can see “hardware” flow control enabled and “software” flow control disabled. The enabling of “hardware” flow control means the serial communications cable must be equipped with the necessary lines to convey these handshaking signals (when needed) between devices. Software flow control tends to be the more popular option, the advantage of this of course being fewer conductors necessary in the serial data cable. The disadvantage of using software flow control over hardware is a slight inefficiency in data throughput, since the XON and XOFF commands require time to be transmitted serially over the same network as the rest of the data.

15.5.5 Channel arbitration

When two or more communication devices exchange data, the directions of their communication may be classified into one of two categories: *simplex* or *duplex*. A “simplex” network is one-way communication only. A sensor outputting digital data to a remotely-located indicator over a digital network would be an example of simplex communication, where the flow of information goes from sensor to indicator, and never the other direction. A public-address (PA) system is an analog example of a simplex communication system, since audio information only goes in one direction (from the person with the microphone to the audience).

“Duplex” communication refers to two-way data exchange. Voice telephony is an analog example of two-way (duplex) communication, where either person at the end of the connection can hear the other person talking. Duplex communication may be further subdivided into *half-duplex* and *full-duplex*, referring to whether or not the two-way communication may be simultaneous. In a “full-duplex” system, both devices may transmit data to each other simultaneously because they have separate channels (separate wires, or optical fibers, or radio frequencies) for their respective transmissions. In a “half-duplex” system, only one device may transmit at any time because the devices must share a common channel. A telephone system is an example of a full-duplex system, although it may be rather difficult for the people to understand each other when they are speaking over one another. A push-to-talk radio system (“walkie-talkie”) is an example of a half-duplex system, where each person must take turns talking.

Most industrial data networks are half-duplex, if only for the reason that most networks consist of more than two devices on a network segment. When more than two devices share a network, there are not enough data channels to allow *all* of the devices to simultaneously transmit and listen to each other. Thus, virtually any network supporting more than two devices will be half-duplex at best, and may even be limited to simplex operation in some cases.

In half-duplex systems, there must be some way for the respective devices to “know” when they are allowed to transmit. If multiple devices sharing one communications channel attempt to transmit simultaneously, their messages will “collide” in such a way that no device on the network will be able to interpret either message. The problem is analogous to two people simultaneously pressing the “talk” buttons on their two-way radio units: neither of the talking people can hear each other, and anyone else on the same channel hears the garbled amalgam of those two peoples’ superimposed transmissions. In order to avoid this scenario in a half-duplex network, there must be some strategy to coordinate transmissions so only one device “talks” at any given time. The problem of deciding “who” gets to “talk” at any given time is generally known as *channel arbitration*. Several strategies for addressing this problem have been developed in the data communications field, a few of which will be described in this subsection.

Master-slave

Our first method works on the principle of having only one device on the network (the “master”) with permission to arbitrarily transmit data. All other devices on the network are “slaves,” which may only respond in direct answer to a query from the master. If the network happens to be simplex in nature, slave devices don’t even have the ability to transmit data – all they can do is “listen” and receive data from the one master device.

For example, in a half-duplex master-slave network, if one slave device has data that needs to be sent to another slave device, the first slave device must wait until it is prompted (“polled”) by the master device before it is allowed to transmit that data to the network. Once the data is transmitted, any and all slave devices may receive that transmission, since they all “listen” to the same communications channel.

An example of an industrial network using master-slave channel arbitration is HART *multidrop*, where multiple HART field instruments are parallel-connected on the same wire pair, and one device (usually a dedicated computer) serves as the master node, polling the field instruments one at a time for their data.

Another example of a master-slave industrial network is a *Modbus* network connecting a programmable logic controller (PLC) to multiple variable-frequency motor drives (VFDs). The master device (the PLC) initiates all communications, with the slave devices (the motor drives) at most replying to the PLC master (and in many cases not replying at all, but merely receiving data from the PLC in simplex mode).

Master-slave arbitration is simple and efficient, but suffers from one glaring weakness: if the master device happens to fail, all communication on the network halts. This means the ability of *any* device on the network to transmit information utterly depends on the proper function of *one* device, representing a high level of dependence on that one (master) device’s function.

Some master-slave networks address this problem by pre-assigning special “back-up” status to one or more slave devices. In the event that the master device fails and stops transmitting for a certain amount of time, the back-up device becomes “deputized” to act as the new master, taking over the role of the old master device by ensuring all slave devices are polled on schedule.

Token-passing

Another method of arbitrating which device gets to transmit on a channel in a half-duplex network is the *token-passing* method. Here, a special data message called the “token” serves as temporary authorization for each device to transmit. Any device in possession of the token is allowed to act as a master device, transmitting at will. After a certain amount of time, that device must relinquish the token by transmitting the token message on the network, complete with the address of the next device. When that other device receives the token message, it switches into master mode and transmits at will. The strategy is not unlike a group of people situated at a table, where only one of them at a time holds some object universally agreed to grant speaking authority to the holder.

Token-passing ensures only one device is allowed to transmit at any given time, and it also solves the problem inherent to master-slave networks of what happens when the master device fails. If one of the devices on a token-passing network fails, its silence will be detected after the last token-holding device transmits the token message to the failed device. After some pre-arranged period of time, the last token-holding device may re-transmit the token message to the *next* device after the one that failed, re-establishing the pattern of token sharing and ensuring all devices get to “speak” their turn once more.

Examples of token-passing networks include the general-purpose Token Ring network standard (IEEE 802.5) and the defunct Token Bus (IEEE 802.4). Some proprietary industrial networks such as Honeywell’s TDC 3000 network (called the *Local Control Network*, or *LCN*) utilize token-passing arbitration.

Token-passing networks require a substantially greater amount of “intelligence” built into each network device than master-slave requires. The benefits, though, are greater reliability and a high level of bandwidth utilization. That being said, token-passing networks may suffer unique disadvantages of their own. For example, there is the question of what to do if such a network becomes severed, so that the one network is now divided into two segments. At the time of the break, only one device will possess the token, which means only one of the segments will possess any token at all. If this state of affairs holds for some time, it will mean the devices lucky enough to be in the segment that still has the token will continue communicating with each other, passing the token to one another over time as if nothing was wrong. The isolated segment, however, lacking any token at all, will remain silent even though all its devices are still in working order and the network cable connecting them together is still functional. In a case like this, the token-passing concept fares no better than a master-slave network. However, what if the designers of the token-passing network decide to program the devices to automatically generate a new token in the event of prolonged network silence, anticipating such a failure? If the network becomes severed and broken into multiple segments, the isolated segments will now generate their own tokens and resume communication between their respective devices, which is certainly better than complete silence as before. The problem now is, what happens if a technician locates the break in the network cable and re-connects it? Now, there will be *multiple* tokens on one network, and confusion will reign!

Another example of a potential token-passing weakness is to consider what would happen to such a network if the device in possession of the token failed before it had an opportunity to relinquish the token to another device. Now, the entire network will be silent, because no device possesses the token! Of course, the network designers could anticipate such a scenario and pre-program the devices to generate a new token after some amount of silence is detected, but then this raises the possibility of the previously-mentioned problem when a network becomes severed and multiple tokens arise in

an effort to maintain communication in those isolated network segments, then at some later time the network is re-connected and now multiple tokens create data collision problems.

TDMA

A method of channel arbitration similar to token-passing is *TDMA*, or “Time Division Multiple Access.” Here, each device is assigned an absolute “time slot” in a repeating schedule when it alone is allowed to transmit. With token-passing, permission to transmit is granted to each device by the previous device as it relinquishes the token. With TDMA, permission to transmit is granted by an appointment on a fixed time schedule. TDMA is less time-efficient than token-passing because devices with no data to transmit still occupy the same amount of time in the schedule as when they have data to transmit. However, TDMA has the potential to be more tolerant of device failure and network segmentation than token-passing because neither the failure of a device nor segmentation of the network can prevent remaining devices from communicating with each other. If a device fails (becomes “silent”) in a TDMA network, that time slot simply goes unused while all other communication continues unabated. If the network becomes severed, each set of devices in the two segments will still follow their pre-programmed time schedule and therefore will still be able to communicate with each other.

Examples of TDMA networks include the *WirelessHART* and *ISA100.11a* radio instrumentation standards. The GSM cell phone network standard also includes TDMA as part of a larger strategy to manage access between multiple cell phones and cell towers. TDMA arbitration works very well for wireless (radio) networks where the communication channel is inherently unreliable due to physical obstacles. If a device on a TDMA wireless network falls out of range or becomes blocked, the rest of the network carries on without missing a step.

Practical TDMA networks are not quite as fault tolerant as the idealized vision of TDMA previously described. Real TDMA networks do depend on some “master” device to assign new time slots and also to maintain synchronization of all device clocks so that they do not “lose their place” in the schedule. If this master device fails, the TDMA network will lose the ability to accept new devices and will (eventually) lose synchronization.

In light of this fact, it might appear at first that TDMA is no better than master-slave arbitration, since both ultimately depend on one master device to manage communication between all other devices. However, TDMA does offer one significant benefit over master-slave, and that is more efficient use of time. In a master-slave network, the master must poll each and every device on the network to check if it has data to transmit. This polling requires additional network time beyond that required by the “slave” devices to report their data. In a TDMA network, the master device need only occupy time transmitting to the network when updating time-slot assignments and when broadcasting time synchronization messages. You can think of TDMA as being a “smarter” version of master-slave arbitration, where the devices need only be told once when they may transmit, rather than having to be told every single time when they may transmit.

CSMA

A completely different method of channel arbitration is where any and all devices have permission to transmit when the network is silent. This is generally called *CSMA*, or “Carrier Sense Multiple Access.” There are no dedicated master and slave devices with CSMA, nor are devices permitted to transmit in a pre-determined order as with token-passing or in a pre-determined schedule as with TDMA. *Any* device on a CSMA network may “talk” in any order and at any time whenever the network is free. This is analogous to an informal conversation between peers, where anyone is permitted to break the silence.

Of course, such an egalitarian form of channel arbitration invites instances where two or more devices begin communicating simultaneously. This is called a *collision*, and must be addressed in some manner in order for any CSMA network to be practical.

Multiple methods exist to overcome this problem. Perhaps the most popular in terms of number of installed networks is *CSMA/CD* (“Carrier Sense Multiple Access with Collision Detection”), the strategy used in Ethernet. With CSMA/CD, all devices are not only able to sense an idle channel, but are also able to sense when they have “collided” with another transmitting device. In the event of a collision, the colliding devices both cease transmission, and set random time-delays to wait before re-transmission. The individual time delays are randomized to decrease the probability that a re-collision between the same devices will occur after the wait. This strategy is analogous to several peers in one group holding a conversation, where all people involved are equally free to begin speaking, and equally deferential to their peers if ever two or more accidentally begin speaking at the same time. Occasional collisions are normal in a CSMA/CD network, and should not be taken as an indication of trouble unless their frequency becomes severe.

A different method of addressing collisions is to pre-assign to each device on the network a priority number, which determines the order of re-transmission following a collision. This is called *CSMA/BA*, or “Carrier Sense Multiple Access with Bitwise Arbitration,” and it is analogous to several people of different social levels in one group holding a conversation. All are free to speak when the room is silent, but if two or more people accidentally begin speaking at the same time, the person of highest “rank” is allowed to continue while the “lower-rank” person(s) must wait. This is the strategy used in DeviceNet, an industrial network based on CAN technology, one of the more popular data networks used in automotive engine control systems.

Some CSMA networks lack the luxury of collision detection, and must therefore strive to prevent collisions rather than gracefully recover from them. Wireless digital networks are an example where collision detection is not an option, since a wireless (radio) device having a single antenna and a single channel cannot “hear” any other devices’ transmissions while it is transmitting, and therefore cannot detect a collision if one were to occur. A way to avoid collisions for such devices is to pre-assign each device on the network with a priority number, which determines how long each device is forced to wait after detecting a “quiet” network before it is allowed to transmit a new message. So long as no two devices on the network have the same “wait” time, there will be no collisions. This strategy is called *CSMA/CA*, or “Carrier Sense Multiple Access with Collision Avoidance,” and is the technique used for WLAN networks (the IEEE 802.11 specification). A consequence of collision avoidance, though, is unequal access to the network. Those devices with higher-priority (shorter wait times) will always have an advantage in transmitting their data over devices of lower priority. The degree of disparity in network access grows as more devices occupy the network. CSMA/CA is analogous to a group of shy people talking, each person afraid to speak at the same time as another,

and so each person waits a different amount of time following the conclusion of the last utterance before daring to speak. This sort of ultra-polite behavior may ensure no one accidentally interrupts another, but it also means the shiest person will hardly ever get a chance to speak.

One characteristic distinguishing all CSMA networks from master-slave, token-passing, and TDMA networks is a lack of *determinism*. “Determinism” is the ability to guarantee communication of a message within a specified maximum time. A master-slave or TDMA network following a repeating schedule guarantees that the time will not take longer than the period of the cycle. A token-passing network following a definite order guarantees that the time will not take longer than the number of devices times the maximum time each device may hold the token. A CSMA network, at least hypothetically, can prevent any device from transmitting its message if it keeps being blocked by other devices transmitting (the one highest-priority device in a CSMA/BA or CSMA/CA network being an exception of course). Determinism is important in industrial control systems where communication delays may adversely affect the stability of a control loop, and it is especially important in safety control systems where fast action is needed to avert catastrophe.

A potential problem in any digital network, but particularly networks employing CSMA arbitration, is something known as *jabbering*. If a network device happens to fail in such a way that it ceaselessly transmits a signal on the network, none of the other CSMA devices will ever be allowed to transmit because they continuously detect a “carrier” signal from the jabbering device³⁶. Some Ethernet components sport *jabber latch* protection circuits designed to detect jabber and automatically cut off the offending device from the network, or employ “store-and-forward” buffering which is able to block jabbered data frames.

³⁶I once encountered this very type of failure on the job, where a copper-to-fiber adapter on a personal computer’s Ethernet port jammed the entire network by constantly spewing a meaningless stream of data. Fortunately, indicator lights on all the channels of the communications equipment clearly showed where the offending device was on the network, allowing us to take it out of service for replacement.

15.5.6 The OSI Reference Model

Digital data communication may be described in many ways. For example, a connection formed between two computers to exchange a text document is a multi-layered activity, involving many steps to convert human language into electrical impulses for transmission, then re-convert those electrical impulses into human language again at the receiving end. Not surprisingly, there usually exist many different ways to perform this same task: different types of networks, different encodings, different communications and presentation software, etc.

To illustrate by analogy, think of all the actions and components necessary to transport items using an automobile. In order to move furniture from an apartment to a house, for example, you would require the following:

- An appropriate vehicle
- Addresses or directions for both locations
- A driver's license and knowledge of driving rules
- Fuel for the vehicle
- Knowledge of how to safely stack furniture for transport
- Knowledge of how the furniture is to be placed in the house

These details may seem trivial to mention, as human beings familiar with the common task of moving personal belongings from one location to another, but imagine having to describe every single action and component to someone from a primitive culture ignorant of vehicles, addresses, maps, driver's licenses, fuel, etc. One way to help describe all this complexity would be to assign different people to different layers of detail. For example, an automotive engineer could discuss the details of how engines burn fuel to do mechanical work (propelling the vehicle) while a furniture loader could describe how furniture is to be loaded and taken off the vehicle. A driving instructor could then explain all the procedures of safely driving the vehicle, while a city planner could explain the organization of streets and addresses in relation to a city map. Finally, an interior decorator could wax eloquent on the proper placement of furniture in the house. Each person would be describing a different aspect of the furniture move, each one of those aspects being important to the overall goal of moving furniture from one location to another.

Moreover, for each one of the aspects described by a specialist, there may exist several different alternatives. For example, there are many different models and designs of vehicle one might use for the job, and there may be different driving rules depending on where the two locations are for the move. Addresses and directions will *certainly* vary from city to city, and even within one city there will be alternative routes between the two locations. Finally, there is virtually no end to arrangements for furniture at the destination house, each one with its own functional and esthetic merits.

By the same token, the task of transporting digital data may be divided into similar categories. In order to move and process data from one computer to another, you need the following:

- An appropriate cable (or radio link) connecting the two computers
- Standardized electrical signals to represent bit states
- Addresses for each computer on the network
- Algorithms specifying how each computer takes turns “talking” on the common network
- Algorithms specifying how to organize packets of data to be sent and received serially
- Software to format the data on the transmitting end and interpret the data on the receiving end

Each of these aspects is important to the overall goal of creating, moving, and interpreting digital data between two or more computers, and there are many alternative methods (standards) for each aspect. We may represent 0 and 1 bits using NRZ (Non-Return to Zero) encoding, Manchester encoding, FSK modulation, etc.; the signals may be electrical or they may be optical or they may even be radio waves; the options for electrical cables and connector types are many. Bits may be framed differently as they are packaged for transmission, and arbitration between devices on the network managed in a variety of different ways. How we address multiple devices on a network so messages get routed to their proper destinations is important as well.

A scheme originally intended as a formal standard, but now widely regarded as a general model to describe the portions of other standards, helps us clarify the complexity of digital communications by dividing communication functions into seven distinct “layers.” Developed by the *ISO* (International Organization for Standards)³⁷ in 1983, the *OSI Reference Model* divides communication functions into the following categories, shown in this table with examples:

Layer 7 Application	This is where digital data takes on practical meaning in the context of some human or overall system function. <i>Examples: HTTP, FTP, HART, Modbus</i>
Layer 6 Presentation	This is where data gets converted between different formats. <i>Examples: ASCII, EBCDIC, MPEG, JPG, MP3</i>
Layer 5 Session	This is where "conversations" between digital devices are opened, closed, and otherwise managed for reliable data flow. <i>Examples: Sockets, NetBIOS</i>
Layer 4 Transport	This is where complete data transfer is handled, ensuring all data gets put together and error-checked before use. <i>Examples: TCP, UDP</i>
Layer 3 Network	This is where the system determines network-wide addresses, ensuring a means for data to get from one node to another. <i>Examples: IP, ARP</i>
Layer 2 Data link	This is where basic data transfer methods and sequences (frames) are defined within the smallest segment(s) of a network. <i>Examples: CSMA/CD, Token passing, Master/Slave</i>
Layer 1 Physical	This is where data bits are equated to electrical, optical, or other signals. Other physical details such as cable and connector types are also specified here. <i>Examples: EIA/TIA-232, 422, 485, Bell 202</i>

The vast majority of digital networking standards in existence address mere portions of the 7-layer model. Any one of the various Ethernet standards, for example, applies to layers 1 and 2, but none of the higher-level layers. In other words, Ethernet is a means of encoding digital information in electronic form and packaging that data in a standard format understandable to other Ethernet devices, but it provides no functionality beyond that. Ethernet does *not* specify how data will be routed over large-area networks, how to manage data-exchange sessions between computers (opening connections, initiating data transfer, closing connections), or how to format the data to represent real-world variables and media. Common industrial network standards such as EIA/TIA-232 and EIA/TIA-485 don't even go that far, being limited mostly to layer 1 concerns (signal voltage levels, wiring, and in some cases types of electrical connectors). For example, EIA/TIA-485 does not specify

³⁷If you are thinking the acronym should be “IOS” instead of “ISO,” you are thinking in terms of English. “ISO” is a non-English acronym!

how to address multiple devices connected to a common electrical network – all it does is specify what voltage levels represent “0” and “1” bits.

By contrast, some other industrial networking standards specify nothing about lower-level layers, instead focusing on high-level concerns. Modbus, for example, is concerned only with layer 7, and not with any of the lower-level layers³⁸. This means if two or more industrial devices on a network (such as programmable logic controllers, or PLCs) use “Modbus” to communicate with each other, it refers only to the high-level programming codes designed to poll and interpret data within those devices. The actual cable connections, electrical signals, and communication techniques used in that “Modbus” network may vary widely. Anything from EIA/TIA-232 to Ethernet to a wireless network such as WLAN may be used to actually communicate the high-level Modbus instructions between PLCs.

Due to the fact that most industrial instrumentation networks span limited physical distances and require low latency (short time delays) in order to exchange process-critical data between devices, the most common layers of the OSI model seen in industrial networks are layers 1, 2, and 7. The middle layers become important when data leaves the plant site and must be routed along major communications networks such as the Internet. Inside an industrial facility, however, the networks must merely encode the data in some physical media (layer 1), package and address the data so it makes it to the appropriate device(s) (layer 2), and give that data contextual meaning (layer 7).

The following sections explore some common networking standards used for industrial instrumentation systems. The OSI Reference Model will be mentioned where appropriate.

15.6 EIA/TIA-232, 422, and 485 networks

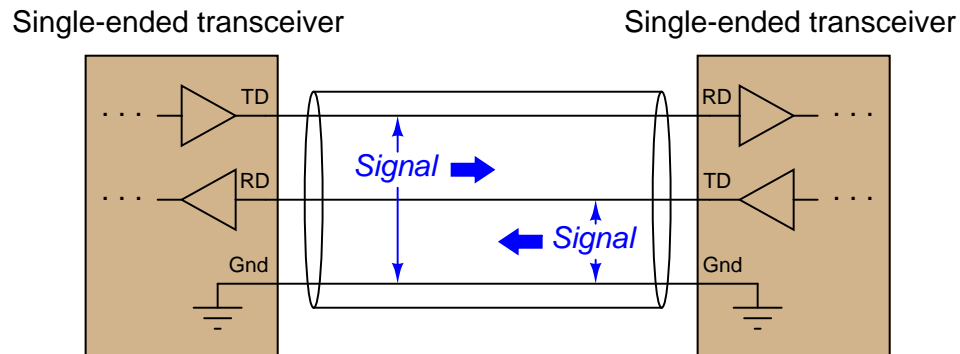
Some of the simplest types of digital communication networks found in industry are defined by the EIA (Electronic Industry Alliance) and TIA (Telecommunications Industry Alliance) groups, under the numerical labels 232, 422, and 485. This section discusses these three network types.

³⁸It should be noted here that some network standards incorporating the name “Modbus” actually do specify lower-level concerns. *Modbus Plus* is a layer 2 standard, for example.

15.6.1 EIA/TIA-232

The EIA/TIA-232C standard, formerly³⁹ known as *RS-232*, is a standard defining details found at layer 1 of the OSI Reference Model (voltage signaling, connector types) and some details found at layer 2 of the OSI model (asynchronous transfer, “flow control” or “handshaking” signals between transmitting and receiving devices). In the early days of personal computers, almost every PC had either a 9-pin or a 25-pin connector (and sometimes multiple of each!) dedicated to this form of digital communication. For a while, it was *the* way peripheral devices such as keyboards, printers, modems, and mice connected to personal computers. USB (Universal Serial Bus) has now all but replaced EIA/TIA-232 for personal computers, but it still lives on in the world of industrial devices.

EIA/TIA-232 networks are point-to-point, intended to connect only two devices⁴⁰. The signaling is *single-ended* (also known as *unbalanced*), which means the respective voltage pulses are referenced to a common “ground” conductor, a single conductor used to transfer data in each direction:




EIA/TIA-232 specifies positive and negative voltages (with respect to the common ground conductor) for its NRZ signaling: any signal more negative than -3 volts detected at the receiver is considered a “mark” (1) and any signal more positive than $+3$ volts detected at the receiver is considered a “space” (0). EIA/TIA-232 transmitters are supposed to generate -5 and $+5$ volt signals (minimum amplitude) to ensure at least 2 volts of noise margin between transmitter and receiver. The voltage limits and NRZ encoding of EIA/TIA-232 comprise the OSI layer 1 elements of the standard.

³⁹The designation of “RS-232” has been used for so many years that it still persists in modern writing and manufacturers’ documentation, despite the official status of the EIA/TIA label. The same is true for EIA/TIA-422 and EIA/TIA-485, which were formerly known as RS-422 and RS-485, respectively.

⁴⁰“Daisy-chain” networks formed of more than two devices communicating via EIA/TIA-232 signals have been built, but they are rarely encountered, especially in industrial control applications.

Cable connectors are also specified in the EIA/TIA-232 standard, the most common being the DE-9⁴¹ (nine-pin) connector. The “pinout” of a DE-9 connector for any *DTE* (Data Terminal Equipment) device at the end of an EIA/TIA-232 cable is shown here:

DE-9 cable connector



Pin number	Assignment	Abbreviation
1	Carrier Detect	CD
2	Received Data	RD
3	Transmitted Data	TD
4	Data Terminal Ready	DTR
5	Signal Ground	Gnd
6	Data Set Ready	DSR
7	Request To Send	RTS
8	Clear To Send	CTS
9	Ring Indicator	RI

Those terminals highlighted in **bold** font represent those connections absolutely essential for any EIA/TIA-232 link to function. The other terminals carry optional “handshaking” (“flow control”) signals⁴² specified for the purpose of coordinating data transactions (these are the OSI layer 2 elements of the EIA/TIA-232 standard).


For *DCE* (Data Communications Equipment⁴³) devices such as modems, which extend the EIA/TIA-232 signal path onward to other devices, the assignments of transmitting and receiving pins are swapped. For example, pin 2 is the Transmitted Data (TD) output while pin 3 is the Received Data (RD) input on a DCE device. This allows straight pin-to-pin cable connections between the DTE and DCE devices, so the transmit pin of the DTE device connects to the receive pin of the DCE, and vice-versa.

⁴¹Often (incorrectly) called a “DB-9” connector.


⁴²The way hardware-based flow control works in the EIA/TIA-232 standard involves two lines labeled RTS (“Request To Send”) and CTS (“Clear To Send”) connecting the two devices together on a point-to-point serial network in addition to the TD (“Transmitted Data”) and RD (“Received Data”) and signal ground lines. Like the TD and RD terminals which must be “crossed over” between devices such that the TD of one device connects to the RD of the other device and vice-versa, the RTS and CTS terminals of the two devices must be similarly crossed. The RTS is an output line while the CTS is an input, on both devices. When a device is able to receive data, it activates its RTS output line to *request* data. A device is not permitted to transmit data on its TD line until it is *cleared* to send data by an active state on its CTS input line.

⁴³Also known by the unwieldy acronym *DCTE* (Data Circuit Terminating Equipment). Just think of “DTE” devices as being at the very end (“terminal”) of the line, whereas “DCE” devices are somewhere in the middle, helping to exchange serial data between DTE devices.

The following diagram shows the minimum cable requirements for an EIA/TIA-232 serial communication link consisting of a pair of DTEs connecting through a pair of DCEs. This minimal point-to-point network assumes the devices are configured for either *software* flow control (i.e. digital codes send over the TD/RD lines requesting the transmitting device to halt and resume) or no flow control at all:




In order to utilize hardware flow control, the EIA/TIA-232 cable(s) must include two additional conductors connecting the RTS and CTS pins between devices to enable them to signal each other with voltage-level states declaring when they are ready to receive more data:



Improper flow control configuration is a common commissioning problem in new serial networks. If the devices are configured for hardware handshaking (i.e. necessitating RTS and CTS lines in the connecting cables) or those devices simply default to hardware handshaking as new, but cables lacking RTS-to-CTS lines are used between these devices, the devices will never communicate because their respective (floating) CTS inputs will remain in the idle state and therefore the devices “think” they do not have permission to send data. *Hardware flow control necessitates serial cables with at least five conductors, while software flow control needs only three (TD, RD, and Ground).*

If one desires to directly connect two DTE devices together using EIA/TIA-232, a special cable called a *null modem* must be used, which swaps the connections between pins 2 and 3 of each device. A “null modem” connection is necessary for the transmit pin of each DTE device to connect to the receive pin of the other DTE device:



The concept of a “null modem” is not unique to EIA/TIA-232 circuits⁴⁴. Any communications standard where the devices have separate “transmit” and “receive” channels will require a “null modem” connection with transmit and receive channels swapped to be able to communicate directly without the benefit of interconnecting DCE devices. Four-wire EIA/TIA-485 and Ethernet over twisted-pair wiring are two other examples of digital communication standards where a “null” style cable is required for two DTE devices to directly connect.

EIA/TIA-232 networks may be simple, but they tend to be rather limited both in data bit rate and distance, those two parameters being inversely related. References to the EIA/TIA-232 standard repeatedly cite a maximum data rate of 19.2 kbps at 50 feet cable rate. Experimental tests⁴⁵ suggest greater rate/distance combinations may be possible in optimum conditions (low cable capacitance, minimum noise, good grounding). Since EIA/TIA-232 was developed to connect peripheral devices to computers (typically within the physical span of one room), and at modest speeds, neither of these limitations were significant to its intended application.

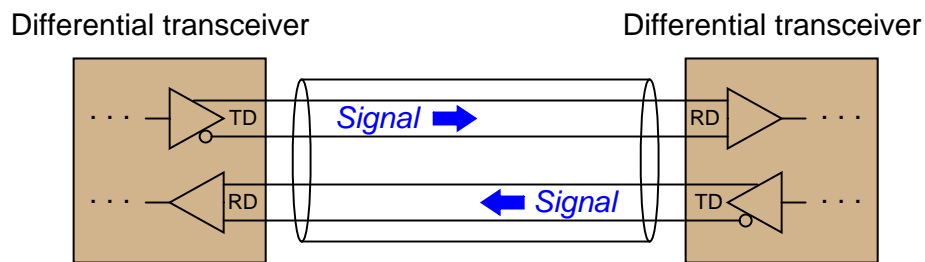
⁴⁴In fact, the concept is not unique to digital systems at all. Try talking to someone using a telephone handset held upside-down, with the speaker near your mouth and the microphone near your ear, and you will immediately understand the necessity of having “transmit” and “receive” channels swapped from one end of a network to the other!

⁴⁵Once I experimented with the fastest data rate I could “push” an EIA/TIA-232 network to, using a “flat” (untwisted, unshielded pair) cable less than ten feet long, and it was 192 kbps with occasional data corruptions. Park, Mackay, and Wright, in their book *Practical Data Communications for Instrumentation and Control* document cable lengths as long as 20 meters at 115 kbps for EIA/TIA-232, and 50 meters (over 150 feet!) at 19.2 kbps: over three times better than the advertised EIA/TIA-232 standard.

15.6.2 EIA/TIA-422 and EIA/TIA-485

The next two network standards⁴⁶ are less comprehensive than EIA/TIA-232, specifying only the electrical characteristics of signaling without any regard for connector types or any layer 2 (handshaking) considerations. Within these domains, the 422 and 485 standards differ significantly from 232, their designs intended to optimize both maximum cable length and maximum data rate.


To begin with, the electrical signaling used for both EIA/TIA-422 and EIA/TIA-485 is *differential* rather than single-ended (*balanced* rather than unbalanced). This means a dedicated *pair* of wires is used for each communications channel rather than a single wire whose voltage is referenced to a common ground point as is the case with EIA/TIA-232:



Using dedicated wire pairs instead of single conductors sharing a common ground means that EIA/TIA-422 and EIA/TIA-485 networks enjoy much greater immunity to induced noise than EIA/TIA-232. Noise induced via electrostatic coupling along the length of a network cable tends to be fairly equal (i.e. common-mode) on all non-grounded conductors of that cable, but since the receivers in EIA/TIA-422 and EIA/TIA-485 networks respond only to differential voltages (not common-mode voltages), induced noise is ignored.


⁴⁶Former labels for EIA/TIA-422 and EIA/TIA-485 were RS-422 and RS-485, respectively. These older labels persist even today, to the extent that some people will not recognize what you are referring to if you say “EIA/TIA-422” or “EIA/TIA-485.”

The advantage differential signaling enjoys over single-ended signaling may be understood by graphical comparison. The first illustration shows how electrical noise imposed on the ungrounded conductor of a simplex communications cable becomes superimposed on the digital data signal, detected at the receiving end. Noise is modeled here as a voltage source in series along the ungrounded conductor, near the receiving end. In reality, it is more likely to be distributed along the bulk of the cable length:



If the superimposed noise voltage detected at the receiver has sufficient peak-to-peak amplitude to push the signal voltage above or below critical threshold levels, the receiver will interpret this as a change of digital state and cause corruptions in the data stream.

By contrast, any noise superimposed on ungrounded conductors in a differential signaling circuit cancel at the receiver, because the close proximity of those two conductors ensures any induced noise will be the same. Since the receiver responds only to *differential* voltage between its two inputs, this common-mode noise cancels, revealing a “clean” data signal at the end:



Both EIA/TIA-422 and EIA/TIA-485 systems use differential signaling, allowing them to operate over much longer cable lengths at much greater cable speeds than EIA/TIA-232 which is single-ended. Other high-speed network standards including Ethernet and USB (Universal Serial Bus) use differential signaling as well.


EIA/TIA-422 is a *simplex* (one-way) communications standard, whereas EIA/TIA-485 is a *duplex* (two-way) standard. Both support more than two devices on a network segment. With EIA/TIA-422, this means one transmitter and multiple receivers. With EIA/TIA-485, this may include multiple *transceivers* (devices capable of both transmitting and receiving at different times: half-duplex). Four wires are necessary to connect two such devices when full-duplex (simultaneous two-way communication) is required, and full-duplex is only practical between two devices (as shown in the previous illustration).

EIA/TIA-422 and EIA/TIA-485 specify positive and negative voltage differences (measured between each dedicated wire pair) for its signaling, both for transmitting devices as well as receiving devices. A receiver must recognize any signal more negative than -200 millivolts as a “mark” (1) and any signal more positive than $+200$ millivolts as a “space” (0). These voltage thresholds are much lower than those specified for EIA/TIA-232 (± 3 volts) due to the relative lack of noise on differential signal lines compared to ground-referenced signal lines. Simply put, less noise voltage on the lines means the signal doesn’t have to be as strong to “swamp” that noise and be reliably detected at the receiver. EIA/TIA-422 transmitters (“drivers”) are supposed to generate -2 and $+2$ volt signals (minimum amplitude) to ensure at least 1.8 volts of noise margin between transmitter and receiver. EIA/TIA-485 drivers are allowed a smaller noise margin, with the minimum signal levels being -1.5 volts and $+1.5$ volts.

The maximum recommended cable length for both EIA/TIA-422 and EIA/TIA-485 networks is 1200 meters, which is greater than half a mile⁴⁷. The maximum data rate is inversely dependent on cable length (just as it is for EIA/TIA-232), but substantially greater owing to the noise immunity of differential signaling. With the long cable lengths and higher data rates made possible by differential signaling, some applications may require *terminating resistors* to eliminate reflected signals. Experiments conducted by Texas Instruments demonstrate acceptable signal integrity at 200 kbps over a cable 100 feet long with no termination resistors. With a termination resistor at the receiver input (for simplex data transmission) in place on the same 100 foot cable, a data rate of 1 Mbps was achieved.

⁴⁷1200 meters is the figure commonly cited in technical literature. However, Park, Mackay, and Wright, in their book *Practical Data Communications for Instrumentation and Control* document EIA/TIA-422 and EIA/TIA-485 networks operating with cable lengths up to 5 km (over 16000 feet!) at data rates of 1200 bps. Undoubtedly, such systems were installed with care, using high-quality cable and good wiring practices to minimize cable capacitance and noise.

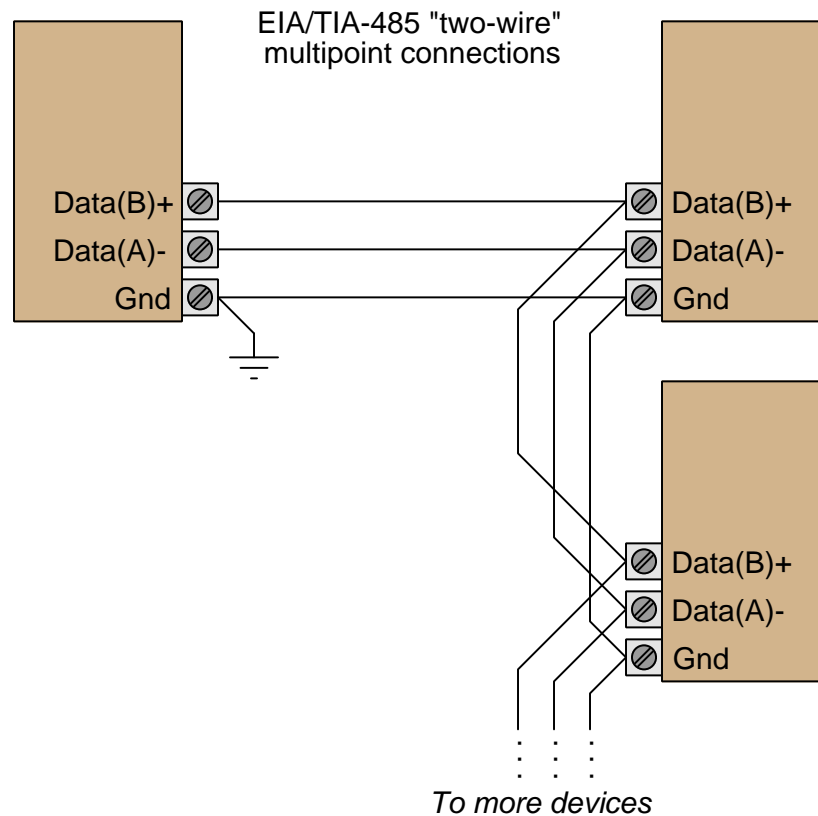
Due to the lack of standardization for cable connectors in EIA/TIA-422 and EIA/TIA-485 networks, there are no established pin numbers or labels for the differential transmit and receive conductors. A common convention seen in industrial devices, though, are the labels “A” and “B”, alternative labeled “-” and “+” or “A-” and “B+” in honor of their idle-state polarities (the “mark” or “1” state). In a 4-wire EIA/TIA-485 network, where full-duplex operation is possible, the terminals and connections will look something like this:



The good news with regard to 422/485 terminal labeling is that you will not harm the electronics by accidentally connecting the wires with incorrect polarity. If, for example, you cannot get a 422/485 receiver to acknowledge data sent by a 422/485 transmitter, you may try swapping the polarity (A/B, or +/- connections) without risking harm to the device and see if that fixes the problem.

Note the use of a ground conductor connecting both devices together. Even though the data signaling is differential and therefore does not theoretically require a common ground connection (since common-mode voltage is ignored), a ground connection helps ensure the common-mode voltage does not become excessive, since *real* receiver circuits have practical limits on the amount of common-mode voltage they can tolerate.

A popular connection scheme for EIA/TIA-485 half-duplex operation is where the Transmitted Data (TD) and Received Data (RD) terminal pairs are combined, so that two-way communication may occur over one pair of wires. With such devices, it is customary to label the terminals simply as “Data” (A– and B+):




The possibility of half-duplex operation raises the question of channel arbitration and device addressing, but since the EIA/TIA-485 standard does not specify anything outside of layer 1 concerns, these matters are left to other networking standards to fulfill. In other words, EIA/TIA-485 is not a complete data communications standard, but merely serves as the layer 1 component of other standards such as Allen-Bradley’s *Data Highway* (DH), Opto 22’s *Optomux*, and others.

Given the potential for high-speed communication along lengthy runs of cable using EIA/TIA-422 or EIA/TIA-485, the potential necessity of terminating resistors to prevent signal “reflection” is very real. Networks operating with short cables, and/or slow data rates, may work just fine without termination resistors⁴⁸. However, the effects of reflected signals grows more pronounced as the reflection time (time-of-flight for the signal to travel “round-trip” from one end of the cable to the other and back) approaches a substantial fraction of the bit time.

⁴⁸In fact, a great many EIA/TIA-485 networks in industry operate “unterminated” with no problems at all.

No network should have more than two termination resistors, one at each (far) end, and care should be taken to limit the lengths of all cable “stubs” or “spurs” branching off of the main “trunk” cable:

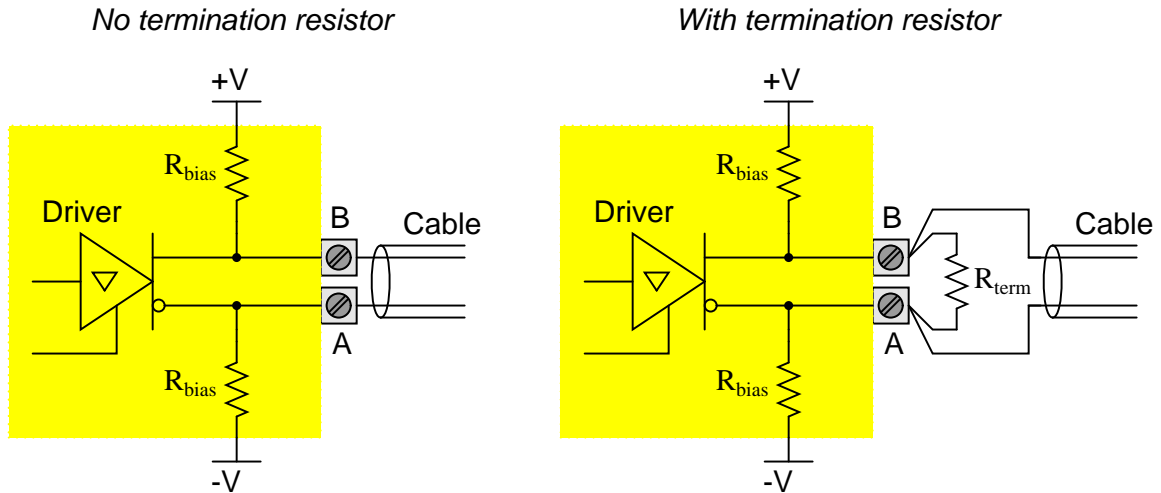


The proper value for these resistors, of course, is equality with the characteristic impedance⁴⁹ of the cable itself. A termination resistor value greater than the cable’s surge impedance will still allow positive reflections of limited amplitude, while a termination resistor value less than the cable’s surge impedance will still allow negative reflections of limited amplitude.

However, the inclusion of resistive loads to an EIA/TIA-422 or EIA/TIA-485 network may cause other problems. Many devices use a pair of internal *biasing resistors* to establish the “mark” state necessary for idle conditions, connecting the “A” terminal to the negative supply voltage rail through a resistor and the “B” terminal to the positive supply voltage rail through another resistor. Connecting a terminating resistor between terminals “A” and “B” will alter the voltage levels normally provided by these biasing resistors, consequently causing problems.

⁴⁹For detailed explanation of how and why this is necessary, refer to section 5.10 beginning on page 475.

The following schematic diagram shows the equivalent circuit of an EIA/TIA-485 transceiver device, with and without a terminating resistor connected:



When the driver is in high-impedance (High-Z) mode, the “idle” state of the wire pair will be established by the bias resistors (equal to the supply voltage so long as there is no loading). However, a terminating resistor will act as a DC load to this biasing network, causing a substantial reduction of the “idle” state voltage toward 0 volts. Recall that -200 millivolts was the receiving threshold value for a “mark” state in both EIA/TIA-422 and EIA/TIA-485 standards (terminal “A” negative and terminal “B” positive). If the presence of a terminating resistor⁵⁰ reduces the idle state voltage to less than 200 millivolts absolute, the receiver(s) will not be able to reliably read the network’s idle state and communication errors will result.

Thus, we see that the inclusion of any terminating resistors must be accompanied by an analysis of the devices’ bias resistor networks if we are to ensure robust network operation. It is foolhardy to simply attach terminating resistors to an EIA/TIA-422 or EIA/TIA-485 network without considering their combined effect on biasing.

⁵⁰Actually *two* terminating resistors in parallel, since one will be at each end of the cable! The actual DC biasing network will be more complicated as well if more than one device has its own set of internal bias resistors.

15.7 Ethernet networks

An engineer named Bob Metcalfe conceived the idea of Ethernet in 1973, while working for the Xerox research center in Palo Alto, California. His fundamental invention was the CSMA/CD method of channel arbitration, allowing multiple devices to share a common channel of communication while recovering gracefully from inevitable “collisions.” In Metcalfe’s vision, all of the “network intelligence” would be built directly into “controller” devices situated between the DTE devices (computers, terminals, printers, etc.) and a completely passive coaxial cable network. Unlike some other networks in operation at the time, Metcalfe’s did not rely on additional devices to help coordinate communications between DTE devices. The coaxial cable linking DTE devices together would be completely passive and “dumb,” performing no task but the conduction of broadcast signals between all devices. In that sense, it served the same purpose as the “luminiferous ether” once believed to fill empty space: conducting electromagnetic waves between separated points.

The CSMA/CD (“Carrier Sense Multiple Access with Collision Detection”) method of bus arbitration works by giving each Ethernet device the ability to sense an idle channel as well as sense if it happens to “collide” with another transmitting device. In the event of a collision, the colliding devices both cease transmission, and set random time-delays to wait before re-transmission. The individual time delays are randomized to decrease the probability that a re-collision between the same devices will occur after the wait. This strategy is analogous to several peers in one group holding a conversation, where all people involved are equally free to begin speaking, and equally deferential to their peers if ever two or more accidentally begin speaking at the same time. Occasional collisions are perfectly normal in an Ethernet network, and should not be taken as an indication of trouble unless their frequency becomes severe.

Metcalfe’s original network design operated at a data rate of 2.94 Mbps, impressive for its time. By 1980, the three American computer companies DEC (Digital Equipment Corporation), Intel, and Xerox had collaborated to revise the Ethernet design to a speed of 10 Mbps, and released a standard called the *DIX Ethernet* standard (the acronym “DIX” representing the first letter of each company’s name). Later, the IEEE Local and Metropolitan Networks Standards Committee codified the DIX Ethernet standard under the numeric label 802.3. At the present time there exist many “supplemental” standards underneath the basic 802.3 definition, a few of them listed here:


- 802.3a-1985 *10BASE2 “thin” Ethernet*
- 802.3d-1987 *FOIRL fiber-optic link*
- 802.3i-1990 *10BASE-T twisted-pair cable Ethernet*
- 802.3u-1995 *100BASE-T “Fast” Ethernet and Auto-Negotiation*
- 802.3x-1997 *Full-Duplex standard*
- 802.3ab-1999 *1000BASE-T “Gigabit” Ethernet over twisted-pair cable*

The IEEE 802.3 standard is limited to layers 1 and 2 of the OSI Reference Model: the “Physical” and “Data link” layers. In the physical layer (1), the various supplements describe all the different ways in which bits are electrically or optically represented, as well as permissible cable and connector types. In the data link layer (2), the IEEE standard describes how devices are addressed (each one with a unique identifier known as a *MAC address*, consisting of a 48-bit binary number usually

divided into six bytes, each byte written as a two-character hexadecimal number), the CSMA/CD channel arbitration protocol, and also how data frames are organized for Ethernet transmissions.

15.7.1 Repeaters (hubs)


Bob Metcalfe's original design for Ethernet consisted of DTE devices connected to a common coaxial cable through the use of "tee" connectors, like this:



This cabling arrangement suffered several problems. First, it was inconvenient to run through an office building, since each DTE device needed to be coupled rather closely to the main "trunk." Short cable segments (called *stubs*, *spurs*, or *drops*) joining the main trunk line to each DTE device could not be too long, or else they could cause multiple signal reflections to occur in the main line. Secondly, the signal strength decreased with each "tee" connector: every time the signal branched, it would lose power. Thirdly, the need for termination resistors at the far ends of the "ether" cable opened up the possibility that those terminators might fail, fall off, or be forgotten during installation or maintenance⁵¹.

⁵¹These very same problems may arise in FOUNDATION Fieldbus networks, for the exact same reason: the cabling is passive (for increased reliability). This makes FOUNDATION Fieldbus instrument systems challenging to properly install for most applications (except in really simple cases where the cable route is straightforward), which in my mind is its single greatest weakness at the time of this writing (2009). I strongly suspect Ethernet's history will repeat itself in FOUNDATION Fieldbus at some later date: a system of reliable "hub" devices will be introduced so that these problems may be averted, and installations made much simpler.


As Ethernet evolved as a practical networking standard, one of the many improvements added to its design was the concept of a *repeating hub*. A “repeater” is an active device designed to re-broadcast a signal, usually to overcome inevitable power losses incurred as that signal propagates along a cable. Repeaters are common in the telecommunications industry, where telephone, television, and computer signals must travel hundreds or thousands of miles between points of transmission and reception. A “repeating hub” is a repeater with multiple ports for many cables to plug into, where any signal entering on any cable is repeated to *all* ports on the device. Thus, a repeating hub (or simply “hub”) allows multiple Ethernet devices to interconnect with no degradation in signal quality:



Not only do hubs improve system performance by boosting signals’ voltage levels, but they also eliminate the need for termination resistors in the network. With a hub-based system, each and every cable terminates at either a DTE or DCE device, which is (now) designed with the proper termination resistance built-in to their internal transceiver circuitry. This means each and every Ethernet cable is automatically terminated with the proper impedance simply by plugging it in to the Ethernet port of *any* device. “Stub” or “spur” cables with their length restrictions are also a thing of the past, since no cable ever splits or branches in a hub-based network system.

Hubs are considered “layer 1” devices, because they operate purely on the physical layer of Ethernet: all they do is receive Ethernet signals and re-broadcast those signals in boosted form to all other devices plugged into the hub. As a piece of interconnecting hardware, a hub is considered a DCE (Data Communications Equipment), as opposed to the end-of-cable devices such as computers and printers which are DTEs (Data Terminal Equipment).

Repeating hubs may be connected together to form larger networks⁵²:




Since hubs are merely “layer 1” devices, mindlessly boosting and re-broadcasting signals received to their ports, their presence does not mitigate collisions between transmitting devices. As far as collisions between those devices is concerned, they might as well be directly connected together on a single piece of coaxial cable. One way to express this concept is to say that all portions of the network are part of the same *collision domain*. In other words, any devices on this network are able to collide with each other, because all transmissions are sensed by all the devices. This is analogous to a small room with several people in it: the room is small enough that everyone can hear everyone else talking, which means only one person in that room is able to speak at a time.

⁵²There are practical limits as to how many hubs may be “daisy-chained” together in this manner, just as there are practical limits to how long a twisted-pair cable may be (up to 100 meters). If too many hubs are cascaded, the inevitable time delays caused by the process of repeating those electrical impulses will cause problems in the network. Also, I have neglected to specify the use of *crossover* cables to connect hubs to other hubs – this is a topic to be covered later in this book!

15.7.2 Ethernet cabling

Along with hubs came another form of Ethernet cable and connector: *unshielded, twisted pair* (UTP) wiring and *RJ-45* “flat” connectors. These cables use multiple twisted pairs of wires instead of the coaxial cable specified in Metcalfe’s original Ethernet. The purpose of using twisted-wire pairs is to reduce magnetic signal coupling (for more information, refer to section 8.3.6 beginning on page 625).



For 10 Mbps Ethernet over UTP cable (called 10BASE-T) and for 100 Mbps Ethernet (called 100BASE-TX), only two⁵³ out of four available wire pairs are used:

Pin number	Assignment	Abbreviation
1	Transmit Data (+)	TD+
2	Transmit Data (–)	TD–
3	Receive Data (+)	RD+
4	(not used)	
5	(not used)	
6	Receive Data (–)	RD–
7	(not used)	
8	(not used)	


It should be noted that 1000 Mbps (“Gigabit”) Ethernet over twisted-wire pairs does in fact use all four pairs in an eight-wire cable, a departure from traditional UTP Ethernet cable wiring:

Pin number	Assignment	Abbreviation
1	Pair “A” (+)	BLDA+
2	Pair “A” (–)	BLDA–
3	Pair “B” (+)	BLDB+
4	Pair “C” (+)	BLDC+
5	Pair “C” (–)	BLDC–
6	Pair “B” (–)	BLDB–
7	Pair “D” (+)	BLDD+
8	Pair “D” (–)	BLDD–

⁵³With only half the available wire pairs used in a standard 10 Mbps or 100 Mbps Ethernet cable, this opens the possibility of routing *two* Ethernet channels over a single four-pair UTP cable and RJ-45 connector. Although this is non-standard wiring, it may be a useful way to “squeeze” more use out of existing cables in certain applications. In fact, “splitter” devices are sold to allow two RJ-45-tipped cables to be plugged into a single RJ-45 socket such that one four-pair cable will then support two Ethernet pathways.


Along with UTP cables and RJ-45 connectors came a significant alteration to the basic electrical scheme of Ethernet. Metcalfe's original design used a simple coaxial cable as the "ether" connecting devices together. Such cables had only two conductors, meaning each device needed to transmit *and* receive data over the same two conductors. With UTP cable's four pairs of conductors, transmission and reception of signals occurs over different wire pairs⁵⁴. This means connections made between Ethernet devices must employ a "swap" between TD and RD wire pairs in order for communication to take place, so that the "receiver" circuitry of one device connects to the "transmitter" circuitry of the other, and vice-versa. This is precisely the same characteristic inherent to EIA/TIA-232 and four-wire EIA/TIA-485 networks, where separate wire pairs are dedicated to "transmit" and "receive" functions.

In a typical Ethernet system, the interconnecting hubs perform this transmit/receive swap. Hubs are considered DCE devices, while computers and other end-of-the-line devices are considered DTE devices. This means the pin assignments of DTE and DCE devices must be different in order to ensure the transmit/receive pin swap necessary for straight-through cables to work. This also means if someone ever wishes to directly connect two Ethernet DTE devices together without the benefit of a hub in between, a special *crossover* cable must be used for the connection, identical in function to the *null modem* cable used to connect two EIA/TIA-232 DTE devices together:




⁵⁴This means modern Ethernet is capable of full-duplex communication between two devices, whereas the original coaxial-based Ethernet was only capable of half-duplex communication.

Furthermore, the same problem exists when multiple hubs are connected to form larger networks. Since each hub is a DCE device, a straight-through cable connecting two hubs together will pass transmitted signals from one hub directly to the “transmit” pins of the other hub, not the “receive” pins as it needs to. Consequently, a “crossover” cable should be used to connect two Ethernet hubs together in order to avoid this problem:



Some early Ethernet hubs provided a different solution to the “crossover” problem, and that was a crossover *switch* built into the hub, allowing a person to manually switch the transmit and receive wire pairs with the push of a button. In this next photograph of a four-port Ethernet hub, you can see the “Normal/Uplink” pushbutton on the right-hand side of the front panel, controlling the furthest-right port of the hub. This switch is supposed to be placed in the “Normal” position if the device plugged into that port is a DTE device, and placed in the “Uplink” position if the device is a DCE device (e.g. another hub):




Note the LED indicator lights by each port on the hub. One LED indicates whether or not the cable is active (when a powered Ethernet DTE device is plugged into that port of the hub), while the other LED indicates traffic on the cable (by blinking). These LEDs are very helpful for identifying a crossover problem. This hub even has an LED indicating the occurrence of collisions (the “Col” LED just below the main power LED), giving simple visual indication of collision frequency.

Newer Ethernet DTE and DCE devices use auto-sensing technology to perform any necessary transmit/receive pin swaps, rendering crossover cables and crossover pushbuttons unnecessary for either DTE-to-DTE or hub-to-hub connections. Auto-sensing is a standard feature of 1000BASE-T (“Gigabit” Ethernet).

15.7.3 Switching hubs


The next evolutionary step in Ethernet network connections is the introduction of a *switching hub*, or simply *switch*. A “switch” looks exactly like a repeating hub, but it contains intelligence to route transmitted signals only to specific ports, rather than broadcasting every received data frame to all ports. What enables this to happen is the information contained in each Ethernet frame transmitted by DTE devices:



Note that part of the frame includes both a source address and a destination address. These refer to the 48-bit “MAC” addresses uniquely identifying each and every Ethernet device. A switching hub “learns” the identities of all devices plugged into each of its ports by remembering the “source” addresses received through those ports. When a switch receives an Ethernet frame with a destination address it recognizes as residing on one of its ports, it *only* repeats that frame to that specific port, and not to the other ports. In other words, an Ethernet switch does not mindlessly broadcast all messages to all of its ports the way an Ethernet hub does. The switch’s targeted direction of messages reduces the amount of “traffic” seen at the other ports, and also avoids unnecessary collisions because messages only get sent to their intended destinations.

If a switch receives a data frame with an unrecognized destination address, it defaults to basic “hub” behavior by broadcasting that frame to all ports. If a device plugged into one of that switch’s ports replies to that data frame, the MAC address of that device is noted for future traffic direction to that port.

The presence of a switching hub in a larger network has the effect of dividing that network into separate collision domains, so that a collision occurring in one domain does not “spill over” into another domain where it would delay communication between those devices:



Of course, collisions between these two domains may still occur, for instance if a device in the first domain tries to transmit to a device in the second domain at the exact same time that a device in the second domain attempts to transmit to a device in the first.

With this added intelligence, switching hubs are considered “layer 2” devices, since they operate not just at the physical layer of electrical impulses, but also at the next layer of device addressing. Since switching hubs add benefit beyond repeating hubs without any drawbacks⁵⁵, most people elect to use switches whenever possible.

⁵⁵Even the cost difference is negligible. It should be noted, though, that switches may exhibit unintended behavior if a cable is unplugged from one of the ports and re-plugged into a different port. Since switches internally map ports to device addresses, swapping a device from one port to another will “confuse” the switch until it re-initializes the port identities. Re-initialization may be forced by cycling power to the switch, if the switch does not do so on its own.

15.8 Internet Protocol (IP)

I remember first learning about the world-wide Internet, and wondering what it actually *looked like*. The first vision entering my mind when people told me about a computer network spanning nearly all of the United States and many other parts of the world was that of a thick cable strung along telephone poles and buried underground, with a big sign on it saying “Internet.” I also remember well the shock of learning that although the Internet made use of several high-capacity networks (called *backbones*) connecting large data centers in different cities, the real “magic” of the Internet did not reside in any particular cable or link. Instead, what made the Internet so widespread and accessible was actually a *protocol* allowing for the free exchange of data along and between disparate systems. This “protocol” allowed digital data to be packaged in such a way that it could be sent along nearly any kind of communications link (from copper wires to fiber-optic to radio waves) – and indeed along multiple pathways between the same two points – while arriving at the destination intact. Thus, the Internet was akin to a random patchwork of existing communications pathways pressed into coordinated service by the sharing of a common “language.” In this section, we will investigate the protocol at the heart of the Internet, appropriately called *Internet Protocol*, or *IP*.

Physical network standards such as Ethernet only define aspects relevant to lower layers of the OSI Reference Model. While these details are essential for communication to occur, they are not enough on their own to support a wide-spread communications system. For this reason, network standards such as EIA/TIA-485 and Ethernet almost always comprise the lower layer(s) of a more complex communications protocol capable of managing higher-order addresses, message integrity, “sessions” between computers, and a host of other details.

Internet Protocol (IP) manages network addresses and data handling over a much larger physical domain than Ethernet is able to. The basic principle of IP is that large digital messages may be broken down into smaller pieces, then each piece buffered with additional data bits to form *packets* specifying (among other things) how the pieces are to be directed to their proper destination(s). The completed packets are then transmitted individually and received individually, where they may be reassembled at the receiver to form the original message in its entirety. An analogy for this process is an author with a printed paper manuscript for a book, who needs to get her manuscript to a print shop across town. Unfortunately, the mail service in this town cannot transport the bulky manuscript in one piece, so the author divides the manuscript into 10-page bundles and mails each of these bundles in its own package to the print shop. The individual packages may not make it to the print shop on the same day, or even in the correct order, but the addressing on each package directs the postal service to deliver each of them to the proper location.

This strategy for transmitting large digital messages is at the heart of the Internet: data sent from one computer to another over the Internet is first broken down into packets, which are then routed over a variety of pathways to their destination. The packets need not take the same route to their destination, nor do they even need to travel along the same *kinds* of networks. The receiving computer must then reassemble those packets in the proper order to re-create the original data. This “packetization” of data allows multiple messages to be interleaved on a network (i.e. the network’s bandwidth being alternately used to convey pieces of completely different messages, rather than being reserved for one whole message at a time) as well as permitting alternate routes that the message may take in order to traverse large physical distances. In a web-shaped network where multiple pathways exist between any two points, the ability to direct packets of data along alternate routes increases the reliability of that network: failure of any one routing node or communications pathway does not necessarily prevent data from reaching its final destination. This fault tolerance

was one of the design criteria for what came to be the Internet when it was first developed by the United States' Department of Defense.

Interestingly, the task of portioning a large block of digital data into packet-sized pieces, and then re-assembling those pieces together in the proper order to form the original data block, is *not* the task of IP, but rather the task of some higher-level protocol such as TCP (Transmission Control Protocol). Internet Protocol (IP) merely specifies how the individual packets are to be marked and routed to their proper destination(s)⁵⁶. To use the manuscript analogy again, IP is the postal service with its system of mailing addresses, postage stamps, and labeling conventions, while TCP (or some other higher-level protocol) is the author and publisher who divide the manuscript into smaller bundles and then reassemble those bundles into the original manuscript, respectively. For this reason, IP is not a complete solution for large-scale network communication on its own. This is why the Internet's central protocol is referred to as TCP/IP, the two protocols working together to ensure coordinated and reliable communication of packetized data over wide areas.

⁵⁶When packets travel between different kinds of networks, the “gateway” devices at those transition points may need to *fragment* large IP packets into smaller IP packets and then re-assemble those fragments at the other end. This fragmentation and reassembly *is* a function of Internet Protocol, but it happens at the packet level. The task of portioning a large data block into packet-sized pieces at the very start and then reassembling those packets into a facsimile of the original data at the very end, however, is beyond the scope of IP.

15.8.1 IP addresses

IP is a “layer 3” technology, being concerned with network-wide addresses for routing information between two different locations. IP is not concerned with the details of communication along any particular wire or fiber-optic cable. It is not “aware” of how bits are represented electrically, or what kind of connectors are used to couple cables together. IP is only concerned with “networks” in the broad sense of the word, as abstract collections of computers that are *somehow* (it doesn’t care exactly how) connected to each other.

Networking equipment (DCE) designed to pay attention to IP addresses for routing purposes are called, not surprisingly, *routers*. Their purpose is to direct packets to their appropriate destinations in the shortest amount of time.


In order for the Internet Protocol to specify where packets are coming from and where they are going to, each source and destination must be marked with its own *IP address*. IP version 4 (IPv4) uses 32-bit addresses, usually expressed as four octets (four bytes) written using decimal numbers. For example:

IP address 00000000 00000000 00000000 00000000 is written as 0.0.0.0

IP address 11111111 11111111 11111111 11111111 is written as 255.255.255.255

IP address 10101001 11111010 00101101 00000011 is written as 169.250.45.3

In order for two inter-connected computers to exchange data using Internet Protocol, each one must have a unique IP address:



At first, this may seem redundant. Doesn’t each and every Ethernet device already have its own unique “MAC address” 48 bits in length to distinguish it from every other Ethernet device in existence? If so, why add *another* set of identifying addresses to the system?

This is true – Ethernet devices are already uniquely addressed – but those MAC addresses serve different purposes than IP addresses. Recall that Ethernet is a standard only at layers 1 and 2, and is not “aware” of any higher-level concerns. Ethernet MAC addresses are useful to switching hubs and other Ethernet DCE devices tasked with management of Ethernet data frames, but those MAC addresses – unique as they may be – have little relevance in the greater picture of IP where we must fragment and reassemble messages over very large-scale networks. More importantly, the reason we need IP addresses is to be able to use interconnecting networks other than Ethernet. For example, two computers may be connected to each other with a simple EIA/TIA-232 cable (or even using radio transceiver units for a “wireless” connection) instead of Ethernet, but still use Internet

Protocol to route packets to their destinations⁵⁷. By having its own dedicated addressing scheme, IP ensures computers can send and receive data packets with no regard to *physical interconnection details, channel arbitration methods, or anything else in between*. In a sense, IP is the “glue” that holds disparate networks together, and makes something like a world-wide Internet possible when so many different network types exist to connect digital devices together. If we attempted to use Ethernet MAC addresses for the same purpose, *the entire Internet would have to consist solely of Ethernet networks!*

A helpful analogy is to think of Ethernet MAC addresses like Social Security numbers for United States citizens, while IP addresses are like street addresses used to route mail. Each US citizen should have their own unique Social Security number, shared by no one else. This number is used for many purposes, including identification on Federal tax documents, to help route specific information (such as income records and Social Security payments) to the proper people. Despite the uniqueness of these numbers, though, people still need separate mailing addresses in order to receive mail through the postal service and other package distribution agencies. The mailing address serves a different purpose than the Social Security “address” each US citizen possesses. Furthermore, the existence of separate mailing addresses ensures even non-citizens living in the United States (e.g. foreign students, ambassadors, etc.) who have no Social Security numbers still have a way to send and receive mail. The mapping of device MAC addresses to IP addresses is handled by a protocol called *ARP* (Address Resolution Protocol) discussed later in this chapter.

The “ping” utility

Computers enabled to communicate using Internet Protocol (IP) are equipped with a utility program named `ping` useful for detecting the presence of other IP-enabled computers connected to the same network. The classic format of this program is execution by typing the word “ping” at the computer’s command-line interface followed by the IP address of the other computer you wish to detect the presence of. For example, if I wished to check for the presence of a computer on the network with an IP address of 133.82.201.5, I would type this command at my computer’s command line and press the “Enter” key:

```
ping 133.82.201.5
```

The `ping` utility works by sending a very short digital message⁵⁸ to the specified IP address, requesting a reply from that computer (usually with multiple attempts). The `ping` command as implemented on the Microsoft Windows (XP) operating system typically makes four attempts before quitting. Some other operating systems’ implementation of `ping` continue indefinitely until halted by the user with the “Control-C” keystroke interrupt combination.

⁵⁷In fact, this is precisely the state of affairs if you use a *dial-up* telephone connection to link your personal computer with the Internet. If you use dial-up, your PC may not use Ethernet at all to make the connection to your telephone provider’s network, but rather it might use EIA/TIA-232 or USB to a modem (modulator/demodulator) device, which turns those bits into modulated waveforms transmittable over a voice-quality analog telephone line.

⁵⁸The “ping” command is technically defined as an “Echo Request” command, which is part of the *Internet Control Message Protocol* (ICMP) suite.

When diagnosing problems with IP-enabled network devices, few utilities are as immediately useful as `ping`. Networking professionals commonly use the word “ping” as a verb, as in “I tried to *ping* that computer, but it gave no response.” There are many reasons why a computer might fail to respond to a `ping` query, but a successful `ping` attempt proves several things:

- The destination device is powered up and its IP functionality is working
- All network devices (DCE) between your computer and the destination device are communicating
- All cables necessary for the communication of data between your computer and the destination are functional
- Both your computer and the destination device are on the same *subnet* (this topic covered in more detail later)

Since `ping` requires the first three layers of the OSI model to properly function (Physical, Data Link, and Network layers), using this as a diagnostic test neatly identifies where in the OSI model a problem exists. If two computers are not communicating with each other as they should but the `ping` utility works between them, the communication fault must lie within one of the upper OSI layers (e.g. Transport, Session, Presentation, or Application). Thus, we see the `ping` utility as a tool for “divide-and-conquer” style troubleshooting, where we may prove good connections between certain devices and thereby narrow the scope of the problem by elimination.

IPv4 address ranges

Given the addressing purpose of Internet Protocol (to designate addresses over an extremely large collection of digital communication devices), addresses must be chosen with care. IP version 4 uses a 32-bit field to designate addresses, limiting its address capacity to 2^{32} unique addresses. As large as this number is, it is not enough to uniquely identify all Internet-capable devices worldwide. The inventors of IP did not dream their Internet would grow to the proportions it has today. Let this be a lesson to all those involved with computers: the future is usually bigger than you think! A variety of clever techniques has been developed to deal with this shortage of IP addresses. One of them is to dynamically assign addresses to Internet-connected computers *only when they are turned on*. This is how most personal Internet connections work: when you power up your personal computer to connect to the Internet, your service provider assigns you a temporary IP address through a protocol called DHCP (Dynamic Host Configuration Protocol). Your provider then forces you to relinquish this temporary IP address when you shut down your computer, so someone else may use it for theirs.

The *Internet Corporation for Assigned Names and Numbers*, or *ICANN*, is the organization responsible⁵⁹ for assigning IP addresses to Internet users worldwide (among other tasks). This group has designated certain IP address ranges specific to internal (i.e. *Local Area Network*, or *LAN*) network devices, which shall never be used “publicly” to address devices on the world-wide Internet. These specially-designated “private” LAN address ranges are as follows:

10.0.0.0 to 10.255.255.255

172.16.0.0 to 172.31.255.255

192.168.0.0 to 192.168.255.255

Additionally, all computers have their own special *loopback* IP address, used to send IP message packets to itself for certain purposes (including diagnostics): 127.0.0.1. This IP address is completely *virtual*, not associated with any network hardware at all⁶⁰. Therefore, the `ping` command executed on any computer should *always* be able to detect address 127.0.0.1, regardless of the status or even existence of actual network hardware (cards or interfaces) on that computer. Failure of the `ping` command to detect the loopback address is a sign that the computer’s operating system is not configured to use Internet Protocol.

A computer’s loopback address may have uses other than diagnostic. Some computer applications are network-oriented by nature, and rely on IP addresses even if the application is performing some local function rather than a function between computers on an actual network. The *X-windows* graphic-user interface (GUI) system popularly used on UNIX operating systems is an example of this, referencing the loopback address to form a connection between client and server applications running on the same computer.

⁵⁹Prior to ICANN’s formation in 1999, the *Internet Assigned Numbers Authority*, or *IANA* was responsible for these functions. This effort was headed by a man named Jon Postel, who died in 1998.


⁶⁰The term “loopback” refers to an old trick used by network technicians to diagnose suspect serial port connections on a computer. Using a short piece of copper wire (or even a paperclip) to “jumper” pins 2 and 3 on an EIA/TIA-232 serial port, any serial data transmitted (out of pin 3) would be immediately received (in pin 2), allowing the serial data to “loop back” to the computer where it could be read. This simple test, if passed, would prove the computer’s low-level communication software and hardware was working properly and that any networking problems must lie elsewhere.

15.8.2 Subnetworks and subnet masks

IP (version 4) addresses are used in conjunction with something called *subnet masks*⁶¹ to divide IP networks into “subnetworks.” A “subnetwork” is a range of IP-addressed devices allowed to communicate with each other. You may think of the subnet mask to be a sort of “filter” used to identify IP addresses belonging to the proper range.

The subnet mask works as a bitwise filter, identifying those bits in the IP address defining the subnetwork. For example, if the subnet mask on a computer is set to 255.0.0.0 (binary 11111111 00000000 00000000 00000000), it means the first 8 bits of the IP address define the subnetwork, and thus the computer is only allowed to communicate with another computer belonging to the same subnetwork (i.e. having the same first octet in its IP address).

A set of examples showing two interconnected computers with differing IP addresses (and in some cases, different masks) illustrates how this works⁶². In the first example, two computers with IP addresses differing in the last two octets are able to communicate because their subnets are the same (169.254):



⁶¹Also called “netmasks” or simply “masks.”

⁶²These are real test cases I performed between two computers connected on a 10 Mbps Ethernet network. The error messages are those generated by the `ping` utility when communication was attempted between mis-matched computers.


We may check to see the IP addresses and subnet masks are correct by using the `ping` command. A screenshot of `ping` being used on a personal computer running the Microsoft Windows XP operating system is shown here:

```

C:\Documents and Settings\btc>ping 169.254.1.1
Pinging 169.254.1.1 with 32 bytes of data:
Reply from 169.254.1.1: bytes=32 time<1ms TTL=128
Reply from 169.254.1.1: bytes=32 time<1ms TTL=128
Reply from 169.254.1.1: bytes=32 time<1ms TTL=128
Reply from 169.254.1.1: bytes=32 time<1ms TTL=128
Ping statistics for 169.254.1.1:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 0ms, Average = 0ms
C:\Documents and Settings\btc>


```

In the next example, we see two computers with the same mask value, but with different address values in the octets designated by their masks. In other words, these two computers belong to different subnets: one to 167.254 and the other to 169.254, and as a result they are not allowed to communicate with each other using Internet Protocol. The resulting error messages generated by the `ping` utility are shown in this diagram:



In the last example, we see two computers having different mask values as well as different IP addresses. The subnet of the left-hand computer is 169.254.10 while the subnet of the right-hand computer is 169.254:

Computers with different subnet masks, on different subnets



The computer on the left may only communicate with IP addresses matching in the first three octets (169.254.10). Seeing that the destination address for the second computer does not match in its third octet, `ping` returns a “Destination host unreachable” error message when executed from the left-hand computer.

When the computer on the right attempts to communicate with (“ping”) the computer on the left, it is allowed to transmit to that computer because its mask only screens for agreement in the first two octets (169.254), which happen to match. However, the computer on the left is not allowed to transmit to the computer on the right because of its more restrictive subnet, and so `ping` running on the right-hand computer returns a “Request timed out” error message because it never receives a reply from the left-hand computer to any of its queries.

With just two computers connected by a single cable, the concept of subnetworks and masks seems useless, and indeed it is on this small of a scale. However, “subnetting” is a useful technique for managing high traffic loads on large networked systems using IP addresses, and so it is commonly seen in many local area networks (LANs) such as those found at industry and commercial sites.

While many IPv4-compliant computers designate both the IP address and the subnet mask values as sets of “dotted-decimal” numbers with each decimal (0-255) representing an “octet” of eight bits in the 32-bit IPv4 address space (e.g. IP = 169.254.5.1 and Mask = 255.255.0.0), a more modern designation for subnets is to append the device’s IP address with a forward slash character and a decimal number specifying how many bits are used to specify the subnet.

To illustrate by example, consider a computer having the IP address 169.254.5.1 and a mask value of 255.255.0.0 (thus specifying that it belongs to the subnetwork 169.254), we could alternatively state that computer’s IP address as 169.254.5.1/16. The “16” means that the first sixteen bits of its IP address define its subnet. To cite another example, a computer with an IP address of 192.168.35.100 and a subnet mask of 255.255.255.0 could have its address alternatively written as 192.168.35.100/24. The “24” tells us the first 24 bits (first three octets) of the IP address define its subnetwork: this computer belongs to the subnetwork 192.168.35.

The `ping` diagnostic utility program may be used to search for unknown IP addresses on a known subnet. This may be done by “pinging” to the *broadcast address* for that subnet: an IP address formed by the known subnet numbers, followed by all binary 1’s filling the unknown bit spaces. For example, you could use `ping` to search for devices on the subnet 156.71 (subnet mask 255.255.0.0) by using the following command:

```
ping 156.71.255.255
```

15.8.3 Routing tables

Devices on an IP-compliant network need to know how to best route IP data packets from one location to another. In the case of special *router* devices managing traffic on the Internet, optimum packet routes are determined by a number of different criteria (e.g. degree of congestion in a route, the fewest number of “hops” from one router to another, geographical distance, etc.), updated continually by sophisticated algorithms operating within the routers. Data for these optimum routes are stored in lists called *routing tables*.

Personal computers also have routing tables, which may be modified by the user. An application where you may need to modify the routing table of a personal computer is the case of enabling that computer to communicate with a brand-new device installed on the industrial network, whose subnetwork ID differs from the other devices on that network. Many network-ready industrial instruments are shipped from the manufacturer with default subnets of 192.168.1. It is quite possible that the subnet of the industrial network you intend to have the device operate on is different from this default factory-configuration. This in itself is not necessarily a problem, as IP addresses and subnet mask values of IP-enabled devices are always user-configurable. However, if the only method of configuring this new device is by communicating to it through an IP network connection, you are faced with a Catch-22: how do you communicate with it to alter its subnet, when its subnet prevents you from communicating with it?

One solution to this Catch-22 dilemma is to temporarily use a personal computer with a subnet matching the new device to configure that new device, then disconnecting the device from the personal computer after it has been configured for the new subnetwork, and subsequently plugging the device into the industrial network where it will function.

Another solution is to use one of the personal computers already residing on the industrial network to configure the new device, and specially enabling that one computer to talk to the device’s default subnet. In this way, the new device may be plugged into the industrial network, configured for a new IP address and subnet while on that network, where it will subsequently communicate with existing devices on the proper subnet. This may be done through the `route` command-line utility. At a command prompt (just like when using the `ping` command), type “route” followed by arguments telling it to add the device’s default address and subnet mask to the computer’s routing table. Supposing our new device has a default IP address of 192.168.1.10 and a default mask of 255.255.255.0, our `route` command would need to be entered as follows:

```
route add 192.168.1.10 mask 255.255.255.0
```

After issuing this command to the personal computer, it may be used to communicate with the new device to change its IP address and subnet mask values to match devices on the industrial network.

15.8.4 IP version 6

The next version of IP (version 6, or IPv6) uses 128-bit addresses, giving 2^{128} address possibilities (in excess of 3.4×10^{38}), in stark contrast to IPv4's paltry 2^{32} address space. To put this enormous quantity into perspective, there are enough IPv6 addresses to designate nearly 57 *billion* of them for each and every gram of the Earth's mass⁶³. While IPv4 addresses are typically written as four octets in decimal form (e.g. 169.254.10.5), this notation would be very cumbersome for writing IPv6 addresses. Thus, IPv6 addresses are written as a set of eight hexadecimal numbers (up to four characters per number) separated by colons, such as `4ffd:522:c441:d2:93b2:f5a:8:101f`. The phase-in of IPv6 to replace IPv4 has already started for certain portions of the Internet, but the full transition to IPv6 is expected to take many years. The IPv6 "loopback" virtual address for computers is `0:0:0:0:0:0:0:1`, or more simply⁶⁴ written as `::1`.

Note the "shorthand" notation used in the previous IPv6 addresses to eliminate extra characters: some of the 16-bit segments are truncated to less than four hexadecimal characters if the preceding (more-significant) characters are zero. Thus, you see `:522:` instead of `:0522:`, and `:d2:` instead of `:00d2:`. The loopback address of `::1` is the ultimate shorthand notation, collapsing all prior segments (which are all zero) into a pair of back-to-back colons.

IP version 6 supports subnetworks just as IPv4 does, but instead of denoting subnet masks in the same colon-delimited fashion as IPv6 addresses, IPv6 subnet masks are simply specified by the number of "1" bits beginning from the first (MSB). The rationale here is that subnet mask bits should be contiguous⁶⁵, with no zero bits separating one bits. This being the case, the subnet mask for any practical IP range may be specified as a simple number of 1's from the MSB down to the LSB⁶⁶.

It should be noted that an updated version of the `ping` command (called `ping6`) is available to help diagnose IPv6 systems.

⁶³According to Douglas Giancoli's *Physics for Scientists and Engineers* textbook, the mass of the Earth is 5.98×10^{24} kg, or 5.98×10^{27} grams. Dividing 2^{128} (the number of unique IPv6 addresses) by the Earth's mass in grams yields the number of available IPv6 address per gram of Earth mass. Furthermore, if we assume a grain of sand has a mass of about 1 milligram, and that the Earth is modeled as a very large collection of sand grains (not quite the truth, but good enough for a dramatic illustration!), we arrive at 57 *million* IPv6 addresses per grain of sand on Earth.

⁶⁴The fully-written loopback address is actually `0000:0000:0000:0000:0000:0000:0000:0001`.

⁶⁵While it is possible to use non-contiguous subnet mask values, the practice is frowned upon by most system administrators.

⁶⁶Indeed, subnet masks for IPv4 can be specified in this manner as well, not just IPv6 subnet masks.


15.8.5 ARP

While Internet Protocol (IP) provides a universal addressing standard for devices operating on large-scale digital networks, individual devices usually possess MAC addresses unique to each device. As mentioned in a previous section, IP addresses are to MAC addresses as mailing addresses are to Social Security numbers: the IP address serves to route information sent over the network, while MAC addresses identify the individual devices themselves. Any digital network system dealing with both types of addresses must somehow “map” each MAC address to a corresponding IP address, and this is handled by a protocol called *Address Resolution Protocol*, or *ARP*.

Every node running the ARP protocol on a digital network maintains a table of equivalent addresses, MAC to IP. This table is called the *ARP cache*, the contents of which may be displayed by running the following command on the device’s command-line interface:

```
arp -a
```

The `arp -a` command instructs the machine to print *all* (`-a`) ARP cache entries to the screen for your viewing. This, of course, only displays what that machine happens to know at that time. If the ARP cache has not been updated recently, addressing data found in the ARP cache may be out of date or even missing. Here is a partial screenshot of the `arp -a` command run on a Microsoft Windows computer, showing each IP (“internet”) address in the ARP cache and its corresponding MAC (“physical”) address:



```

C:\Windows\system32\cmd.exe
Microsoft Windows [Version 6.1.7601]
Copyright (c) 2009 Microsoft Corporation. All rights reserved.

C:\Users\tkuphald>arp -a

Interface: 172.16.4.228 --- 0xb
Internet Address      Physical Address      Type
172.16.0.6            00-17-c5-16-2c-ec     dynamic
172.16.0.19           00-15-5d-00-e8-16     dynamic
172.16.0.29           00-15-5d-00-19-07     dynamic
172.16.0.40           d8-cb-8a-36-87-f2     dynamic
172.16.0.65           00-15-5d-00-f6-1b     dynamic
172.16.0.68           00-21-5e-52-37-38     dynamic
172.16.0.69           00-15-5d-00-dd-04     dynamic
172.16.0.84           00-21-5e-52-37-38     dynamic
172.16.0.85           00-15-5d-00-f2-18     dynamic
172.16.0.86           a4-ba-db-0b-ef-71     dynamic
172.16.0.99           00-15-5d-00-1b-04     dynamic
172.16.0.103          00-21-5a-50-4f-bb     dynamic
172.16.0.143          00-23-ae-78-40-9d     dynamic
172.16.0.144          00-1a-a0-a0-b8-cf     dynamic
172.16.0.145          00-19-b9-40-10-4f     dynamic
172.16.0.163          bc-30-5b-d2-11-42     dynamic
172.16.0.220          00-15-5d-00-dd-0a     dynamic
172.16.0.236          00-15-5d-00-dd-0c     dynamic
172.16.0.240          00-15-5d-00-f6-0d     dynamic
172.16.0.242          00-1b-21-d7-aa-da     dynamic
172.16.0.250          00-22-19-02-39-1h     dynamic
172.16.0.251          00-24-e8-61-f8-a3     dynamic
172.16.2.5            00-19-b9-47-f8-c2     dynamic

```

One way to update the ARP cache on a machine with a command-line interface is to first issue a broadcast⁶⁷ ping request. Responses from active nodes on the network will populate the machine’s

⁶⁷The “ping” command is often used to test the response of a single IP node on a network, by issuing the command followed by the IP address of interest (e.g. `ping 192.168.35.70`). By contrast, a “broadcast” ping request attempts to contact a range of IP addresses within a subnet. For example, if we wished to ping all the IP addresses beginning with 192.168.35, we would issue the command with all 1’s in the last octet of the IP address field (e.g. `ping 192.168.35.255`).

ARP cache with address information, after which you may run the `arp -a` command to display those cache entries.

15.8.6 DNS

The acronym *DNS* actually stands for two related things: *Domain Name System* and *Domain Name Server*. The first meaning of “DNS” refers to the system of exchanging numerical IP addresses with alphanumeric *Uniform Resource Locators (URLs)* which are easier for human beings to remember. When you use web browser software to navigate to a web site on the Internet, you have the option of entering the URL *name* of that site (e.g. `www.google.com`) or a numerical IP address (e.g. `75.125.53.104`). Special computers connected to the Internet called *Domain Name Servers*, and *Domain Name Resolvers (DNRs)* use the *Address Resolution Protocol (ARP)* to convert your target web site name to its actual IP address so that a connection may be made between that computer and yours.

ICANN, the same organization responsible for allotting IP addresses, also maintains databases for all registered domain names.

15.8.7 Command-line diagnostic utilities

In addition to `ping` and `arp`, another utility program useful for troubleshooting network connections from a Microsoft Windows computer’s command line interface⁶⁸ is `ipconfig`. When executed, `ipconfig` returns a listing of all available (configured and operating) network interfaces on that computer:

```
C:\Documents and Settings\htc>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix . . . : 
    IP Address . . . . . : 169.254.1.2
    Subnet Mask . . . . . : 255.255.0.0
    Default Gateway . . . . . : 

Ethernet adapter Wireless Network Connection:

    Media State . . . . . : Media disconnected

C:\Documents and Settings\htc>
```

⁶⁸In UNIX-based operating systems the program used to access the command line is often called `terminal` or `xterm`. In Microsoft Windows systems it is simply called `cmd`.

The equivalent command for UNIX operating systems is `ifconfig`, shown in this screenshot:

```
root@Renegade2:/home# ifconfig
eth0      Link encap:Ethernet  HWaddr 00:13:20:08:ec:e6
          inet addr:192.168.0.64  Bcast:192.168.0.255  Mask:255.255.255.0
          inet6 addr: fe80::213:20ff:fe08:ece6/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:170901 errors:0 dropped:0 overruns:0 frame:0
          TX packets:107550 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:212178154 (212.1 MB)  TX bytes:14005068 (14.0 MB)

eth1      Link encap:Ethernet  HWaddr 00:0e:35:a2:1b:7f
          inet6 addr: fe80::20e:35ff:fea2:1b7f/64 Scope:Link
          UP BROADCAST MULTICAST  MTU:1500  Metric:1
          RX packets:441 errors:0 dropped:0 overruns:0 frame:0
          TX packets:570 errors:0 dropped:6 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:0 (0.0 B)  TX bytes:0 (0.0 B)
          Interrupt:18 Base address:0x2000 Memory:48005000-48005fff

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:8 errors:0 dropped:0 overruns:0 frame:0
          TX packets:8 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:516 (516.0 B)  TX bytes:516 (516.0 B)

root@Renegade2:/home# █
```

Some of the more interesting details contained in the output from `ifconfig` are the IPv6 addresses (in addition to IPv4 addresses), Ethernet MAC addresses (listed as “hardware addresses” or *HWaddr*), Ethernet performance data (e.g. number of collisions), IP performance data (e.g. number of IP packets received and transmitted), and details on the “loopback” address (IPv4 127.0.0.1 or IPv6 ::1).

A utility intended to reveal the DNS name of a computer given its IP address, or visa versa, is `nslookup`. The same command works on Microsoft Windows and UNIX operating systems alike. Here, we see the UNIX version used to identify four IP addresses of the popular Google search engine web site, followed by the Microsoft Windows version:

```
root@Renegade2:/home# nslookup www.google.com
Server:          192.168.0.1
Address:         192.168.0.1#53

Non-authoritative answer:
www.google.com canonical name = www.l.google.com.
Name:   www.l.google.com
Address: 74.125.53.103
Name:   www.l.google.com
Address: 74.125.53.147
Name:   www.l.google.com
Address: 74.125.53.99
Name:   www.l.google.com
Address: 74.125.53.104
```

```
Microsoft Windows [Version 5.2.3790]
(C) Copyright 1985-2003 Microsoft Corp.

U:\>nslookup www.google.com
Server:   btc2000-dc1.bellingham-tech.edu
Address:  172.16.0.240

Non-authoritative answer:
Name:     www.l.google.com
Addresses: 209.85.173.104, 209.85.173.103, 209.85.173.147, 209.85.173.99
Aliases:  www.google.com

U:\>_
```

Another utility used to explore network connections is `traceroute` (spelled `tracert` on Microsoft Windows operating systems). This utility sends a test packet to the designated destination address, returning information on all the “hops” the IP packet takes between computers along the network to reach its destination and the amount of time taken to make the trip. Execution of `traceroute` on a UNIX computer and `tracert` on a Microsoft Windows computer are shown here:

```

root@Renegade2:/home# traceroute www.google.com
traceroute to www.google.com (74.125.53.147), 30 hops max, 40 byte packets
 1 home (192.168.0.1) 4.763 ms 4.803 ms 4.784 ms
 2 tukw-dsl-gw22-214.tukw.qwest.net (63.231.10.214) 57.927 ms 59.993 ms 61.916 ms
 3 tukw-agw1.inet.qwest.net (71.217.184.169) 63.924 ms 65.905 ms 67.882 ms
 4 sea-core-01.inet.qwest.net (67.14.1.194) 71.775 ms 71.784 ms 73.609 ms
 5 sea-brdr-01.inet.qwest.net (205.171.26.54) 75.642 ms 77.442 ms 79.421 ms
 6 63.146.26.198 (63.146.26.198) 81.438 ms 67.052 ms 68.856 ms
 7 sl-gw20-sea-0-0-0.sprintlink.net (144.232.6.8) 70.633 ms 56.617 ms 60.219 ms
 8 sl-googl13-199181-0.sprintlink.net (144.224.13.138) 62.133 ms 64.301 ms 66.162
 9 209.85.249.32 (209.85.249.32) 68.140 ms 209.85.249.34 (209.85.249.34) 70.028 ms
10 216.239.46.204 (216.239.46.204) 79.865 ms 81.739 ms 85.587 ms
11 64.233.174.121 (64.233.174.121) 249.193 ms 64.233.174.129 (64.233.174.129) 113.
12 72.14.232.70 (72.14.232.70) 93.458 ms 72.14.232.10 (72.14.232.10) 99.574 ms 72.
13 72.14.232.6 (72.14.232.6) 68.876 ms 72.14.232.2 (72.14.232.2) 70.251 ms pw-in-f
root@Renegade2:/home#

```

```

U:\>tracert www.google.com

Tracing route to www.l.google.com [209.85.173.99]
over a maximum of 30 hops:
  1  4294964928 ms  4294964927 ms  4294964927 ms  134.39.250.1
  2  4294964927 ms  4294964927 ms  4294964927 ms  bellingham-2691.ctc.edu [192.6
4.1.105]
  3  4294964931 ms  4294964931 ms  4294964930 ms  ge-0-1-0--941.seawescarl.infra
.wa-k20.net [68.179.207.210]
  4  4294964931 ms  4294964931 ms  4294964930 ms  ge-3-0-3--0.seawescor1.infra.
wa-k20.net [68.179.203.26]
  5  4294964931 ms  4294964931 ms  4294964931 ms  ge-2-2-0--311.iccr-sttlwa01-02
.infra.pnw-gigapop.net [209.124.188.182]
  6  4294964931 ms  4294964930 ms  4294964931 ms  pnwgp-cust.tr01-sttlwa01.trans
itrail.net [137.164.131.186]
  7  4294964931 ms  4294964931 ms  4294964931 ms  te4-3--301.tr01-sttlwa01.trans
itrail.net [137.164.131.185]
  8  4294964931 ms  4294964931 ms  4294964931 ms  137.164.130.158
  9  4294964931 ms  4294964931 ms  4294964931 ms  209.85.249.32
 10 4294964935 ms  4294964938 ms  4294964938 ms  216.239.46.208
 11 4294964936 ms  4294964937 ms  4294964937 ms  64.233.174.127
 12 4294964937 ms  4294964936 ms  4294964937 ms  209.85.251.149
 13 4294964943 ms  4294964937 ms  4294964941 ms  209.85.251.145
 14 4294964941 ms  4294964944 ms  4294964937 ms  mh-in-f99.google.com [209.85.1
73.99]

Trace complete.

U:\>

```

15.9 Transmission Control Protocol (TCP) and User Datagram Protocol (UDP)

At the next OSI Reference Model layer (layer 4) is a set of protocols specifying how reliable communication “connections” should be established between devices on a digital network. Rather than specifying addresses for routing data packets on a large network (OSI layer 3), layer 4 works on the level of virtual data “ports” at the transmitting and receiving devices. The purpose of these virtual ports is to manage multiple types of data transactions to and from the same IP address, such as in the case of a personal computer accessing a web page (using HTTP) and sending an email message (using SMTP) at the same time. An analogy to help understand the role of ports is to think of multiple packages delivered to different people at a common address such as a business office. The mailing address for the office is analogous to the IP address of a computer exchanging data over a network: it is how other computers on the network “find” that computer. The personal or department names written on the different packages are analogous to virtual ports on the computer: “places” within that node where specific messages are directed once they arrive.

Transmission Control Protocol (TCP) and *User Datagram Protocol (UDP)* are two methods used to manage data flow through “ports” on a DTE device, with TCP being the more complex and robust of the two. Both TCP and UDP rely on IP addressing to specify which devices send and receive data, which is why you will often see these protocols listed in conjunction with IP (e.g. TCP/IP and UDP/IP). TCP and UDP are both useless on their own: a protocol specifying port locations without an IP address would be as meaningless as a package placed in the general mail system with just a name but with no street address. Conversely, IP anticipates the presence of a higher-level protocol such as TCP or UDP by reserving a portion of its “datagram” (packet) bit space for a “protocol” field⁶⁹ to specify which high-level protocol generated the data in the IP packet.

TCP is a complex protocol specifying not only which virtual “ports” will be used at the sending and receiving devices, but also how data integrity will be guaranteed. Data communicated via TCP are sent in blocks of bits called *segments*. In the same way that the IP algorithm encapsulates data into self-contained “packets” with the necessary routing data to ensure proper delivery to the destination, the TCP algorithm encapsulates data with “header” bits specifying such details as sequence number, acknowledgment identification, checksum (for error-detection), urgency of the message, and optional data. TCP takes responsibility for ensuring both devices on the network are ready to exchange data, breaks the data block into segments of an appropriate size to be encapsulated inside IP packets on the transmitting end, checks each segment for corruption (using the CRC method) upon receipt, reassembles those segments into the whole data block at the receiving end, and terminates the connection between devices upon completion. If a TCP segment successfully arrives at the receiving device, the TCP protocol running at the receiving device acknowledges this to the transmitting device. If the transmitting device fails to receive an acknowledgment within a specified time, it automatically re-sends the segment. If a TCP segment arrives corrupted, the TCP protocol running at the receiving device requests a re-send from the transmitting device. TCP thus fulfills the task of the author and publisher in the manuscript analogy, handling the disassembly and reassembly of the manuscript. In rigorously checking the integrity of the data transfer, TCP’s functionality is akin to sending each manuscript package via *certified mail* where the author may

⁶⁹Both IPv4 and IPv6 reserve eight bits for this purpose.

track the successful delivery of each and every package.

Another feature of TCP is end-to-end *flow control*, enabling the receiving device to halt the incoming data stream for any reason, for example if its buffer memory becomes full. This is analogous to the XON/XOFF control signals used in simple serial communication protocols such as EIA/TIA-232. In the case of large networks, TCP flow control manages the flow of data segments even when the communication paths are dynamic and liable to change.

If IP (Internet Protocol) is the “glue” that holds the Internet together, TCP (Transmission Control Protocol) is what makes that glue strong enough to be practically useful. Without TCP, data communication over the wildly disparate physical networks that comprise the world-wide Internet would be far less reliable.

UDP is a much simpler protocol, lacking not only the segment sequencing of TCP but also lacking many of the data-integrity features of TCP. Unlike TCP, the UDP algorithm does not sequence the data block into numbered segments at the transmitting end or reassemble those numbered segments at the receiving end, instead relegating this task to whatever application software is responsible for generating and using the data. Likewise, UDP does not error-check each segment and request re-transmissions for corrupted or lost segments, once again relegating these tasks to the application software running in the receiving computer. It is quite common to see UDP applied in industrial settings, where communication takes place over much smaller networks than the world-wide Internet, and where the IP data packets themselves tend to be much smaller as well. Another reason UDP is more common in industrial applications is that it is easier to implement in the “embedded” computer hardware at the heart of many industrial devices. The TCP algorithm requires greater computational power and memory capacity than the UDP algorithm, and so it is much easier to engineer a single-chip computer (i.e. microcontroller) to implement UDP than it would be to implement TCP.

Information-technology professionals accustomed to the world of the Internet may find the use of UDP rather than TCP in industrial networks alarming. However, there are good reasons why industrial data networks do not necessarily need all the data-integrity features of TCP. For example, many industrial control networks consist of seamless Ethernet networks, with no routers (only hubs and/or switches) between source and destination devices. With no routers to go through, all IP packets communicated between any two specified devices are guaranteed to take the exact same data path rather than being directed along alternate, dynamic routes. The relative stability and consistency of communication paths in such industrial networks greatly simplifies the task of packet-reassembly at the receiving end, and makes packet corruption less likely. In fact, many industrial device messages are short enough to be communicated as single IP packets, with no need for fragmentation or re-assembly at all!

A noteworthy exception to the use of UDP in industrial applications is FOUNDATION Fieldbus HSE (High-Speed Ethernet), which uses TCP rather than UDP to ensure reliable transmission of data from source to destination. Here, the designers of FOUNDATION Fieldbus HSE opted for the increased data integrity offered by TCP rather than use UDP as may other industrial IP-based network protocols do.

Using another utility program on a personal computer called `netstat` (available for both Microsoft Windows and UNIX operating systems) to check active connections⁷⁰, we see the various IP addresses and their respective port numbers (shown by the digits following the colon after the IP address) as a list, organized by TCP connections and UDP connections:

```
C:\Documents and Settings\btc>netstat -an
Active Connections

Proto Local Address           Foreign Address         State
TCP   0.0.0.0:135              0.0.0.0:0               LISTENING
TCP   0.0.0.0:445              0.0.0.0:0               LISTENING
TCP   0.0.0.0:2869            0.0.0.0:0               LISTENING
TCP   127.0.0.1:1025          0.0.0.0:0               LISTENING
TCP   127.0.0.1:5152          0.0.0.0:0               LISTENING
TCP   169.254.1.2:23          169.254.1.1:1116       ESTABLISHED
TCP   169.254.1.2:139        0.0.0.0:0               LISTENING
UDP   0.0.0.0:445             *:*                     *:*
UDP   0.0.0.0:500            *:*                     *:*
UDP   0.0.0.0:1062           *:*                     *:*
UDP   0.0.0.0:4500           *:*                     *:*
UDP   0.0.0.0:7725           *:*                     *:*
UDP   127.0.0.1:123          *:*                     *:*
UDP   127.0.0.1:1063         *:*                     *:*
UDP   127.0.0.1:1066         *:*                     *:*
UDP   127.0.0.1:1900         *:*                     *:*
UDP   169.254.1.2:123        *:*                     *:*
UDP   169.254.1.2:137        *:*                     *:*
UDP   169.254.1.2:138        *:*                     *:*
UDP   169.254.1.2:1900      *:*                     *:*
C:\Documents and Settings\btc>
```

An interesting distinction between TCP and UDP is evident in this screenshot. Note how each of the TCP connections has an associated “state” (either LISTENING or ESTABLISHED), while none of the UDP connections has a state associated with it. Recall that TCP is responsible for initiating and terminating the connection between network devices, and as such each TCP connection must have a descriptive state. UDP, on the other hand, is nowhere near as formal as TCP in establishing connections or ensuring data transfer integrity, and so there are no “states” to associate with any of the UDP connections shown. Data arriving over a UDP connection simply shows up unannounced, like an impolite guest. Data arriving over a TCP connection is more like a guest who announces in advance when they will arrive, and also says “good bye” to the host as they depart.

Many different port numbers have been standardized for different applications at OSI Reference Model layers above 4 (above that of TCP or UDP). Port 25, for example, is always used for SMTP (Simple Mail Transfer Protocol) applications. Port 80 is used by HTTP (HyperText Transport Protocol), a layer-7 protocol used to view Internet “web” pages. Port 443 is used by HTTPS, an encrypted (secure) version of HTTP. Port 107 is used by TELNET applications, a protocol whose purpose it is to establish command-line connections between computers for remote administrative work. Port 22 is used by SSH, a protocol similar to TELNET but with significantly enhanced security. Port 502 is designated for use with Modbus messages communicated over TCP/IP.


⁷⁰In this particular case, I typed `netstat -an` to specify *all* (a) ports with *numerical* (n) IP addresses and port numbers shown.

15.10 The HART digital/analog hybrid standard

A technological advance introduced in the late 1980's was *HART*, an acronym standing for **H**ighway **A**ddressable **R**emote **T**ransmitter. The purpose of the HART standard was to create a way for instruments to digitally communicate with one another over the same two wires used to convey a 4-20 mA analog instrument signal. In other words, HART is a *hybrid* communication standard, with one variable (channel) of information communicated by the analog value of a 4-20 mA DC signal, and another channel for digital communication whereby many other variables could be communicated using pulses of current to represent binary bit values of 0 and 1. Those digital current pulses are *superimposed* upon the analog DC current signal, such that the same two wires carry both analog and digital data simultaneously.

15.10.1 Basic concept of HART

Looking at a standard loop-powered (2-wire) process transmitter circuit, we see the transmitter, a DC power supply (voltage source), and usually a 250 ohm resistor to create a 1 to 5 volt signal readable by any voltage-sensing indicator, controller, or recorder:




The transmitter's primary function in this circuit is to regulate current to a value representative of the measured process variable (e.g. pressure, temperature, flow, etc.) using a range of 4 to 20 mA, while the DC voltage source provides power for the transmitter to operate. Loop-powered instruments are very common in industrial instrumentation because they allow both power and (analog) data to be conveyed on the same pair of wires.

With the advent of microprocessor-based process transmitters, it became possible for instrument technicians to digitally configure parameters inside the transmitter (e.g. range values, damping values) and also query the transmitter for self-diagnostic alarms. In order to make full use of this digital functionality, though, there had to be some way to *communicate* digital data to and from the process transmitter over the same two wires used to convey the 4-20 mA analog signal. Otherwise, the only way to access this rich array of digital data inside the transmitter would be to connect a communicator device to some data port located on the transmitter itself, which is inconvenient due to the nature of how these transmitters are used in industry (located in dirty places, often hard to access while carrying a personal computer or other communication device).

Thus the *HART* communication protocol was born to address this need. HART communicates digital data along the loop conductors in the form of AC signals (audio-frequency tones) *superimposed* on the 4-20 mA DC current signal. A modem built into the smart transmitter translates these AC signals into binary bits, and vice-versa. Now, instrument technicians could "talk" with the new microprocessor-based transmitters simply by connecting a HART communications device at any point along the two-wire cable, even at the far end where the cable terminates at the control system hardware (panel-mounted controller, PLC, DCS, etc.).

Being able to communicate digital data over the same wire pair as the DC power and analog signal opens a whole new range of possibilities. Now, the field-mounted transmitter can communicate self-diagnostic information, status reports, alarms, and even multiple process variables to the control system in addition to the original analog signal representing the (main) process variable. With digital communication, the only data limitation is speed (data rate), not quantity. The control system may even communicate information *to* the transmitter using the same digital protocol, using this digital data channel to switch between different measurement range sets, activating special features (e.g. square-root characterization, damping, etc.), automatically and remotely.


With HART, the traditional series-connected circuit configuration of transmitter, DC power supply, and resistor remains unchanged. A HART-enabled transmitter, though, is equipped with a built-in digital microcontroller managing its functions, and this miniature computer is able to send and receive digital data as AC signals (current pulses in sending mode, voltage pulses in receiving mode) superimposed on the same two wires carrying the 4 to 20 mA analog signal and DC power. Any other computer device equipped with a HART modem, the necessary configuration software, and an appropriate device description⁷¹ for that particular instrument may communicate with the HART transmitter if connected in parallel with the transmitter's loop power terminals:



This external computer – through the use of HART data transmission – now has the ability to monitor details of the transmitter's operation, configure the transmitter, update its measurement ranges, and perform a host of additional functions.

⁷¹A *Device Description*, or *DD* (DD) file, is analogous to a “driver” file used to instruct a personal computer how to communicate with a printer, scanner, or any other complex peripheral device. In this case, the file instructs the HART configuration computer on how it should access parameters inside the field instrument's microcontroller. Without an appropriate DD file loaded on the configuration computer, many of the field instrument's parameters may be inaccessible.

The connection between the HART modem and the HART-enabled transmitter need not be made directly at the transmitter's terminals. *Any* set of points in the circuit electrically parallel to the transmitter's terminals are permissible as connection points for the HART modem. This flexibility is a great advantage in loop circuits spanning long distances, enabling the technician to connect their HART configuration equipment at the most physically convenient location:



A convenient alternative to a personal computer and HART modem is a special hand-held device called a *HART communicator*. Two different models of HART communicator are shown in the following photographs, a Rosemount model 268 on the left and an Emerson model 375 on the right:



HART communicators are battery-powered, portable devices built specifically for configuring HART-enabled field instruments. Like personal computers, they must be updated with DD files⁷² to be able to communicate with the latest models of HART-enabled field instruments.

Perhaps the single greatest disadvantage of HART data communication is its slow speed. Not only is the bit rate slow by modern standards – only 1200 bits per second – but the encoding of HART data is laden with “overhead” bits⁷³ required to synchronize devices, specify device addresses on a network, check for errors, and other functions necessary for making HART a reliable data-communications protocol. As a result, the typical time required for a complete HART transaction (one device requesting data, and the other device replying) is half a second! This means on average no more than two complete messages per second may be communicated by HART. This limits the use of HART data communication to non-routine purposes (e.g. error messages sent by field

⁷²A “DD” file, or *Device Descriptor* file, is akin to a *driver* file used in a personal computer to allow it to communicate data with some peripheral device such as a printer. DD files basically tell the HART communicator how it needs to access specific data points within the HART field instrument.

⁷³Every byte (8 bits) of actual HART data is sent as an asynchronous serial frame with a start bit, parity bit, and stop bit, so that 11 bits’ worth of time are necessary to communicate 8 bits of real data. These “byte frames” are then packaged into larger message units called HART *telegrams* (similar to Ethernet data frames) which include bits for synchronizing receiving devices, specifying device addresses, specifying the length of the data payload, communicating device status, etc.

instruments, range changes made by technicians) rather than continuous process measurement and control in all but the slowest process applications.

HART technology has breathed new life into the venerable 4-20 mA analog instrumentation signal standard. It has allowed new features and capabilities to be added on to existing analog signal loops without having to upgrade wiring or change all the instruments in the loop. Some of the features of HART are listed here:

- Diagnostic data may be transmitted by the field device (self-test results, out-of-limit alarms, preventive maintenance alerts, etc.)
- Field instruments may be re-ranged remotely through the use of HART communicators
- Technicians may use HART communicators to force field instruments into different “manual” modes for diagnostic purposes (e.g. forcing a transmitter to output a fixed current so as to check calibration of other loop components, manually stroking a valve equipped with a HART-capable positioner)
- Field instruments may be programmed with identification data (e.g. tag numbers corresponding to plant-wide instrument loop documentation)


It should be mentioned that HART communication over 4-20 mA signal wires is a legacy technology. At the time of this writing (2011), HART protocol is still the most popular form of wired digital field instrument communication in industrial use. However, more modern digital standards such as Profibus and FOUNDATION Fieldbus deliver all the benefits of HART technology and more. It seems that wired-HART will remain in wide industrial use for many years to come, but it is really just the beginning of digital field instrument technology and does not represent the state of the art. For more information regarding modern “fieldbus” digital instrumentation, refer to chapter 16 on the FOUNDATION Fieldbus standard beginning on page 1133.

An important addition to the HART standard introduced with version 7 is *wireless* (radio) communication capability. This portion of the standard describes how HART data may be communicated via radio waves instead of audio-frequency AC signal bursts superimposed on DC signal wires. Here, HART is the sole method of exchanging process data between field instruments rather than a secondary method to a 4-20 mA DC analog signal. Communicating over radio waves eliminates the theoretical speed barrier faced by wired-HART (1200 bits per second) while still permitting HART-compliant computer software and field communicator devices to work with these wireless instruments. For more information regarding the *WirelessHART* standard, refer to section 17.2 beginning on page 1232.

15.10.2 HART physical layer


The HART standard was developed with existing installations in mind. The signals had to be robust enough to travel over twisted-pair cables of very long length and unknown characteristic impedance. This meant that the data communication rate for the digital data had to be very slow, even by 1980's standards, in order to avoid problems created by reflections along unterminated cabling. The HART standard is concerned only with three layers of the OSI Reference model: layer 1 (FSK modulation, ± 0.5 mA transmitter signaling), layer 2 (Master-slave arbitration, data frame organization), and layer 7 (specific commands to read and write device data). Layers 3 through 6 of the OSI model are irrelevant to the HART standard.

Digital data is encoded in HART using the Bell 202 modem standard: two audio-frequency "tones" (1200 Hz and 2200 Hz) are used to represent the binary states of "1" and "0," respectively, transmitted at a rate of 1200 bits per second. This is known as *frequency-shift keying*, or *FSK*. The physical representation of these two frequencies is an AC current of 1 mA peak-to-peak superimposed on the 4-20 mA DC signal. Thus, when a HART-compatible device "talks" digitally on a two-wire loop circuit, it produces tone bursts of AC current at 1.2 kHz and 2.2 kHz. The receiving HART device "listens" for these AC current frequencies and interprets them as binary bits. The following photograph shows a HART waveform captured by a digital oscilloscope:




Each "1" bit in this HART waveform is a single cycle of 1200 Hz tone, while each "0" bit is a double-cycle of 2200 Hz tone. The waveform shown here is an alternating series of "1" and "0" bits.

An important consideration in HART current loops is that the total loop resistance (precision resistor values plus wire resistance) must fall within a certain range: 250 ohms to 1100 ohms. Most 4-20 mA loops (containing a single 250 ohm resistor for converting 4-20 mA to 1-5 V) measure in at just over 250 ohms total resistance, and work quite well with HART. Even loops containing two 250 ohm precision resistors meet this requirement. Where technicians often encounter problems is when they set up a loop-powered HART transmitter on the test bench with a lab-style power supply and *no* 250 ohm resistor anywhere in the circuit:



The HART transmitter may be modeled as two parallel current sources: one DC and one AC. The DC current source provides the 4-20 mA regulation necessary to represent the process measurement as an analog current value. The AC current source turns on and off as necessary to “inject” the 1 mA P-P audio-frequency HART signal along the two wires. Inside the transmitter is also a HART modem for interpreting AC voltage tones as HART data packets. Thus, data transmission takes place through the AC current source, and data reception takes place through a voltage-sensitive modem, all inside the transmitter, all “talking” along the same two wires that carry the DC 4-20 mA signal.

For ease of connection in the field, HART devices are designed to be connected in parallel with each other. This eliminates the need to break the loop and interrupt the DC current signal every time we wish to connect a HART communicator device to communicate with the transmitter. A typical HART communicator may be modeled as an AC voltage source⁷⁴ (along with another HART voltage-sensitive modem for receiving HART data). Connected in parallel with the HART transmitter, the complete circuit looks something like this:




With all these sources in the same circuit, it is advisable to use the *Superposition Theorem* for analysis. This involves “turning off” all but one source at a time to see what the effect is for each source, then superimposing the results to see what all the sources do when all are working simultaneously.

⁷⁴The HART standard specifies “master” devices in a HART network transmit AC voltage signals, while “slave” devices transmit AC current signals.


We really only need to consider the effects of either AC source to see what the problem is in this circuit with no loop resistance. Consider the situation where the transmitter is sending HART data to the communicator. The AC current source inside the transmitter will be active, injecting its 1 mA P-P audio-frequency signal onto the two wires of the circuit. The AC voltage source in the communicator will disconnect itself from the network, allowing the communicator to “listen” to the transmitter’s data.

To apply the Superposition Theorem, we replace all the other sources with their own equivalent internal resistances (voltage sources become “shorts,” and current sources become “opens”). The HART communicator will be modeled as an “open” even though it is technically a voltage source because it must turn itself off (i.e. switch to high-impedance mode) in order for any field device to communicate to it:




The HART communicator is “listening” for those audio tone signals sent by the transmitter’s AC source, but it “hears” nothing because the DC power supply’s equivalent short-circuit prevents any significant AC voltage from developing across the two wires. This is what happens when there is no loop resistance: no HART device is able to receive data sent by any other HART device.

The solution to this dilemma is to install a resistance of at least 250 ohms but not greater than 1100 ohms between the DC power source and all other HART devices, like this:




Loop resistance must be at least 250 ohms to allow the 1 mA P-P AC signal to develop enough voltage to be reliably detected by the HART modem in the listening device. The upper limit (1100 ohms) is not a function of HART communication so much as it is a function of the DC voltage drop, and the need to maintain a minimum DC terminal voltage at the transmitter for its own operation. If there is too much loop resistance, the transmitter will become “starved” of voltage and act erratically. In fact, even 1100 ohms of loop resistance may be too much if the DC power supply voltage is modest.

Loop resistance is also necessary for the HART transmitter to receive data signals transmitted by the HART communicator. If we analyze the circuit when the HART communicator's voltage source is active (replacing the DC power supply with a short and the transmitter current sources with opens), we get this result:



Without the loop resistance in place, the DC power supply would “short out” the communicator’s AC voltage signal just as effectively as it shorted out the transmitter’s AC current signal. The presence of a loop resistor in the circuit prevents the DC power supply from “loading” the AC voltage signal by the communicator. This AC voltage is seen in the diagram as being directly in parallel with the transmitter, where its internal HART modem receives the audio tones and processes the data packets.

Manufacturers' instructions generally recommend HART communicator devices be connected directly in parallel with the HART field instrument, as shown in the previous schematic diagrams. However, it is also perfectly valid to connect the communicator device directly in parallel with the loop resistor like this:



Connected directly in parallel with the loop resistor, the communicator is able to receive transmissions from the HART transmitter just fine, as the DC power source acts as a dead short to the AC current HART signal and passes it through to the transmitter.


This is nice to know, as it is often easier to achieve an alligator-clip connection across the leads of a resistor than it is to clip in parallel with the loop wires when at a terminal strip or at the controller end of the loop circuit.

15.10.3 HART multidrop mode

The HART standard also supports a mode of operation that is totally digital, and capable of supporting multiple HART instruments on the same pair of wires. This is known as *multidrop mode*, where field instruments are slaves and the communicator (or control system) is the master.

Every HART instrument has an *address* number, which is typically set to a value of zero (0). A network address is a number used to distinguish one device from another on a broadcast network, so messages broadcast across the network may be directed to specific destinations. When a HART instrument operates in digital/analog hybrid mode, where it must have its own dedicated wire pair for communicating the 4-20 mA DC signal between it and an indicator or controller, there is no need for a digital address. An address becomes necessary only when multiple devices are connected to the same network wiring, and there arises a need to digitally distinguish one device from another on the same network.

This is a functionality the designers of HART intended from the beginning, although it is rarely used in industry. Multiple HART instruments may be connected directly in parallel with one another along the same wire pair, and information exchanged between those instruments and a host system, if the HART address numbers are set to non-zero values (between 1 and 15):




Setting an instrument's HART address to a non-zero value is all that is necessary to engage multidrop mode. The address numbers themselves are irrelevant, as long as they fall within the range of 1 to 15 and are unique within that network. Once a HART instrument is configured for multi-drop mode (i.e. given a non-zero HART address), its current becomes fixed at 4 mA and no longer varies with the process variable measurement.

The major disadvantage of using HART instruments in multidrop mode is its slow speed. Due to HART's slow data rate (1200 bits per second), it may take several seconds to access a particular instrument's data on a multidropped network. For some applications such as temperature measurement, this slow response time may be acceptable. For inherently faster processes such as liquid flow control, it would not be nearly fast enough to provide up-to-date information for the control system to act upon.

15.10.4 HART multi-variable transmitters and burst mode

Some “smart” instruments have the ability to report multiple process variables. A good example of this is Coriolis-effect flowmeters, which by their very nature simultaneously measure the density, flow rate, and temperature of the fluid passing through them. A single pair of wires can only convey *one* 4-20 mA analog signal, but that same pair of wires may convey *multiple* digital signals using HART protocol.

If the host system receiving the transmitter’s signal(s) is HART-ready, it may digitally poll the transmitters for all variables. If, however, the host system does not “talk” using the HART protocol, some other means must be found to decode the wealth of digital data digitally offered by the multi-variable transmitter. One such device is Rosemount’s model 333 HART “Tri-Loop” demultiplexer shown in the following photograph:



This device uses a feature of HART called *burst mode* where one HART instrument repeatedly transmits information rather than waiting to be polled by a master HART device (such as a control system or a hand-held communicator)⁷⁵. The Rosemount 333 receives these bursts of data, converting as many as three digitally-encoded HART variables into independent 4-20 mA analog output signals, which any suitable analog indicator or controller device may then receive.

It should be noted that the same caveat applicable to multidrop HART systems (i.e. slow speed) applies to multi-variable HART communication as well. Even though burst mode eliminates the need for master-device communications and therefore occupies less time than normal (polled) HART communication, typical wait times for HART data may still be too long for some industrial process applications. Typical HART burst rate is approximately three HART “telegram” packets per second, as opposed to two complete transmissions per second in the case of polled HART where

⁷⁵Truth be told, HART instruments configured to operate in burst mode are still able to respond to queries from a master device, just not as often. Between bursts, the HART slave device waits a short time to allow any master devices on the network to poll. When polled, the slave device will respond as it normally would, then resumes its bursts of unpolled data once again. This means that normal master/slave communication with a HART instrument set for burst mode will occur at a slower pace than if the instrument is set for normal mode.

the master requests data and the slave replies. If we configure a process transmitter to burst-communicate three different variables, for example, the update rate for each one of those variables will be approximately once per second (three variables communicated at a total rate of three HART telegrams per second), which is simply too slow for many industrial applications (e.g. closed-loop liquid flow control). In applications where speed is not a concern, however, HART communication is a very practical solution for acquiring multiple channels of data from one instrument over a single pair of wires.

15.11 Modbus

Developed by the Modicon company (the original manufacturer of the *Programmable Logic Controller*, or *PLC*) in 1979 for use in its industrial control products, *Modbus* is a protocol designed specifically for exchanging process data between industrial control devices. The Modbus standard does not specify any details of physical networking, and thus may be deployed on many different types of physical networks. In other words, Modbus primarily falls within layer 7 of the OSI Reference Model (the so-called “Application Layer”) and therefore is compatible⁷⁶ with any lower-level communication protocols including EIA/TIA-232, EIA/TIA-485, Ethernet (the latter via TCP/IP), and a special token-passing network also developed by Modicon called *Modbus Plus*. The Modbus standard primarily defines the *meaning* of various Modbus commands, the addressing scheme used to place data within devices, and the formatting of the data.

Modbus consists of a set of standardized digital codes intended to read data from and write data to industrial devices. A Modbus-compliant industrial device has been programmed to understand these codes and respond to them appropriately when received. The simplest Modbus codes read and write single bits of data in the device’s memory, for example the status of a PLC input channel, PLC output channel, or status bit within a PLC program. Other Modbus codes operate on 16-bit words of data, useful for reading and writing counter and timer accumulated values, operands for mathematical instructions, converted analog signals, etc.


Early implementations of Modbus used EIA/TIA-485 as the network physical layer, which is strictly a layer 1 protocol. This meant that Modbus needed to specify a channel arbitration scheme in order to negotiate communications with multiple devices on a network. The arbitration chosen was master/slave, where one PLC functioned as the master Modbus device and all other devices functioned as Modbus slaves.

Interestingly, this vestige of master/slave arbitration survives to this day, even when Modbus commands are communicated via networks with their own differing arbitration methods. For example, Modbus commands communicated over Ethernet still reference “slave” addresses even though the Ethernet network those messages are sent over uses CSMA/CD arbitration. In other words, there is a hint of OSI layer 2 embedded within Modbus messages that still dictates which Modbus devices may issue commands and which must obey commands.

⁷⁶These Modbus data frames may be communicated directly in serial form, or “wrapped” in TCP segments and IP packets and Ethernet frames, or otherwise contained in any form of packet-based protocol as needed to transport the data from one device to another. Thus, Modbus does not “care” how the data is communicated, just what the data means for the end-device.

15.11.1 Modbus overview


We may begin our exploration of Modbus by first considering an example of a PLC-controlled motor system that does *not* employ Modbus. Here, the PLC sends individually-wired Forward, Reverse, and Stop, and speed-control command signals to a variable-frequency drive (VFD) which then sends three-phase power of varying frequency to an electric motor to do some useful task:



The discrete commands (e.g. Stop, Forward, Reverse) are nothing more than on/off contact closures provided by the PLC's output channels to the VFD's input terminals. When the PLC commands the VFD to run in the Reverse direction, it simply activates output channel 0/1 which closes a relay contact inside the PLC to connect the VFD's "Rvs" terminal to the VFD's "Com" terminal. The VFD detects this electrical continuity, and responds by running the motor in its reverse direction. Motor speed is commanded by an analog voltage signal (typically 0 to 10 volts DC) output by the PLC, with 0 volts representing zero speed and 10 volts representing full speed. The VFD receives this analog voltage signal and responds to it by outputting the appropriate frequency of three-phase AC power to the induction motor.

While this system is certainly functional, it may be improved upon through the use of Modbus.


Now consider this updated motor control system, where the only connecting wires between the PLC and VFD is a single two-conductor cable between the Modbus/RS-485 terminals of both devices. The PLC functions as a Modbus master device while the VFD functions as a Modbus slave:



By using appropriate Modbus commands transmitted to the VFD, the PLC is able to issue all the same commands (e.g. Stop, Forward, Reverse, speed control) as before but using far fewer wires. For example Modbus command code 05 writes a single bit of data to the receiving device, allowing the PLC to send discrete-signal commands to the VFD one at a time. When the PLC commands the VFD to run in the Reverse direction, it issues a 05 command followed by a “1” data bit addressed to the appropriate memory location inside the VFD reserved for the “Reverse” command bit. When the PLC commands the VFD to change motor speed, it issues an 06 Modbus code (“write register”) followed by a 16-bit number representing the desired motor speed and the appropriate address within the VFD reserved for speed command.

Not only can the PLC issue all the same commands as before, but it may also *read* data from the VFD which it could not do before. For example, if the VFD provides a memory location for storing fault codes (e.g. motor overcurrent, bus undervoltage, etc.), the PLC may be programmed to issue an 03 Modbus code to read a single register (16 bit binary number) from that memory location within the VFD, and thereby monitor the status of the VFD to alert human technicians of potential problems, and/or to modify its own controlling of the motor.

Another advantage of the Modbus communication standard is that it is designed to address multiple devices on the same network. This means our hypothetical PLC is not limited to controlling and monitoring just one motor, but *up to 247 separate Modbus slave devices* on the same two-wire communication cable! The following illustration shows how this might work for multiple motors:



Each VFD is given its own Modbus network *slave address*, so that the PLC is able to distinguish between the two drives when communicating on the same wire pair. Every Modbus code transmitted by the PLC contains this address as a single byte (8 bits) of data in order to make the receiving VFD aware that the code applies to it and not to any other Modbus device on the network. In

this example, we may wish to address VFD #1 with Modbus address 1, and VFD #2 with Modbus address 2. The Modbus standard provides a “broadcast address” of 0 which addresses *all* devices on the network simultaneously. For example, if the PLC needed to start all motors in the same direction at once, it could issue a Modbus code 05 (write a single bit) to the same address inside each VFD representing the command bit for the correct direction of motor rotation. So long as the VFDs are identically configured, the data will be received and interpreted by each VFD identically which will cause them to both start up in the same direction.

The only disadvantages to using Modbus as opposed to dedicated wires for each sensing and control function are *speed* and *reliability*. Modbus is necessarily slower than dedicated wire control because the PLC cannot simultaneously issue different commands on the network. For example, if the PLC needed to tell a VFD to begin turning its motor in the forward direction at 1050 RPM, the Modbus-based system would need to issue two separate Modbus codes whereas the individually-wired system could issue these commands all at once. This disadvantage, however, is hardly worth considering if the Modbus network communicates at reasonably high speed (thousands of bits per second). The disadvantage of reliability may be readily perceived if we consider how each system would respond to a wire fault (e.g. one wire coming loose and disconnected from a screw terminal). In the individually-wired system, one wire fault disables that one motor-control function but not necessarily any of the other functions. In the Modbus-based system, one wire fault disables *everything* because any Modbus communication requires full function of that two-conductor communication cable. The problem is even larger when multiple devices are controlled by the same Modbus cable: if a fault occurs between the controlling PLC and all the field devices, the PLC will lose control (and monitoring) for every one of those field devices! This is a factor worth considering when deciding whether or not to use any digital communication method for monitoring and control of multiple devices.

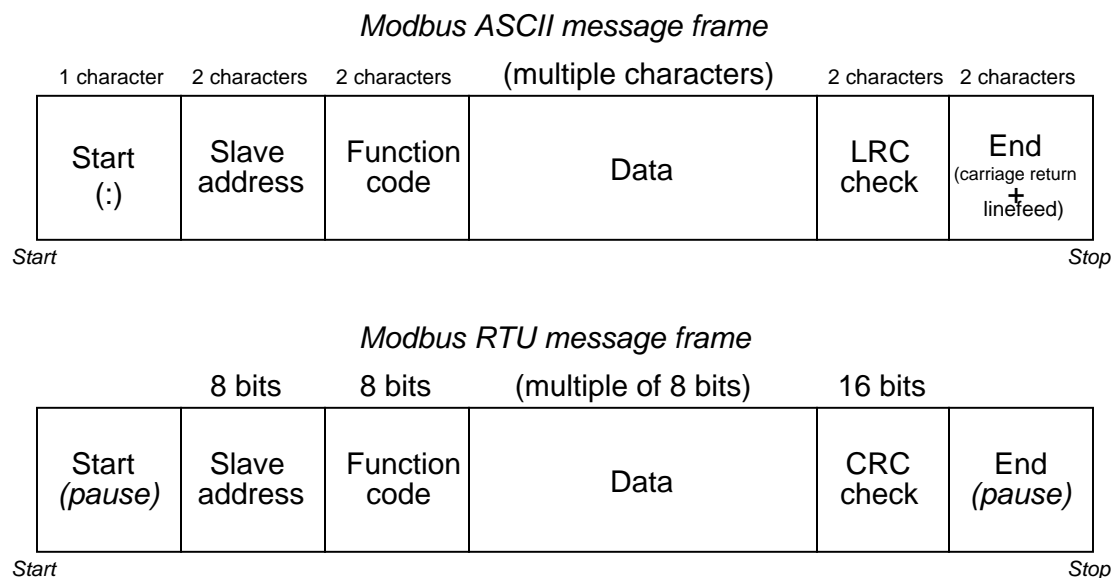
Modbus, especially when implemented over simple serial networks such as EIA/TIA-232 and EIA/TIA-485, is a rather primitive protocol. The seemingly arbitrary decimal codes used to issue commands and specify addresses is antiquated by modern standards. For better or for worse, though, a great many digital industrial devices “speak” Modbus, even if they are also capable of communicating via other network protocols. Using Modbus to communicate with modern control equipment is therefore an act of homage to 1970’s-era telecommunications: all participating devices in a Modbus network essentially behave the same as a 1970’s vintage Modicon PLC for the sake of exchanging information, even if their processing capabilities enable communications far more sophisticated than the Modbus protocol. A Modbus device querying another Modbus device does not “know” how modern or antiquated that other device is, because the basic Modbus standard has remained fixed for all this time.

The rest of this section explores details of the Modbus standard: its command vocabulary, addressing scheme, and some examples of read/write operations.

15.11.2 Modbus data frames

The Modbus communication standard defines a set of commands for reading (receiving) and writing (transmitting) data between a master device and one or more slave devices connected to the network. Each of these commands is referenced by a numerical code, with addresses of the master and slave devices' internal registers (data sources and data destinations) specified along with the function code in the Modbus frame.

Two different formats are specified in the Modbus standard: *ASCII* and *RTU*. The difference between these two modes is how addresses, function codes, data, and error-checking bits are represented. In Modbus ASCII mode, all slave device addresses, function codes, and data are represented in the form of ASCII characters (7 bits each), which may be read directly by any terminal program (e.g. `minicom`, `Hyperterminal`, `kermit`, etc.) intercepting the serial data stream. This makes troubleshooting easier: to be able to directly view the Modbus data frames in human-readable form. In Modbus RTU mode, all slave device addresses, function codes, and data are expressed in raw binary form. Different error-checking techniques are used for ASCII and RTU modes as well. The following diagram compares data frames for the two Modbus modes:



As you can see from a comparison of the two frames, ASCII frames require nearly twice⁷⁷ the number of bits as RTU frames, making Modbus ASCII slower than Modbus RTU for any given data

⁷⁷Recall that each ASCII character requires 7 bits to encode. This impacts nearly every portion of the Modbus data frame. Slave address and function code portions, for example, require 14 bits each in ASCII but only 8 bits each in RTU. The data portion of a Modbus ASCII frame requires one ASCII character (7 bits) to represent each hexadecimal symbol that in turn represents just 4 bits of actual data. The data portion of a Modbus RTU frame, by contrast, codes the data bits directly (i.e. 8 bits of data appear as 8 bits within that portion of the frame). Additionally, RTU data frames use quiet periods (pauses) as delimiters, while ASCII data frames use three ASCII characters in total to mark the start and stop of each frame, at a "cost" of 21 additional bits. These additional delimiting bits do serve a practical purpose, though: they format each Modbus ASCII data frame as its own line on the screen of a terminal program.

rate (bits per second).

The contents of the “Data” field vary greatly depending on which function is invoked, and whether or not the frame is issued by the master device or from a slave device. More details on Modbus “Data” field contents will appear in a later subsection.

Since Modbus is strictly a “layer 7” protocol, these message frames are usually embedded within other data frames specified by lower-level protocols. For example, the *Modbus TCP* standard encapsulates individual Modbus data frames as TCP/IP packets, which are then (usually) encapsulated again as Ethernet packets to arrive at the destination device. This “multi-layered” approach inherent to Modbus being such a high-level protocol may seem cumbersome, but it offers great flexibility in that Modbus frames may be communicated over nearly any kind of virtual and physical network type.

15.11.3 Modbus function codes and addresses

A listing of commonly-used Modbus function codes appears in the following table:

Modbus code (decimal)	Function
01	Read one or more PLC output “coils” (1 bit each)
02	Read one or more PLC input “contacts” (1 bit each)
03	Read one or more PLC “holding” registers (16 bits each)
04	Read one or more PLC analog input registers (16 bits each)
05	Write (force) a single PLC output “coil” (1 bit)
06	Write (preset) a single PLC “holding” register (16 bits)
15	Write (force) multiple PLC output “coils” (1 bit each)
16	Write (preset) multiple PLC “holding” registers (16 bits each)

Live data inside of any digital device is always located at some *address* within that device’s random-access memory (RAM). The Modbus “984” addressing standard defines sets of fixed numerical addresses where various types of data may be found in a PLC or other control device. The absolute address ranges (according to the Modbus 984 scheme) are shown in this table, with each address holding 16 bits of data:

Modbus codes (decimal)	Address range (decimal)	Purpose
01, 05, 15	00001 to 09999	Discrete outputs (“coils”), <i>read/write</i>
02	10001 to 19999	Discrete inputs (“contacts”), <i>read-only</i>
04	30001 to 39999	Analog input registers, <i>read-only</i>
03, 06, 16	40001 to 49999	“Holding” registers, <i>read/write</i>

Note how all the Modbus address ranges begin at the number one, not zero as is customary for so many digital systems. For example, a PLC with sixteen analog input channels numbered 0 through 15 by the manufacturer may “map” those input registers to Modbus addresses 30001 through 30016, respectively.

While this fixed addressing scheme was correct for the original PLCs developed by Modicon, it almost never corresponds directly to the addresses within a modern Modbus master or slave device. Manufacturer’s documentation for Modbus-compatible devices normally provide Modbus “mapping” references so technicians and engineers alike may determine which Modbus addresses refer to specific bit or word registers in the device. In some cases the configuration software for a Modbus-compatible device provides a utility where you may assign specific device variables to standard Modbus register numbers. An example of a Modbus variable mapping page appears in this screenshot taken from the configuration utility for an Emerson Smart Wireless gateway, used to “map” data from variables within *WirelessHART* radio-based field instruments to Modbus registers within the gateway device where other devices on a wired network may read that data:

The screenshot displays the 'Modbus Register Map' configuration page in the Emerson Smart Wireless Gateway software. The page features a navigation sidebar on the left with categories like Diagnostics, Setup, Network, Security, and HART. The main content area shows a table of Modbus registers mapped to device variables. The table has four columns: Register, Point Name, State, and Invert. The registers listed are 30001 (TT-101.PV), 30011 (TT-101.QV), 30021 (TT-101.ACTIVE_NEIGHBORS), 30031 (TT-101.BURST_178_RELIABILITY), 30033 (TT-101.RSSI), 30041 (LSL-78.SV), 30051 (TT-ORANGE.PV), and 30061 (TT-101.PV_STATUS). A warning 'Point does not exist' is displayed for register 30033. The interface includes a search bar, a 'New entry' button, and a 'Delete selected' button. The page is titled 'Modbus Register Map' and shows 'Page 1 of 1'.

Register	Point Name	State	Invert
<input type="checkbox"/> 30001	TT-101.PV		<input type="checkbox"/>
<input type="checkbox"/> 30011	TT-101.QV		<input type="checkbox"/>
<input type="checkbox"/> 30021	TT-101.ACTIVE_NEIGHBORS		<input type="checkbox"/>
<input type="checkbox"/> 30031	TT-101.BURST_178_RELIABILITY		<input type="checkbox"/>
<input type="checkbox"/> 30033	TT-101.RSSI		<input type="checkbox"/>
Point does not exist			
<input type="checkbox"/> 30041	LSL-78.SV		<input type="checkbox"/>
<input type="checkbox"/> 30051	TT-ORANGE.PV		<input type="checkbox"/>
<input type="checkbox"/> 30061	TT-101.PV_STATUS		<input type="checkbox"/>

As you can see here, the primary variable within temperature transmitter TT-101 (TT-101.PV) has been mapped to Modbus register 30001, where any Modbus master device on the wired network will be able to read it. Likewise, the secondary variable within level switch LSL-78 (LSL-78.SV) has been mapped to Modbus register 30041.

It is important to note that Modbus registers are 16 bits each, which may or may not exactly fit the bit width of the device variable in question. If the device variable happens to be a 32-bit floating point number, then *two* contiguous Modbus registers must be used to hold that variable, only the first of which will likely appear on the Modbus mapping page (i.e. the Modbus map will only show the *first* Modbus register of that pair). If the device variable happens to be a boolean (single bit), then it is likely only one bit within the 16-bit Modbus register will be used, the other 15 bits being “wasted” (unavailable) for other purposes. Details such as this may be documented in the manual for the device performing the Modbus mapping (in this case the Emerson Smart Wireless Gateway), or you may be forced to discover them by experimentation.

15.11.4 Modbus relative addressing

An interesting idiosyncrasy of Modbus communication is that the address values specified within Modbus data frames are *relative* rather than *absolute*. Since each Modbus read or write function only operates on a limited range of register addresses, there is no need to specify the entire address in the data frame. For example, Modbus function code 02 reads discrete input registers in the device with an absolute address range of 10001 to 19999 (i.e. all the addresses beginning with the digit “1”). Therefore, it is not necessary for the “read” command function 02 to specify the first digit of the register address. Instead, the read command only needs to specify a four-digit “relative address” specifying how far up from the beginning of the address range (in this case, from 10001) to go.

An analogy to aid your understanding of relative addressing is to envision a hotel building with multiple floors. The first digit of every room number is the same as the floor number, so that the first floor only contains rooms numbered in the 100’s, the second floor only contains rooms numbered in the 200’s, etc. With this very orderly system of room numbers, it becomes possible to specify a room’s location in more than one way. For example, you could give instructions to go to room 314 (an *absolute* room number), or alternatively you could specify that same room as “number 14 (a *relative* room number) on the third floor”. To a hotel employee who only works on the third floor, the shortened room number might be easier to remember.

In Modbus, relative addresses are just a little bit more complicated than this. Relative addresses actually span a range beginning at zero, while absolute addresses begin with “1” as the least-significant digit. This means there is an additional offset of 1 between a Modbus relative address and its corresponding absolute address. Returning to the hotel analogy, imagine the very first room on the third floor was room 301 (i.e. there was no room 300) and that the relative address represented the number of rooms *past* that first room. In this unintuitive scheme, room 314 could be specified as “the 13th room after the starting room on the third floor”. If this seems needlessly confusing, you are not alone. Welcome to Hotel Modbus.

A few examples are given here for illustration:

- Read the content of contact register 12440: *Modbus read function 02; relative address 2439*
- Read the content of analog input register 30050: *Modbus read function 04; relative address 49*
- Read the content of holding register 41000: *Modbus read function 03; relative address 999*
- Write multiple output coils in register 00008: *Modbus write function 15; relative address 7*

In each case, the pattern is the same: the relative address gets added to the first address of that range in order to arrive at the absolute address within the Modbus device. Referencing the first example shown above: 2439 (relative address) + 10001 (first address of register range) = 12440 (absolute address).

Thankfully, the only time you are likely to contend with relative addressing is if you program a computer using some low-level language such as assembly or C++. Most high-level industrial programming languages such as Function Block or Ladder Diagram make it easy for the end-user by allowing absolute addresses to be directly specified in the read and write commands. In a typical PLC program, for example, you would read contact register 12440 by simply specifying the number 12440 within the address field of a “read 02” instruction.

The following listing shows code (written in the C language) utilizing the open-source `libmodbus` function library instructing a computer to access 16-bit integer data from four Modbus “holding” registers (absolute addresses 49001 through 49004) via Modbus/TCP. The device’s IP address is 192.169.0.10 and port 502 is used for the TCP connection:

C code listing

```
#include <stdio.h>
#include <modbus.h>

modbus_t *Device;

int main (void)
{
    int read_count;
    uint16_t inreg_word [4];

    Device = modbus_new_tcp ("192.168.0.10", 502);
    modbus_set_error_recovery (Device, MODBUS_ERROR_RECOVERY_LINK);

    read_count = modbus_read_registers (Device, 9000, 4, inreg_word);

    printf("Number of registers read = %i\n", read_count);
    printf("Value of register 49001 = %i\n", inreg_word [0]);
    printf("Value of register 49002 = %i\n", inreg_word [1]);
    printf("Value of register 49003 = %i\n", inreg_word [2]);
    printf("Value of register 49004 = %i\n", inreg_word [3]);

    modbus_close (Device);
    modbus_free (Device);

    return read_count;
}
```

Note how the starting address passed to the read function is specified in relative form (9000), when in fact the desired absolute starting address inside the device is 49001. The result of running this code is shown here, the Modbus device in question being an Emerson Smart Wireless Gateway at 4:00 PM (i.e. 16:00 military time) on March 22, 2016. These four registers (49001 through 49004) happen to contain date and time information (year, month, day, and hour) stored in the device:

```
Number of registers read = 4
Value of register 49001 = 2016
Value of register 49002 = 3
Value of register 49003 = 22
Value of register 49004 = 16
```

This next listing shows similar code (also written in the C language⁷⁸) accessing 16-bit integer data from three Modbus “analog input” registers (absolute addresses 30015 through 30017) via Modbus/TCP from the same device as before:

C code listing

```
#include <stdio.h>
#include <modbus.h>

modbus_t *Device;

int main (void)
{
    int read_count;
    uint16_t inreg_word [3];

    Device = modbus_new_tcp ("192.168.0.10", 502);
    modbus_set_error_recovery (Device, MODBUS_ERROR_RECOVERY_LINK);

    read_count = modbus_read_input_registers (Device, 14, 3, inreg_word);

    printf("Number of registers read = %i\n", read_count);
    printf("Value of register 30015 = %i\n", inreg_word [0]);
    printf("Value of register 30016 = %i\n", inreg_word [1]);
    printf("Value of register 30017 = %i\n", inreg_word [2]);

    modbus_close (Device);
    modbus_free (Device);

    return read_count;
}
```

Note once again how the relative starting address specified in the code (14) maps to the absolute Modbus register address 30015, since analog input registers begin with the address 30001 and relative addresses begin at 0.

⁷⁸This C-language code is typed and saved as a plain-text file on the computer, and then a *compiler* program is run to convert this “source” code into an “executable” file that the computer may then run. The compiler I use on my Linux-based systems is `gcc` (the GNU C Compiler). If I save my Modbus program source code to a file named `tony_modbus.c`, then the command-line instruction I will need to issue to my computer instructing GCC to compile this source code will be `gcc tony_modbus.c -lmodbus`. The argument `-lmodbus` tells GCC to “link” my code to the code of the pre-installed `libmodbus` library in order to compile a working executable file. By default, GCC outputs the executable as a file named `a.out`. If I wish to rename the executable something more meaningful, I may either do so manually after compilation, or invoke the “outfile” option of `gcc` and specify the desired executable filename: (e.g. `gcc -o tony.exe tony_modbus -lmodbus`). Once compiled, the executable file may be run and the results of the Modbus query viewed on the computer’s display.

When using the `libmodbus` C/C++ library, the distinction between reading “analog input” registers (address range 30001 to 39999) and “holding” registers (address range 40001 to 49999) is made by the particular `libmodbus` function called. To read “analog input” registers in the 3XXXX address range, you use the `modbus_read_input_registers()` function. To read “holding” registers in the 4XXXX address range, you use the `modbus_read_registers()` function. This subtle difference in function names is important. Refer back to the two previous code examples to verify for yourself which function call is used in each of the register-reading applications.

15.11.5 Modbus function command formats

Every Modbus data frame, whether ASCII or RTU mode, has a field designated for “data.” For each Modbus function, the content of this “data” field follows a specific format. It is the purpose of this subsection to document the data formats required for common Modbus functions, both the “Query” message transmitted by the Modbus master device to a slave device, and the corresponding “Response” message transmitted back to the master device by the queried slave device.

Since each Modbus data frame is packaged in multiples of 8 bits (RTU), they are usually represented in text as individual bytes (two hexadecimal characters). For example, a 16-bit “word” of Modbus data such as 1100100101011011 would typically be documented as C9 5B with a deliberate space separating the “high” (C9) and “low” (5B) bytes.

Function code 01 – Read Coil(s)

This Modbus function reads the statuses of slave device discrete outputs (“coils”) within the slave device, returning those statuses in blocks of eight (even if the “number of coils” specified in the query is not a multiple of eight!). Relevant Modbus addresses for this function range from 00001 to 09999 (decimal) but the starting address is a hexadecimal number representing the $(n - 1)^{th}$ register from the beginning of this range (e.g. decimal address 00100 would be specified as hexadecimal 00 63).

Query message (Function code 01)

Start	Slave address XX	Function code 01	Data				Error check XX	End
			Starting address		Number of coils			
			Hi	Lo	Hi	Lo		

Start *Stop*

Response message (Function code 01)

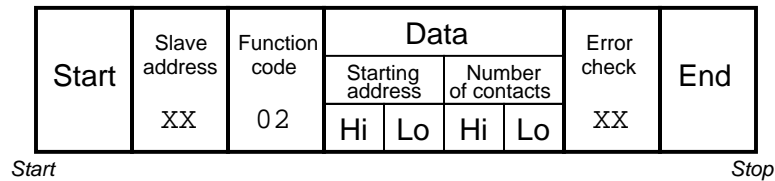

Start	Slave address XX	Function code 01	Data			Error check XX	End
			Number of bytes	First byte (8 coils)	Second byte (8 coils)		

Start *Stop*

Note that the second and third bytes representing coil status are shown in grey, because their existence assumes more than one byte worth of coils has been requested in the query.

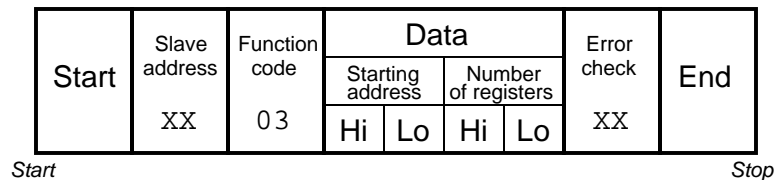

Function code 02 – Read Contact(s)

This Modbus function reads the statuses of slave device discrete inputs (“contacts”) within the slave device, returning those statuses in blocks of eight (even if the “number of contacts” specified in the query is not a multiple of eight!). Relevant Modbus addresses for this function range from 10001 to 19999 (decimal), but the starting address is a hexadecimal number representing the $(n - 1)^{th}$ register from the beginning of this range (e.g. decimal address 10256 would be specified as hexadecimal 00 FF).

Query message (Function code 02)**Response message (Function code 02)**

Function code 03 – Read Holding Register(s)

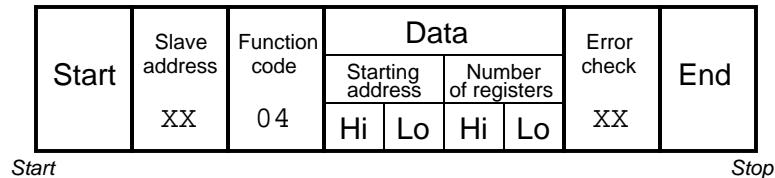

This Modbus function reads the statuses of “holding” registers within the slave device, with the size of each register assumed to be two bytes (16 bits). Relevant Modbus addresses for this function range from 40001 to 49999 (decimal), but the starting address is a hexadecimal number representing the $(n - 1)^{th}$ register from the beginning of this range (e.g. decimal address 40980 would be specified as hexadecimal 03 D3).

Query message (Function code 03)**Response message (Function code 03)**

Note that since the query message specifies the number of registers (each register being two bytes in size), and the response message replies with the number of *bytes*, the response message’s “number of bytes” field will have a value twice that of the query message’s “number of registers” field. Note also that the maximum number of registers which may be requested in the query message (65536) with “high” and “low” byte values grossly exceeds the number of bytes the response message can report (255) with its single byte value.

Function code 04 – Read Analog Input Register(s)

This Modbus function is virtually identical to 03 (Read Holding Registers) except that it reads “input” registers instead: addresses 30001 through 39999 (decimal). As with all the Modbus relative addresses, the starting address specified in both messages is a hexadecimal number representing the $(n - 1)^{th}$ register from the beginning of this range (e.g. decimal address 32893 would be specified as hexadecimal 0B 4C).

Query message (Function code 04)**Response message (Function code 04)**

Note that since the query message specifies the number of registers (each register being two bytes in size), and the response message replies with the number of *bytes*, the response message’s “number of bytes” field will have a value twice that of the query message’s “number of registers” field. Note also that the maximum number of registers which may be requested in the query message (65536) with “high” and “low” byte values grossly exceeds the number of bytes the response message can report (255) with its single byte value.

Function code 05 – Write (Force) Single Coil

This Modbus function writes a single bit of data to a discrete output (“coil”) within the slave device. Relevant Modbus addresses for this function range from 00001 to 09999 (decimal) but the starting address is a hexadecimal number representing the $(n - 1)^{th}$ register from the beginning of this range (e.g. decimal address 07200 would be specified as hexadecimal 1C 1F).

Query/Response message (Function code 05)

Start	Slave address	Function code	Data				Error check	End
			Coil address		Force data			
			Hi	Lo	Hi	Lo		
	XX	05					XX	

Start *Stop*

The “force data” for a single coil consists of either 00 00 (force coil off) or FF 00 (force coil on). No other data values will suffice – anything other than 00 00 or FF 00 will be ignored by the slave device.

A normal response message will be a simple echo (verbatim repeat) of the query message.

Function code 06 – Write (Preset) Single Holding Register

This Modbus function writes data to a single “holding” register within the slave device. Relevant Modbus addresses for this function range from 40001 to 49999 (decimal) but the starting address is a hexadecimal number representing the $(n - 1)^{th}$ register from the beginning of this range (e.g. decimal address 40034 would be specified as hexadecimal 00 21).

Query/Response message (Function code 06)



Start	Slave address	Function code	Data				Error check	End
			Register address		Preset data			
			Hi	Lo	Hi	Lo		
	XX	06					XX	

Start *Stop*

A normal response message will be a simple echo (verbatim repeat) of the query message.

Function code 15 – Write (Force) Multiple Coils

This Modbus function writes multiple bits of data to a set of discrete outputs (“coils”) within the slave device. Relevant Modbus addresses for this function range from 00001 to 09999 (decimal) but the starting address is a hexadecimal number representing the $(n - 1)^{th}$ register from the beginning of this range (e.g. decimal address 03207 would be specified as hexadecimal 0C 86).

Query message (Function code 15)**Response message (Function code 15)**

Note that the query message specifies both the number of coils (bits) and the number of bytes.

Function code 16 – Write (Preset) Multiple Holding Register

This Modbus function writes multiple words of data to a set of “holding” registers within the slave device. Relevant Modbus addresses for this function range from 40001 to 49999 (decimal) but the starting address is a hexadecimal number representing the $(n - 1)^{th}$ register from the beginning of this range (e.g. decimal address 47441 would be specified as hexadecimal 1D 10).

Query message (Function code 16)

Start	Slave address XX	Function code 10	Data								Error check XX	End
			Starting address		Number of registers		Number of bytes	Preset data first register		Preset data second register		
			Hi	Lo	Hi	Lo		Hi	Lo			

Start Stop

Response message (Function code 16)

Start	Slave address XX	Function code 10	Data				Error check XX	End
			Starting address		Number of registers			
			Hi	Lo	Hi	Lo		

Start Stop

Note that the query message specifies both the number of registers (16-bit words) and the number of bytes, which is redundant (the number of bytes must *always* be twice the number of registers, given that each register is two bytes⁷⁹ in size). Note also that the maximum number of registers which may be requested in the query message (65536) with “high” and “low” byte values grossly exceeds the number of bytes the response message can report (255) with its single byte value.

⁷⁹Even for devices where the register size is less than two bytes (e.g. Modicon M84 and 484 model controllers have 10 bits within each register), data is still addressed as two bytes' worth per register, with the leading bits simply set to zero to act as placeholders.

15.11.6 Floating-point values in Modbus

The ANSI/IEEE 754-1985 standard for floating-point data representation specifies 32 bits for the most basic (“single-precision”) form of floating-point number. Modbus input and holding registers, however, are specified by the Modbus standard to be 16 bits each. Therefore, if we wish to read or write a floating-point value within a Modbus slave device, we must issue a Modbus command to read or write *two* 16-bit registers representing the one 32-bit floating-point value. The task of piecing together two 16-bit numbers into a single 32-bit number, or splitting apart one 32-bit number into two 16-bit numbers, is left to the master device. If the master device in question is a modern PLC or HMI unit, this 16/32 bit conversion is most likely handled by the Modbus read/write instruction, so that all you must do is specify the first Modbus address for the pair of registers and the read/write instruction takes care of all other details. If, however, you are programming a computer using a low-level language, you must contend with these details in your own code.

A significant problem here is a lack of standardization among Modbus device manufacturers regarding exactly how 32-bit floating-point numbers are to be split up into two 16-bit register values. Some manufacturers simply take the 32 bits of the floating-point number and break them up into two sequential 16-bit numbers in order (denoted “ABCD” ordering, with each letter representing one byte of the original 32-bit floating-point number). Others reverse the order of the first and second 16-bit pieces (i.e. “CDAB” byte ordering). Others yet treat the 32-bit floating-point value as a set of four independent bytes which may be shuffled in any of several different orderings (e.g. “BADC”, “DCBA”, etc.). The Modbus standard offers no guidance on this matter, leaving the choice up to device manufacturers.

When programming in the C or C++ computer languages, a convenient strategy for splicing or splitting these different bit-length numbers is to make use of the *union* structure. A “union” in these languages is a reserved space in the computer’s memory which may be addressed by elements of different bit-length. For example, the following snippet of C code shows how to declare a union called `junk` which is 32 bits in size, and which may be addressed as a single 32-bit floating-point value called `junk.fp`, or as two 16-bit integer values called `junk.intg[0]` and `junk.intg[1]`, or as four 8-bit values called `junk.by[0]` through `junk.by[3]`:

C code listing

```
union {
    float fp;
    uint16_t intg[2];
    uint8_t by[4];
} junk;
```

This union could be written with 32 bits of data (in one step, as a floating-point value) and then read as either two 16-bit values and/or as four 8-bit values. The union structure gives any software developer the ability to reference a common area of computer memory as different number types.

The following code is a complete program reading two consecutive 16-bit “analog input” registers at addresses 30020 and 30021 and combining them into a single 32-bit floating-point value with “CDBA” ordering. Both the original 16-bit register values as well as the final floating-point value are displayed on the computer’s screen upon execution:

C code listing

```

#include <stdio.h>
#include <modbus.h>

modbus_t *Device;

int main (void)
{
    int read_count;

    union {
        uint16_t word[2];
        uint8_t byte[4];
    } in;

    union {
        float real;
        uint8_t byte[4];
    } out;

    Device = modbus_new_tcp ("192.168.0.10", 502);
    modbus_set_error_recovery (Device, MODBUS_ERROR_RECOVERY_LINK);

    read_count = modbus_read_input_registers (Device, 19, 2, in.word);

    printf("Value_of_16-bit_register_30020_=%i_\n", in.word[0]);
    printf("Value_of_16-bit_register_30021_=%i_\n", in.word[1]);

    out.byte[0] = in.byte[2];
    out.byte[1] = in.byte[3];
    out.byte[2] = in.byte[0];
    out.byte[3] = in.byte[1];

    printf("Value_of_32-bit_floating-point_number_=%f_\n", out.real);

    modbus_close (Device);
    modbus_free (Device);

    return read_count;
}

```

This program utilizes a pair of 32-bit unions (one called `in` and the other called `out`) to do the byte-swapping. First, the two 16-bit registers read by the `modbus_read_input_registers()` function are stored in the `in` structure as two 16-bit “words” addressed `in.word[0]` and `in.word[1]`. Then, those two words’ worth of data are addressed as four bytes, each one written to a different place within the `out` union by the four assignment statements. Note how `out.byte[0]` is assigned the value stored within `in.byte[2]` and so on: this is how the CDBA ordering is specified. One could specify ABCD ordering or DCBA ordering or any other combination of those four bytes by assigning the four `out` bytes to values of different `in` bytes, and the code would be just as straightforward to understand.

If you are fortunate enough, the Modbus library you are using will come complete with functions designed to take pairs of 16-bit registers and convert them into single floating-point numbers. At the time of this writing (2016), the free `libmodbus` library offers such functions. One of those functions (`modbus_get_float()`) is shown here for illustrative purposes, reading the contents of analog input registers 32999 and 33000 and displaying the converted (“ABCD”-ordered) floating-point value:

C code listing

```
#include <stdio.h>
#include <modbus.h>

modbus_t *Device;

int main (void)
{
    int read_count;
    uint16_t word[2];
    float real;

    Device = modbus_new_tcp ("192.68.0.10", 502);
    modbus_set_error_recovery (Device, MODBUS_ERROR_RECOVERY_LINK);

    read_count = modbus_read_input_registers (Device, 2998, 2, word);

    printf("Value of 16-bit register 32999 = %i \n", word[0]);
    printf("Value of 16-bit register 33000 = %i \n", word[1]);

    real = modbus_get_float(word);

    printf("Value of 32-bit floating-point number = %f \n", real);

    modbus_close (Device);
    modbus_free (Device);

    return read_count;
}
```

Sadly, you will often find that the arbitrary byte ordering used by any particular Modbus slave device manufacturer is poorly documented, if at all. This means you may have to experiment with different byte orderings before achieving success reading or writing floating-point data. Some Modbus device manufacturers are thoughtful enough to actually provide configurable options for how their 32-bit floating-point values will be represented within the slave device. Likewise, many Modbus master device manufacturers (e.g. PLCs and HMIs) provide options for how to read and write floating-point values in their Modbus read and write instructions, because they do not know whose Modbus slave device you may be communicating with.

15.12 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Analog vs. digital signals:** analog signals have infinite resolution but are susceptible to corruption by noise. Digital signals have limited resolution but are tolerant of any noise measuring less than the difference in thresholds between the high and low states.
- **Superposition theorem:** any linear, bilateral electrical network with multiple sources may be analyzed by taking one source at a time (while replacing all other sources with their internal impedance values) and analyzing all voltages and currents, then superimposing (summing) those voltage and current values to obtain the voltages and currents with all sources active. Relevant to analyzing DC and AC signals in HART instrument circuits.
- **Transmission lines:** short-duration (pulsed) electrical signals travel along a cable at nearly the speed of light, reflecting off the end of that cable if not properly terminated. Relevant to signal cables carrying high-frequency signals.
- **Self-balancing opamp circuits:** all self-balancing operational amplifier circuits work on the principle of negative feedback maintaining a nearly zero differential input voltage to the opamp. Making the "simplifying assumption" that the opamp's differential input voltage is exactly zero assists in circuit analysis, as does the assumption that the input terminals draw negligible current.

References

“422 and 485 Standards Overview and System Configurations” Application Report SLLA070C, Texas Instruments Incorporated, Dallas, TX, 2002.

“B&B Converters for the Industrial Bus World” Technical Article 13, B&B Electronics Manufacturing Company, Ottawa, IL, 2000.

Floyd, Thomas L., *Digital Fundamentals*, 6th edition, Prentice-Hall, Inc., Upper Saddle River, NJ, 1997.

“FOUNDATION Fieldbus System Engineering Guidelines” (AG 181) Revision 2.0, The Fieldbus Foundation, 2004.

“FOUNDATION Specification System Architecture” (FF 581) Revision FS 1.1, The Fieldbus Foundation, 2000.

“Fundamentals of RS-232 Serial Communications” Application Note 83 (AN83), Maxim Integrated Products, 2001.

Giancoli, Douglas C., *Physics for Scientists & Engineers*, Third Edition, Prentice Hall, Upper Saddle River, NJ, 2000.

Graham, Frank D., *Audels New Electric Library, Volume IX*, Theo. Audel & Co., New York, NY, 1942.

HART Communications, Technical Information L452 EN; SAMSON AG, 1999.

Hecker, Chris, “Let’s Get To The (Floating) Point”, *Game Developer*, pp. 19-24, February/March 1996.

Horak, Ray, *Telecommunications and Data Communications Handbook*, John Wiley & Sons, Inc., New York, NY, 2007.

Horak, Ray, *Webster’s New World Telecom Dictionary*, Wiley Publishing, Inc., Indianapolis, IN, 2008.

Hutchinson, Chuck, *The ARRL Handbook For Radio Amateurs*, 2001 edition, The American Radio Relay League, CT, 2000.

Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Software and Digital Networks*, Third Edition, CRC Press, New York, NY, 2002.

“Modbus Application Protocol Specification”, version 1.1b, Modbus-IDA, Modbus Organization, Inc., 2006.

“Modbus Messaging on TCP/IP Implementation Guide”, version 1.0b, Modbus-IDA, Modbus Organization, Inc., 2006.

“Modicon Modbus Protocol Reference Guide”, (PI-MBUS-300) revision J, Modicon, Inc. Industrial Automation Systems, North Andover, MA, 1996.

Newton, Harry, *Newton’s Telecom Dictionary*, CMP Books, San Francisco, CA, 2005.

Overton, Michael L., “Floating Point Representation”, 1996.

Park, John; Mackay, Steve; Wright, Edwin; *Practical Data Communications for Instrumentation and Control*, IDC Technologies, published by Newnes (an imprint of Elsevier), Oxford, England, 2003.

Polybius, *Histories*, cited from pages 49-50 of *Lapham’s Quarterly* Volume IX, Number 1, Winter 2016, American Agora Foundation, New York, NY, 2016.

Postel, John, *Internet Protocol – DARPA Internet Program Protocol Specification*, RFC 791, Information Sciences Institute, University of Southern California, Marina Del Ray, CA, September 1981.

Postel, John, *Transmission Control Protocol – DARPA Internet Program Protocol Specification*, RFC 793, Information Sciences Institute, University of Southern California, Marina Del Ray, CA, September 1981.

Postel, John, *User Datagram Protocol*, RFC 768, Information Sciences Institute, University of Southern California, Marina Del Ray, CA, August 1980.

“Recommendation ITU-R M.1677 International Morse Code”, ITU Radiocommunication Assembly, 2004.

Rector, B.E. et al., *Industrial Electronics Reference Book*, Westinghouse Electric Corporation, John Wiley & Sons Inc., New York, NY, 1948.

“Selecting and Using RS-232, RS-422, and RS-485 Serial Data Standards” Application Note 723 (AN723), Maxim Integrated Products, 2000.

Skoog, Douglas A.; Holler, F. James; Nieman, Timothy A.; *Principles of Instrumental Analysis*, fifth edition, Saunders College Publishing, Harcourt Brace & Company, Orlando, Florida, 1998.

Smith, Steven W., *The Scientist and Engineer’s Guide to Digital Signal Processing*, California Technical Publishing, San Diego, CA, 1997.

Smith, W. W., *The “Radio” Handbook*, Sixth Edition, Radio Ltd., Santa Barbara, CA, 1939.

Spurgeon, Charles E., *Ethernet: The Definitive Guide*, O’Reilly Media, Inc., Sebastopol, CA, 2000.

Svacina, Bob, *Understanding Device Level Buses: A Tutorial*, InterlinkBT, LLC, Minneapolis, MN, 1998.

Weiguo, Lin, “Address Resolution Protocol (ARP), RFC 826 ”, College of Computing, CUC, 2009-

2013.

Welsh, Matt and Kaufman, Lar, *Running Linux*, Second Edition, O'Reilly & Associates, Sebastopol, CA, 1996.

Chapter 16

FOUNDATION Fieldbus instrumentation


FOUNDATION Fieldbus is a standard for digital field instrumentation enabling field instruments to not only communicate with each other digitally, but also to execute all continuous control algorithms (such as PID, ratio control, cascade control, feedforward control, etc.) traditionally implemented in dedicated control devices. In essence, FOUNDATION Fieldbus extends the general concept of a distributed control system (DCS) all the way to the field devices themselves. In this way, FOUNDATION Fieldbus sets itself apart as more than just another digital communication “bus” for industry – it truly represents a new way to implement measurement and control systems. This chapter is devoted to a discussion of FOUNDATION Fieldbus instrumentation, building on general concepts of digital data acquisition and communication previously explored in this book.

For brevity, “FOUNDATION Fieldbus” will be abbreviated as *FF* throughout the rest of this chapter.

This particular industrial network standard was first proposed as a concept in 1984, and officially standardized by the Fieldbus Foundation (the organization overseeing all FF standards and validation) in 1996. To date, adoption of FF has been somewhat slow, mostly limited to new construction projects. One of the “selling points” of FF is decreased installation time, which makes it a more attractive technology for brand-new installations than for retrofit projects.


16.1 FF design philosophy

To understand just how different FOUNDATION Fieldbus is from other digital instrument systems, consider a typical layout for a distributed control system (DCS), where all the calculations and logical “decisions” are made in dedicated *controllers*, usually taking the form of a multi-card “rack” with processor(s), analog input cards, analog output cards, and other types of I/O (input/output) cards:




Information is communicated in analog form between the DCS controllers and the field instruments. If equipped with the proper types of I/O cards, the DCS may even communicate digitally with some of the field instruments using *HART* protocol. This allows remote configuration and diagnostic testing of field instruments from the host system, or from anywhere along the cable when using a hand-held HART communicator.

It is even possible to build a control system around a DCS using all digital field instruments, using a protocol such as *Profibus PA* to exchange process variable (PV) and manipulated variable (MV) signals to and from the DCS controllers at digital speeds far exceeding that of HART:



Now, multivariable field instruments have the ability to quickly exchange their data with the DCS, along with maintenance-related information (calibration ranges, error messages, and alarms). Each "fieldbus" cable is a multi-drop digital network, permitting multiple field devices per cable and consequently reducing total cable length and connection count. *Coupling devices* may be used in lieu of terminal blocks to conveniently connect multiple instruments together on common networks leading to the DCS. Still, however, all the automatic control algorithms are implemented in the DCS.

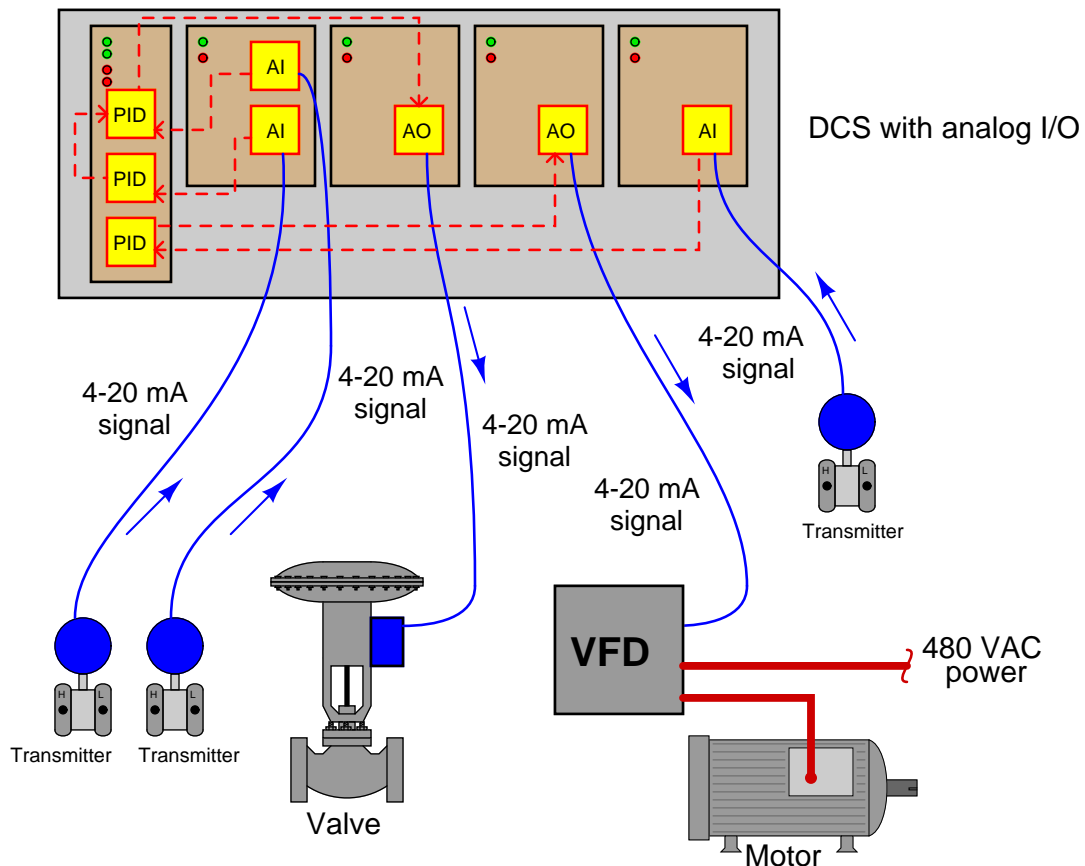
A FOUNDATION Fieldbus system goes one step further by allowing all control algorithms to be executed within the field instruments rather than relying on the DCS controllers to make automatic control “decisions.” In fact, the DCS would not even be necessary if not for the need of operations personnel to monitor and alter control system status:



That being said, it is possible (and in fact common) for control algorithms to be placed in the DCS controllers in addition to algorithms executed by FF field devices.


The locations of the control algorithms – those microprocessor instructions dictating how the loop will be controlled – in these different systems deserves further elaboration. To show this, I will make use of *function block* notation to show where the algorithms are executed in each system type, each function block shown as a yellow box on the diagram.

First, the DCS with analog I/O (inputs/outputs):



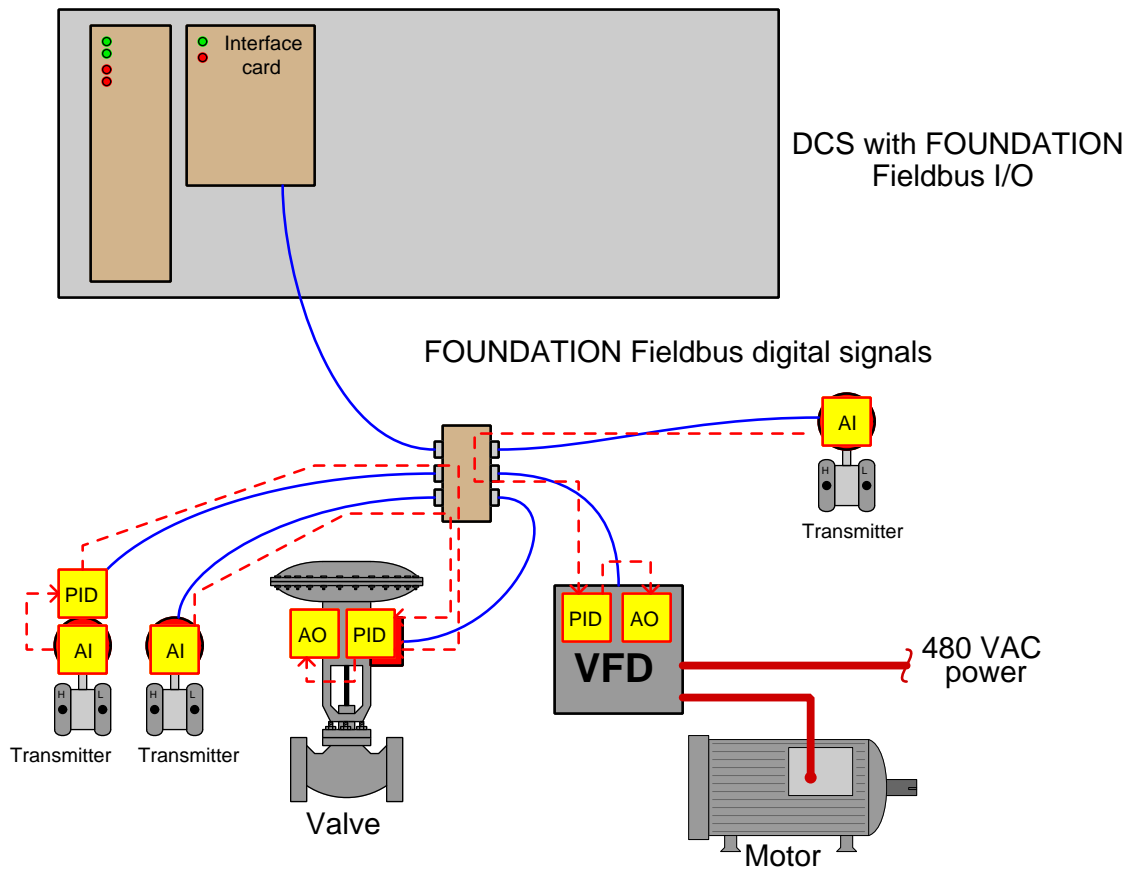
The conversion of 4-20 mA signals from transmitters into scaled digital number values inside the DCS takes place inside “analog input” (AI) function blocks programmed into the DCS. These converted values then pass to PID function blocks where the arithmetic for control loop decisions takes place. Finally, the digital output values of the PID blocks get passed on to “analog output” (AO) function blocks where those values are converted back into 4-20 mA analog signals to drive control valves, variable-frequency drives (VFDs), and other final control elements. Each “function block” is nothing more than a segment of programming code instructing the DCS’s microprocessor what to do with the signal values. Function blocks are usually selected and arranged by engineers and technicians maintaining the DCS using graphical programming software, allowing function blocks to be placed onto a “palette” and connected with lines to show where their signals come from and go to.

Now let us examine Profibus PA again. Here, the field instruments are entirely digital, communicating with each other via digital signals over a network cable to the DCS. This means none of the cables carry analog signals anymore, necessitating the A/D and D/A conversion take place inside the field devices themselves. It also means we may do away with the old analog I/O cards in the DCS rack, replacing them with a single network interface card:



Control decisions still take place within the DCS microprocessor, which is why the PID function blocks are still shown inside the processor card. The analog-digital signal conversions and scaling operations, however, occur within the field instruments themselves. Such is the nature of digital networks that multiple instruments may share the same communication cable back to the DCS, with each instrument “taking turns” communicating in time.

FOUNDATION Fieldbus, by contrast, allows even the control decisions to take place within the field devices, unburdening the DCS to perform higher-level tasks as necessary:



With each evolutionary step in system design, the trend has been to “push” the control algorithms further into the field, away from the central control system. FOUNDATION Fieldbus is the ultimate realization of this trend, where the field instruments themselves can do all necessary control functions. Here, the only necessary purposes served by the DCS are:

- Initial configuration and maintenance tool for the FF instruments
- Provide operators with an interface allowing indication and adjustment of control parameters
- Record long-term “historical” data on the process being controlled

In fact, given the right FF system design, the DCS could even be *disconnected* from the FF network, and the FF instruments would continue to control the process as they did before!

This is not to say that all control algorithms *must* be executed within the field instruments in a FOUNDATION Fieldbus control system. In fact it is quite common to find FF control systems implemented with the host (DCS) performing most of the control. FOUNDATION Fieldbus permits but does not mandate that all control tasks reside “in the field”.

When the FF standard was being designed, two different network levels were planned: a “low speed” network for the connection of field instruments to each other to form network *segments*, and a “high speed” network for use as a plant-wide “backbone” for conveying large amounts of process data over longer distances. The low-speed (field) network was designated *H1*, while the high-speed (plant) network was designated *H2*. Later in the FF standard development process, it was realized that existing Ethernet technology would address all the basic requirements of a high-speed “backbone,” and so it was decided to abandon work on the H2 standard, settling on an extension of 100 Mbps Ethernet called *HSE* (“High Speed Ethernet”) as the backbone FF network instead.

The bulk of this chapter will focus on H1 rather than HSE.

16.2 H1 FF Physical layer

Layer 1 of the OSI Reference Model is where we define the “physical” elements of a digital data network. The H1 FF network exhibits the following properties:


- Two-wire (ungrounded) network cable
- 100 ohm (typical) characteristic impedance
- DC power is conveyed over the same two wires as digital data
- 31.25 kbps data rate
- Differential voltage signaling (0.75 volts peak-to-peak transmit minimum ; 0.15 volts peak-to-peak receive threshold minimum)
- Manchester encoding

Since DC power is conveyed over the same two wires as the digital data, it means each device only needs to connect to two wires in order to function on an H1 network segment. The choice of a (relatively) slow 31.25 kbps data rate allows for imperfect cables and terminations which would otherwise plague a faster network. Manchester encoding embeds the network clock pulse along with the digital data, simplifying synchronization between devices.

As you can see, the layer 1 design parameters were chosen to make FF H1 networks easy to build in unforgiving industrial environments. The physical layer of FOUNDATION Fieldbus happens to be identical to that of Profibus-PA, further simplifying installation by allowing the use of common network validation tools and connection hardware.

16.2.1 Segment topology

A minimal FF H1 segment consists of a DC power supply, a “power conditioner,” exactly two terminator resistors¹ (one at each extreme end of the cable), a shielded and twisted-pair cable, and of course at least two FF instruments to communicate with each other. The cable connecting each instrument to the nearest junction is called a *spur* (or sometimes a *stub* or a *drop*), while the cable connecting all junctions to the main power source (where a host DCS would typically be located) is called a *trunk* (or sometimes a *home run* for the section leading directly to a host system):



The power conditioner shown in this diagram is a simplified model of the actual device, the function of which being to filter out digital data pulses from reaching the DC power supply. Commercially-available Fieldbus power conditioners are complex electronic circuits rather than passive filter networks.


Normally, we would find more than two FF devices connected to a trunk cable, as well as a “host” system such as a DCS FF card for accessing FF instrument data, performing maintenance tasks, and integrating with other control loops. Regardless of how many (or how few) FF devices connect to an H1 segment, though, there should always be *exactly two* terminating resistors in each segment – one at each end² of the trunk cable. These resistor/capacitor networks serve the sole purpose of eliminating signal reflections off the ends of the trunk cable, making the cable look infinitely long from the perspective of the propagating pulse signals. Missing terminators will result in signal

¹Each FF terminator resistor is actually a series resistor/capacitor network. The capacitor blocks direct current, so that the 100 Ω resistor does not impose a DC load on the system. The substantial current that would be drawn by a 100 ohm resistor across 24 VDC source if not blocked by a series capacitor ($24\text{ V} / 100\text{ ohms} = 240\text{ mA}$) would not only waste power (nearly 6 watts per resistor!) but that much current would cause an unnecessary degradation of supply voltage at the field device terminals due to voltage drop along the length of the segment cable’s conductors.

²Be sure to check the specifications of the host system H1 interface card, because many are equipped with internal terminating resistors given the expectation that the host system will connect to one far end of the trunk!

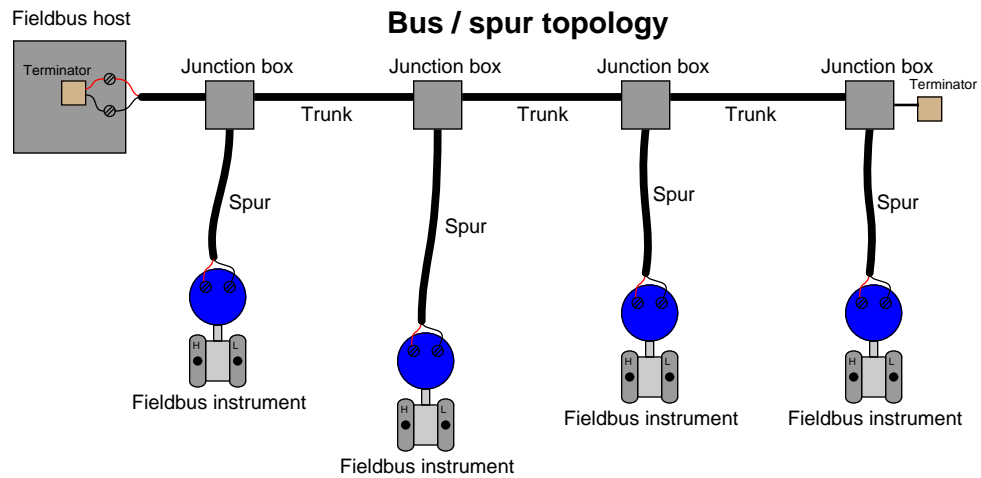
reflections off the unterminated line end(s), while extra terminators have the equally deleterious effect of attenuating signal strength (as well as potentially causing signal reflections of opposite phase).

All H1 networks are essentially parallel electrical circuits, where the two connection terminals of each field instrument are paralleled to each other. The physical arrangement of these transmitters, though, may vary substantially. The simplest way to connect FF H1 devices together is the so-called “daisy-chain” method, where each instrument connects to two cable lengths, forming an uninterrupted “chain” network from one end of the segment to the other:




As simple as this topology is, it suffers from a major disadvantage: it is impossible to disconnect any device in the segment without interrupting the network’s continuity. Disconnecting (and reconnecting for that matter) any device necessarily results in all “downstream” devices losing signal, if only for a brief time. This is an unacceptable liability in most industrial applications, as it complicates maintenance and servicing of individual instruments on the segment.

An alternative topology is the *bus* layout, where short “spur” cables connect instruments to a longer “trunk” cable. Terminal blocks – or even quick-disconnect couplings – within each junction box provide a convenient means of disconnecting individual devices from the segment without interrupting data communication with the other devices:



The ideal arrangement for a “bus” network is to minimize the length of each spur cable, so as to minimize the delay of reflected signals off the unterminated ends of the drops. Remember that only *two* termination resistors are allowed in any electrically continuous network segment, and so this rule forbids the addition of terminators to the end of each spur cable.

Yet another alternative topology for H1 networks is the so-called *chicken-foot* arrangement, where a long trunk cable terminates at a multi-point junction along with several field devices and their spur cables:



Most FF systems resemble a combination of “bus” and “chicken-foot” topologies, where multiple junction devices serve as connection points for two or more field instruments per junction.

16.2.2 Coupling devices

In order to simplify the task of connecting Fieldbus devices to such a network segment, multiple manufacturers sell *coupling devices* (often informally referred to as *bricks*) with quick-disconnect electrical fittings so the end-user does not have to build and commission junction boxes using standard terminal blocks. A photograph of a Turck brand Fieldbus coupling device appears here, showing multiple spur cables plugged into it:



Coupling devices are highly recommended for all industrial fieldbus systems, FF or otherwise. Not only do these devices provide a convenient means of forming highly reliable connections between field instruments and the trunk cable, but many of them are equipped with features such as short-circuit protection (so that a shorted spur cable or field instrument does not cause the entire segment to stop communicating) and LED indication of spur status.

Cables connecting to a coupling device must be equipped with special plugs matching the sockets on the coupler. This presents a bit of a problem when attempting to pull such a cable through electrical conduit: the bulky plug requires either over-sized conduit to accommodate the plug's width, or requires the plug be installed on the cable after pulling through the conduit. Both approaches are expensive, the first in terms of capital cost and the second in terms of installation labor. For this reason, many installers abandon electrical conduit altogether in favor of *ITC* ("Instrument Tray Cable").

A wider-angle photograph of the coupling device previously shown reveals many ITC cables and their routing through wire “basket” style trays among process instruments and vessels:




As evident in this photograph, ITC is obviously rated for continuous exposure to direct sunlight and moisture, as well as a certain amount of physical distress (abrasion, high and low temperatures, etc.). Article 727 of the National Electrical Code (NEC) defines the acceptable uses and installations of ITC³.

It should be noted that while a properly shielded and grounded FF cable is quite resistant to radio-frequency interference, coupling devices may present “weak spots” where radio interference may find its way onto the segment. Different styles of coupling devices offer differing levels of immunity to RF (Radio Frequency) noise. Those made of metal and properly bonded to ground will be well-shielded, while those made of plastic having exposed connection terminals offer little or no protection. In any case, it is a good practice to avoid “keying” any portable radio transmitter in the near vicinity of a Fieldbus coupling device.

³You should consult an NEC code book regarding specific limitations of ITC wiring. Some of the main points include limiting individual ITC cable lengths to a maximum of 50 feet, and mechanically securing the cable at intervals not to exceed 6 feet.

Not all Fieldbus couplers are rated for outdoor installation. Some are intended for mounting inside electrical enclosures, such as this Pepperl+Fuchs model shown mounted on a DIN rail:




This Fieldbus coupling device is aptly labeled a *segment protector*, for it not only couples spurs to the main trunk of the Fieldbus segment, but it also guards against short-circuits in the spur cables and devices from interrupting communication on the rest of the segment. If you look closely at the upper-left of the coupling device, you will see a black plastic square with two leads inserted into screw terminals: this is one of two *terminating resistors* found in this Fieldbus segment, meaning this particular coupling device is at the “end of the line” of the network segment.

Not only do enclosure-protected coupling devices eliminate the need for special weather-proof connectors and instrument tray cable, but they also enjoy the radio interference immunity⁴ granted by being inside a metal cocoon.

⁴Provided the metal enclosure’s door is left in the closed position at all times! Keying a radio transmitter near such a coupling device while the enclosure door is open invites trouble.

16.2.3 Electrical parameters

FOUNDATION Fieldbus H1 networks use Manchester encoding to represent bit states: a “high-to-low” transition represents a logical zero (0), while a “low-to-high” transition represents a logical one (1). The following illustration shows how the data stream 00100 would be represented in Manchester encoding:



FF devices must be able to correctly distinguish between rising- and fall-edge signals in order to properly interpret the bit states of a Manchester-encoded signal. Any device interpreting these pulse edges “backwards” will invert every single bit! Thankfully, this problem is easy to avoid because the DC power supplied by the H1 segment wiring provides a “key” to identifying which wire is which, and therefore which pulses are rising-edge versus which pulses are falling-edge. For this reason, many (but not all!) FF devices are polarity-insensitive, automatically detecting the polarity of the network segment and compensating accordingly.

Every FF device draws at least 10 mA of current from the segment, and this current does not vary in the same manner that an analog (4-20 mA) device draws differing amounts of current under different operating conditions. Always remember that a Fieldbus device signals its variable(s) digitally, not by varying current. Old habits (and thought patterns) die hard, and so Fieldbus systems present challenges to technicians familiar with the behavior of analog current loop instrumentation. The amount of current drawn by any particular FF device depends on that device’s functionality – obviously, some will require more current⁵ for their operation than others. 10 mA to 30 mA should

⁵Perusing documentation on an assortment of Emerson/Rosemount FF products, I found the following data: model 752 indicator = 17.5 mA, model 848L logic = 22 mA, model 848T temperature = 22 mA maximum, model 3244MV temperature = 17.5 mA typical, model DVC6000f valve positioner = 18 mA maximum, model 848L logic = 22 mA,

be considered a general range of current drawn by each FF device.

The standard operating voltage range for FF devices is between 9 and 32 volts DC. It is important to note, however, that not all manufacturers' devices are in full compliance with the Fieldbus Foundation standard, and as such some may not operate properly at low voltages (near 9 volts DC)! The most common DC operating voltage for a FF network segment is 24 VDC (typical).

The minimum transmission voltage of a FF device is 750 millivolts peak-to-peak, while the minimum signal level for reception by a FF device is 150 millivolts peak-to-peak. This represents an acceptable attenuation of 5:1, or -14 dB between any two devices.


model 848T temperature = 22 mA maximum, model 3244MV temperature = 17.5 mA typical, model 5500 guided-wave radar level = 21 mA, model 3095MV flow (differential pressure) = 17 mA approximate, model DVC6000f valve positioner = 18 mA maximum.

16.2.4 Cable types


Fieldbus cable is rated according to a four-level code (A, B, C, or D), each successive letter representing a cable of lower quality⁶. The following table gives *minimum specifications* for each FF cable type:

Cable Type	Type A	Type B	Type C	Type D
Wire size	AWG 18	AWG 22	AWG 26	AWG 16
Char. Impedance	$100 \Omega \pm 20\%$	$100 \Omega \pm 30\%$	–	–
Shielding	1 for each pair	1 for entire cable	none	none
Twisted pairs	Yes	Yes	Yes	No
Max. length	1900 m	1200 m	400 m	200 m

Bear in mind that the maximum length given for each cable type is the *total length* of all cables in a segment, trunk length plus all spur lengths. As a general rule, spur lengths should be kept as short as possible. It is better to route the trunk cable in a serpentine fashion to locate coupling devices close to their respective instruments than it is to streamline the trunk cable routing. The following illustrations contrast the two approaches:



⁶I have successfully built several “demonstration” FF systems using cables of questionable quality, including lamp (“zip”) cord, with no termination resistors whatsoever! If the distances involved are short, just about any cable type or condition will suffice. When planning the installation of any real Fieldbus installation, however, you should never attempt to save money by purchasing lesser-grade cable. The problems you will likely encounter as a consequence of using sub-standard cable will more than offset the initial cost saved by its purchase.




If greater lengths are required for a network segment, devices known as *repeaters* may be added which sense and re-broadcast the Manchester-encoded FF signal between trunk cables. A maximum of four repeaters may be used to extend any H1 segment.

As always, neat wiring practices help make an instrument system easier to maintain and to diagnose when things go wrong. The following photograph shows a triad of FOUNDATION Fieldbus junction boxes and (orange) network cables. Coupling devices located inside each enclosure link each spur cable to the trunk:



16.2.5 Segment design

In addition to maximum (total) cable length and repeater count, a host of other details⁷ conspire to limit how any particular H1 segment is wired. To help engineers and technicians alike deal with these details, manufacturers often provide free *segment design tool* software to pre-validate a segment design on computer before purchasing components and installing them in the field. A screenshot taken from Emerson's offering shows what a typical FF segment layout might look like:



A very nice feature of these segment design packages is their built-in database of FF components. Every time you “pick” a particular component to place in your simulated segment, the program references data for that device’s current draw and other electrical parameters relevant to the performance of the segment. Of course, each manufacturer will tend to feature their own devices more prominently, and so these software tools sometimes feel like a promotional advertisement. Despite the commercial aspect of their design, however, they are extremely useful in the planning stages of a FF network, and should be used whenever possible.

Another reason to use segment design tool software is to document the wiring of each FF segment. One of the casualties of the new Fieldbus paradigm is the traditional *loop diagram* (or “loop sheet”), the purpose of which is to document the signal wiring dedicated for each measurement and control

⁷Total device current draw, spur length versus number, intrinsic safety voltage and current limitations, etc.

loop. In FOUNDATION Fieldbus, the control “loop” is virtual rather than physical, being comprised of digital data sent between field instruments, the path of which being defined by the instruments’ programming. The only physical wiring entity to document in a FF system is the segment, and each segment most likely hosts more than one measurement and/or control loop. Unless and until a standardized documentation format⁸ is invented for Fieldbus network segments, the graphic image provided by segment design tool software is as good as anything.

⁸At the time of this writing (2009), the ISA has yet to standardize new methods of FF documentation in the style of loop sheets and P&IDs. This is one of those circumstances where technology has outpaced convention.

16.3 H1 FF Data Link layer

Like so many other industrial data networks, FOUNDATION Fieldbus is an “unswitched” or “broadcast” type of network. This means all data transmissions by all devices on a network are sensed by all the other devices. In other words, there are no private messages between two devices on a shared network: every device “hears” every transmission from every other device. This means devices must take turns communicating, with no simultaneous transmissions. Layer 2 of the OSI Reference Model is where we define the “data link” elements of a digital data network, describing how individual devices negotiate for the right to transmit on the network. Here is a list of some layer-2 properties of H1 FF networks:

- Master/slave network behavior for cyclic communications (i.e. one device polls the others, and the others merely respond)
- Delegated token network behavior for acyclic communications (i.e. devices serially granted time to broadcast at will)
- Dedicated “scheduler” device for coordinating all segment communications
- 8-bit address field (0 through 255 possible)
- Maximum of 32 “live” devices on a segment

On an operating H1 segment, one device called the *Link Active Scheduler* (abbreviated LAS) functions as the “master” device for coordinating all network communications, analogous to a police officer directing traffic in a road intersection. The LAS device may be a regular field instrument (e.g. transmitter, valve positioner) or it may be the host system (i.e. the H1 segment interface card of a DCS). The FF standard allows for one operating LAS device, with multiple back-up LAS devices waiting to take over if the primary LAS happens to fail for any reason.

One of the tasks of the LAS is to “compel” the various field instruments to transmit their process control data (process variables, PID control output values, and other variables essential for loop monitoring and control), while the devices immediately respond in answer to the LAS’s “compel data” command. These critical communications occur on a regular schedule, and therefore are referred to as *scheduled* or *cyclic* communications. Cyclic communication operates in a “master-slave” fashion, with the LAS acting as the master (commanding slave devices to broadcast specific data), and all other devices responding only when called upon by the LAS. This form of communication is analogous to a traffic policeman specifically directing one vehicle at a time to drive through an intersection in a prescribed manner.

Periods of time in between these critical transmissions on an H1 network are used for device’s internal processing (e.g. PID algorithm execution, diagnostic checking) and also for less-critical data transmission. It is during these *unscheduled* or *acyclic* times that devices are sequentially given permission by the LAS to broadcast data of less importance such as operator setpoints, PID tuning constant updates, alarm acknowledgments, and diagnostic messages. Acyclic communication operates in a manner similar to “token-passing,” with the LAS issuing time-limited tokens to the other devices in sequence permitting them to freely broadcast whatever other data they have to share. This form of communication is analogous to a traffic policeman directing an entire lane of vehicles to enter the intersection at will.

The scheduled nature of cyclic communication guarantees a certain maximum response time to critical control functions, an important property of control networks called *determinism*. Without determinism, a control system cannot be relied upon to perform critical regulatory functions in a timely⁹ manner, and sequencing¹⁰ of control functions such as PID, summers, subtractors, ratio multipliers, and the like may be compromised. Thus, all the critical variables of a FF H1 loop are communicated between devices this way.

⁹While many industrial control systems have been built using networks that are not strictly deterministic (e.g. Ethernet), generally good control behavior will result if the network latency time is arbitrarily short. Lack of “hard” determinism is more of a problem in safety shutdown systems where the system *must* respond within a certain amount of time in order to be effective in its safety function. An industrial example of a safety system requiring “hard” determinism is compressor surge control. An automotive example requiring “hard” determinism is anti-lock brake control.

¹⁰By “sequencing,” I mean the execution of all antecedent control functions prior to “downstream” functions requiring the processed data. If in a chain of function blocks we have some blocks lagging in their execution, other blocks relying on the output signals of those lagging blocks will be functioning on “old” data. This effectively adds dead time to the control system as a whole. The more antecedent blocks in the chain that lag in time behind the needs of their consequent blocks, the more dead time will be present in the entire system. To illustrate, if block *A* feeds data into block *B* which feeds data into block *C*, but the blocks are executed in reverse order (*C*, then *B*, then *A*) on the same period, a lag time of *three whole execution periods* will be manifest by the A-B-C algorithm.

16.3.1 Device addressing

FOUNDATION Fieldbus devices (also called *nodes*) are addressed by an eight-bit binary number when functioning on an H1 segment. This binary number field naturally supports a maximum addressing range of 0 to 255 (decimal), or 00 to FF hexadecimal. This address range is divided into the following sub-ranges by the Fieldbus Foundation:

Address range (decimal)	Address range (hexadecimal)	Allocation
0 through 15	00 through 0F	Reserved
16 through 247	10 through F7	Permanent devices
248 through 251	F8 through FB	New or decommissioned devices
252 through 255	FC through FF	Temporary (“visitor”) devices

Devices are usually assigned addresses to function on the segment by the host system (typically a DCS with FF capability), although it is possible to order FF instruments pre-configured at the factory with addresses specified by the customer upon order. Host systems are generally configured to automatically determine device addresses rather than require the technician or engineer to manually assign each address. This makes the commissioning process more convenient.

The maximum number of “permanent” devices (installed field instruments) allowed on an H1 segment for operational reasons is 32, and as you can see the addressing scheme offers far more valid addresses than that. One of the many tasks given to a segment’s Link Active Scheduler (LAS) device is to probe for new devices connected to the segment. This is done on a one-at-a-time basis, with the LAS sequentially polling for uncommissioned addresses within the valid address range. Obviously, this can be a waste of time with only 32 addresses capable of active service at any given time and over 200 valid address numbers. A practical solution to this problem is to specify an “unused” address range for the LAS to skip, so it does not waste time probing for devices (nodes) within a certain range. This address range is specified as a set of two numbers: one for the First Unused Node (abbreviated *FUN*), and another specifying the Number of Unused Nodes (abbreviated *NUN*). For example, if one wished to have the LAS on a particular H1 segment skip device addresses 40 through 211, one would configure the FUN to equal 40 and the NUN to equal 172, since the address range 40 through 211 is one hundred seventy two addresses (inclusive of both 40 and 211).

Even with a maximum operational limit of 32 devices to an H1 segment, it is rare to find segments operating with more than 16 devices. One reason for this is speed: with additional devices requiring time to broadcast and process data, the total *macrocycle* time (the time period between guaranteed delivery of the same process data from any one device – the determinism time) must necessarily increase. According to the Fieldbus Foundation’s engineering recommendations guide, there must be no more than twelve devices on a segment (including no more than two final control elements) in order to achieve a 1-second or less macrocycle time. For half-second update times, the recommended maximum is six devices (with no more than two final control elements). For quarter-second update times, the limit drops to a total of three devices, with no more than one final control element. Macrocycle time is essentially dead time, which is worse than lag time for any form of feedback control. When controlling certain fast processes (such as liquid pressure or flow rate), dead times on the order of one second are a recipe for instability.

Another limitation to the number of operational addresses on an H1 segment is current draw. FF devices draw 10 mA of current *minimum*. A FF segment with sixteen parallel-connected devices

would see a total current of 160 mA minimum, with a more realistic value being in excess of 300 mA.

In addition to network addresses, each FF device bears an absolutely unique identifier (a 32-byte binary number) to distinguish it from any other FF device in existence. This identifier serves much the same purpose as a *MAC address* on an Ethernet device. However, the identifier field for FF devices allows a far greater instrument count than Ethernet: 32 *bytes* for FF instruments versus 48 bits for Ethernet devices. While the Ethernet MAC address field only allows for a paltry 2.815×10^{14} unique devices, the FF identifier allows 1.158×10^{77} devices! The distinction between a FF device's network address and the device's identifier is virtually identical to the distinction between an Ethernet device's IP address assigned by the end-user and its MAC address number assigned by the manufacturer.

This identifier value is usually expressed as 32 ASCII-encoded characters for brevity (one alphanumeric character per byte), and is subdivided into byte groups as follows:

First 6 bytes	Middle 4 bytes	Last 22 bytes
Manufacturer code	Device type code	Serial number

For example, the identifiers for all *Fisher* brand devices begin with the first six characters 005100. The identifiers for all *Smar* devices begin with the characters 000302. The identifiers for all *Rosemount*¹¹ brand devices begin with 001151. A typical identifier (this particular one for a Fisher model DVC5000f valve positioner) appears here:

005100 0100 FISHERDVC0440761498160

Normally, these identifiers appear as 32-character strings, without spaces at all. I have inserted spaces within this string to make the character groupings easier to see.

16.3.2 Communication management

In a FF network segment, the Link Active Scheduler (LAS) device coordinates all communications between segment devices. Among the many responsibilities the LAS is tasked with are the following:

- Commands non-LAS devices to broadcast data to the segment with “Compel Data” (CD) messages, issued at regular time intervals to specific devices (one at a time)
- Grants permission for non-LAS devices to communicate with “Pass Token” (PT) messages, issued during unscheduled time slots to specific devices (one at a time, in ascending order of address number)
- Keeps all segment devices synchronized with a regular “Time Distribution” (TD) message
- Probes for new devices on the segment with a “Probe Node” (PN) message
- Maintains and publishes a list of all active devices on the network (the *Live List*)

¹¹The engineers there are not without a sense of humor, choosing for their manufacturer code the same model number as the venerable model 1151 differential pressure transmitter, perhaps the most popular Rosemount industrial instrument in the company's history!

Scheduled versus unscheduled communication

As previously mentioned, Fieldbus H1 network communication may be divided into two broad categories: *scheduled* (cyclic) and *unscheduled* (acyclic). Scheduled communication events are reserved for exchanging critical control data such as process variable measurements, cascaded setpoints, and valve position commands. These scheduled communications happen on a regular, timed schedule so that loop determinism is guaranteed. Unscheduled communications, by contrast, are the way in which all other data is communicated along an H1 segment. Manual setpoint changes, configuration updates, alarms, and other data transfers of lesser importance are exchanged between devices in the times between scheduled communication events.

Both forms of communication are orchestrated by the Link Active Scheduler (LAS) device, of which there is only one active at any given time¹² on an H1 segment. The LAS issues “token” messages to non-LAS devices commanding (or merely authorizing) them to broadcast to the segment one at a time. Each token message issued by the LAS grants transmission rights to an FF device either for a limited purpose (i.e. the precise message to be transmitted) or for a limited time (i.e. giving that device the freedom to transmit whatever data it desires for a short duration), after which transmission rights return to the LAS. CD tokens are message-specific: each one issued by the LAS commands a single device to immediately respond with a broadcast of some specific data. This is how scheduled (cyclic) communication is managed. PT tokens are time-specific: each one issued by the LAS grants a single device free time to transmit data of lesser importance. This is how unscheduled (acyclic) communication between devices is managed.

The LAS also issues a third type of token message: the “Probe Node” (PN) token intended to elicit a response from any new devices connected to the network segment. Probe Node tokens are issued one at a time to each uncommitted device address in search of any new devices.

In addition to transmitting tokens – which by definition are messages granting another device permission to transmit to the network – the LAS also broadcasts other messages necessary for the function of an H1 segment. For example, the “Time Distribution” (TD) message regularly broadcast by the LAS keeps all devices’ internal clocks synchronized, which is important for the coordinated transfer of data.

One of the “internal” tasks of the LAS not requiring network broadcasts is the maintenance of the *Live List*, which is a list of all known devices functioning on the network segment. New devices responding to “Probe Node” messages will be added to the Live List when detected. Devices failing to return or use PT tokens issued to them are removed from the Live List after a number of attempts. When “backup” LAS devices exist on the segment, the LAS also publishes updated copies of the Live List to them, so they will have the most up-to-date version should the need arise to take over for the original LAS (in the event of an LAS device failure).

In “busy” H1 segments where multiple devices are exchanging data with each other, a heavy traffic load of scheduled communications (CD tokens and their responses) makes it difficult for substantial unscheduled (acyclic) data exchanges to occur. For example, if a device happens to be maintaining a lengthy list of client/server requests in its queue, which it may address only during its allotted acyclic time slots (i.e. when it has been given the PT token from the LAS), it is quite possible the PT token

¹²In addition to the main LAS, there may be “backup” LAS devices waiting ready to take over in the event the main LAS fails for any reason. These are Link Master devices configured to act as redundant Link Active Schedulers should the need arise. However, at any given time there will be only *one* LAS. In the event of an LAS device failure, the Link Master device with the lowest-number address will “step up” to become the new LAS.

will expire before all the device's transactions have been completed. This means the device will have to wait for the next acyclic period before it can complete all the unscheduled communication tasks in its queue. The Fieldbus Foundation recommends new H1 segments be configured for no more than 30% scheduled communications during each macrocycle (70% unscheduled time). This should leave plenty of "free time" for all necessary acyclic communications to take place without having to routinely wait multiple macrocycles.

Virtual Communication Relationships

A term you will frequently encounter in FF literature is *VCR*, or “Virtual Communication Relationships.” There are three different types of VCRs in FF, describing three different ways in which data is communicated between FF devices:

- **Publisher/Subscriber** (scheduled), otherwise known as Buffered Network-scheduled Unidirectional (BNU)
- **Client/Server** (unscheduled), otherwise known as Queued User-triggered Bidirectional (QUB)
- **Source/Sink** (unscheduled), otherwise known as Queued User-triggered Unidirectional (QUU)

Publisher/Subscriber: this VCR describes the action of a Compel Data (CD) token. The Link Active Scheduler (LAS) calls upon a specific device on the network to transmit specific data for a time-critical control purpose. When the addressed device responds with its data, multiple devices on the network “subscribing” to this published data receive it simultaneously. This is how process-control variables (PV, PID output, etc.) are communicated between instruments comprising a FF control loop. The publisher/subscriber VCR model is highly deterministic because all such communications occur on a precisely defined schedule.

Client/Server: this VCR describes one class of unscheduled communications, permitted when a device receives a Pass Token (PT) message from the LAS. Each device maintains a queue (list) of data requests issued by other devices (clients), and responds to them in order as soon as it receives the Pass Token. By responding to client requests, the device acts as a server. Likewise, each device can use this time to act as a client, posting their own requests to other devices, which will act as servers when they receive the PT token from the LAS. This is how non-critical messages such as maintenance and device configuration data, operator setpoint changes, alarm acknowledgments, PID tuning values, etc. are exchanged between devices on an H1 segment. Trend data (process variables recorded over time and displayed in time-domain graph form) may also be communicated using this type of VCR, with a “burst” of collected samples communicated per server message¹³. Client/server communications are checked for data corruption by their receivers, to ensure reliable data flow.

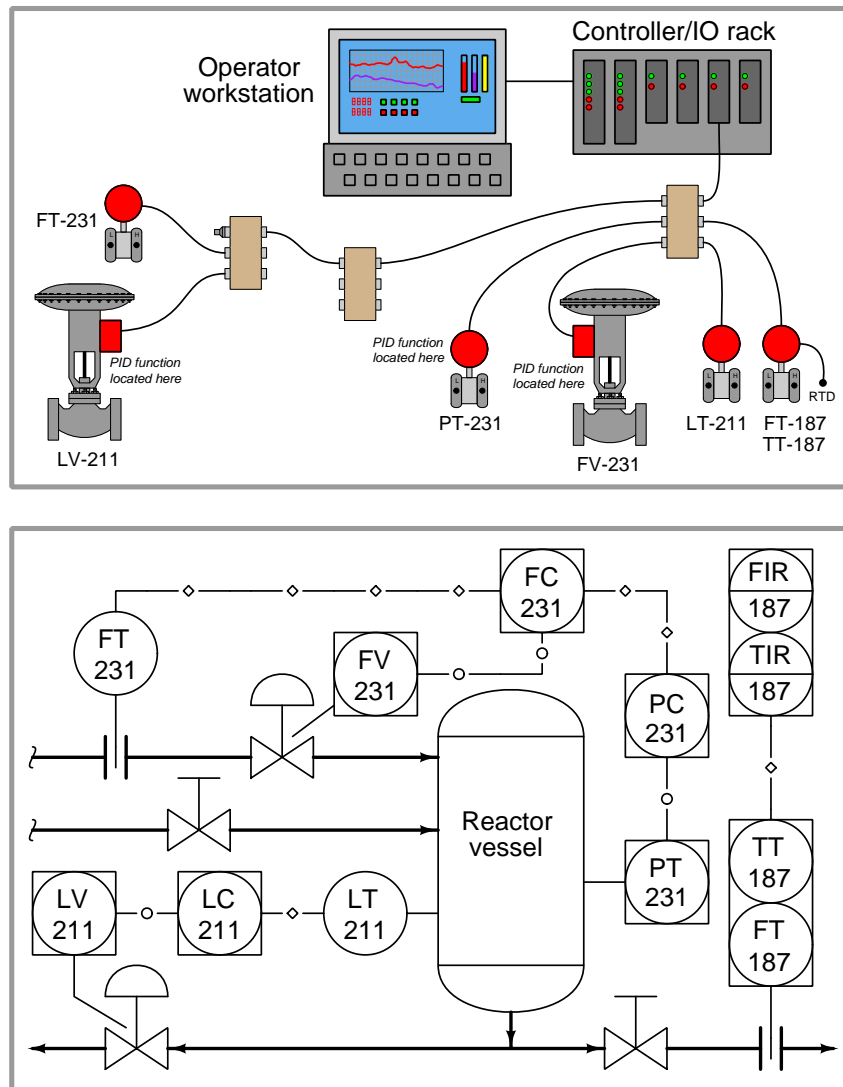
Source/Sink (also called Report Distribution): this VCR describes another class of unscheduled communications, permitted when a device receives a Pass Token (PT) message from the LAS. This is where a device broadcasts data out to a “group address” representing many devices. Source/sink communications are not checked for data corruption, as are client/server communications. Examples of messages communicated in a FF segment using the source/sink VCR include trend reports and alarms.

An analogy for making sense of VCRs is to imagine lines drawn between FF devices on a segment to connect their various messages to other devices. Each line represents an individual transmission

¹³The Source/Sink VCR is the preferred method for communicating trend data, but trends may be communicated via any of the three VCR types. All other factors being equal, acyclic communication (either Source/Sink or Client/Server) of trend data occupies less network bandwidth than cyclic communication (Publisher/Subscriber).

which must take place some time during the macrocycle. Each line is a VCR, some managed differently than others, some more critical than others, but all of them are simply communication events in time. A specific example of this is in the function block diagrams for a FF control system, where connecting lines between function blocks residing in different devices represent messages sent by the Publisher/Subscriber VCR method. Each line connecting function blocks between different devices is a message in response to a CD (Compel Data) token issued by the LAS, ensuring the deterministic transfer of critical control data between function blocks necessary for the control system to reliably function.

For example, consider this H1 segment connected to an interface card on a DCS rack, followed by a P&ID showing the relationships between the instruments:

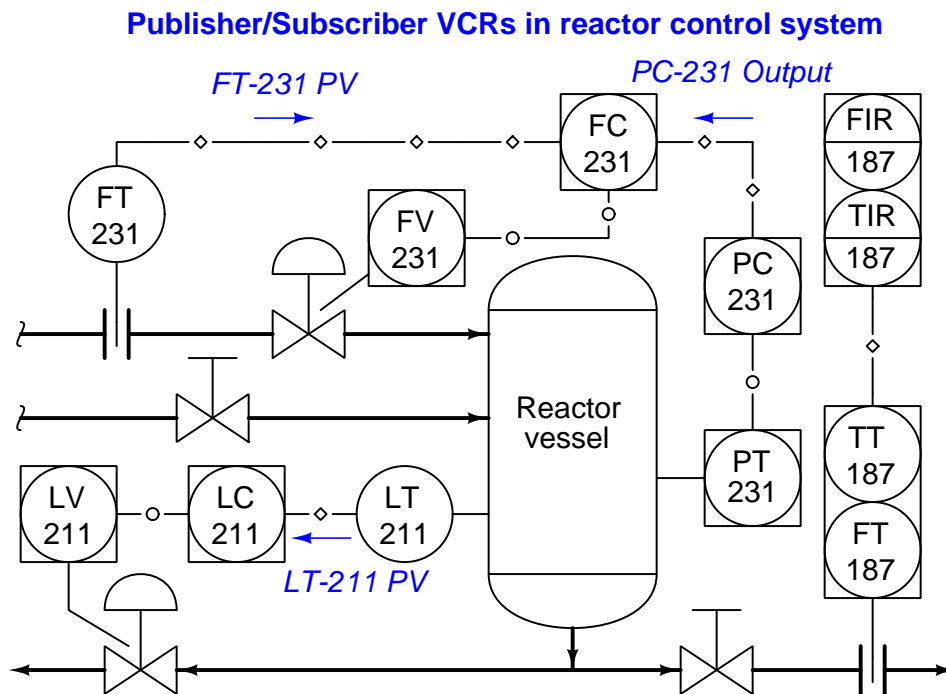


Loop 211 is a simple PID level control, regulating liquid level in the reactor vessel by releasing liquid from the bottom. Loop 187 is a simple indicating/recording system for temperature and flow, the signals coming from a multivariable transmitter. Loop 231 is a cascaded pressure/flow control system, with reactor pressure as the master variable and feed flow as the slave variable: the pressure controller (residing inside pressure transmitter PT-231) provides remote setpoint values to the flow controller (residing in the flow control valve FV-231), which then adjusts the position of the valve to achieve the desired feed flow rate into the reactor until reactor pressure stabilizes at setpoint.

Note the different line types used to represent digital signals in the P&ID: lines with diamond symbols represent data sent over Fieldbus cable, while lines with hollow circles represent data sent between functions within the same physical device. These latter “internal” data links help the reader discern which functions reside in which physical instruments. Functions residing within the same FF device must also share the same loop number. These standards for P&ID notation come from the Fieldbus Foundation’s *System Engineering Guidelines* document (revision 2.0, page 72) and from the ANSI/ISA-5.1-2009 “Instrumentation Symbols and Identification” standard.

For example, the PID control function represented by FC-231 resides within the valve positioner (FV-231), because those two bubbles share the same loop number and are connected with lines having hollow circles (which means they are parts of one homogeneous system rather than independent instruments). Likewise, the same line symbology tells us that pressure control PID function PC-231 resides within the pressure transmitter PT-231.


Control-critical variables communicated over the segment between devices include¹⁴ the output value of PC-231 (FC-231’s remote setpoint value), flow transmitter FT-231’s process variable measurement, and the process variable from level transmitter LT-211. These are all Publisher/Subscriber VCRs, transmitted at the request of a Compel Data (CD) token issued by the LAS device on a tightly controlled schedule:




¹⁴One other variable communicated in Publisher/Subscriber mode is the “back calculation” signal sent from slave (flow) controller FC-231 to master (pressure) controller PC-231. This particular signal is beyond the immediate scope of the discussion, and therefore is mentioned here only as a footnote.

Messages such as operator-adjusted setpoint values and maintenance tasks occur as Client/Server VCRs, done during the “unscheduled” communication times in the LAS’s sequence. The LAS device issues Pass Token (PT) messages to each device, giving permission for each device (one at a time) to transmit such information as necessary. Examples of such non-critical messages in our reactor control system are shown here:

Client/Server VCRs in reactor control system



Finally, our third VCR (Source/Sink) finds application in the reactor control system for flow transmitter FT-187, broadcasting its flow trend data during “unscheduled” periods in the LAS’s cycle, as well as for instrument alarm messages. Like the Client/Server messages, this one is prompted when the device receives a special Pass Token (PT) signal from the LAS, giving temporary permission for that device to broadcast its data:



16.3.3 Device capability

Not all FF devices are equally capable in terms of Data Link (layer 2) functions. The FF standard divides data link device functionality into three distinct groups, shown here in order of increasing capability:

- Basic devices
- Link Master devices
- Bridge devices

A *Basic* device is one capable of receiving and responding to tokens issued by the Link Active Scheduler (LAS) device. As discussed previously, these tokens may take the form of Compel Data (CD) messages which command immediate response from the Basic device, or Pass Token (PT) messages which grant the Basic device time-limited access to the segment for use in broadcasting data of lesser importance.

A *Link Master* device is one with the ability to be configured as the LAS for a segment. Not all FF devices have this ability, due to limited processing capability, memory, or both¹⁵.

A *Bridge* device links multiple H1 segments together to form a larger network. Field instruments are never Bridge devices – a Bridge is a special-purpose device built for the express purpose of joining two or more H1 network segments.

16.4 FF function blocks


Data-processing modules within FF systems are known as *function blocks*. Sometimes these blocks serve merely to catalogue data, while in other instances the blocks execute specific algorithms useful for process measurement and control. These “blocks” are not physical entities, but rather abstract software objects – they exist only as bits of data and instructions in computer memory. However, the blocks are represented on FF computer configuration displays as rectangular objects with input ports on the left-hand side and output ports on the right-hand side. The construction of a working control system comprised of FF devices consists of linking the outputs of certain function blocks with the inputs of other function blocks via configuration software and computer-based tools. This usually takes the form of using a computer to draw connecting lines between the output and input ports of different function blocks.

¹⁵Some FF devices capable of performing advanced function block algorithms for certain process control schemes may have the raw computational power to be an LAS, but the manufacturer has decided not to make them Link Master capable simply to allow their computational power to be devoted to the function block processing rather than split between function block tasks and LAS tasks.

16.4.1 Analog function blocks versus digital function blocks

Function-block programming in general strongly resembles the design philosophy of legacy analog-based computer systems, where specific functions (addition, subtraction, multiplication, ratio, time-integration, limiting, and others) were encapsulated in discrete operational amplifier circuits, and whole systems were built by connecting function blocks together in whatever patterns were desired to achieve a design goal. Here with Fieldbus programming, the function blocks are virtual (bits and data structures in digital memory) rather than real analog circuits, and the connections between blocks are merely pointer assignments in digital memory rather than actual “patch cable” connections between circuit boards.


An example contrasting analog circuit design with Fieldbus function-block design appears here, both systems selecting the *greatest* temperature signal to be the output. The system on the left-hand side receives analog voltage signals from three temperature sensors, using a network of operational amplifiers, diodes, and resistors to select the greatest voltage signal to be the output. The system on the right-hand side uses three Fieldbus transmitters to sense temperature, the greatest temperature signal selected by an algorithm (the ISEL function block) running in a Fieldbus device. The device running the ISEL function could be one of the three FF temperature transmitters, or another device on the segment:




Instead of analog voltage signals sent by wire to special-function circuit modules, FOUNDATION Fieldbus uses digital messages sent over an H1 network segment to special-function software “blocks” running inside ordinary Fieldbus devices. The lines connecting different function blocks together in a FOUNDATION Fieldbus system show the sources and destinations of these digital messages. If two FF function blocks reside in different FF devices, the connecting lines represent publisher/subscriber communication assignments coordinated by the Link Active Scheduler (LAS) device.

16.4.2 Function block location

There is usually some freedom of choice in where various function blocks may be located in a FF segment. Take for example the following flow control loop, where a flow transmitter feeds measured flow data into a PID control function block, which then drives a control valve to whatever position necessary to regulate flow. The actual physical device layout might look something like this:



The function block connections necessary for this control scheme to work are shown in the next diagram, coupling the AI (analog input) block located in the transmitter to a PID control block to an AO (analog output) block located in the valve positioner:




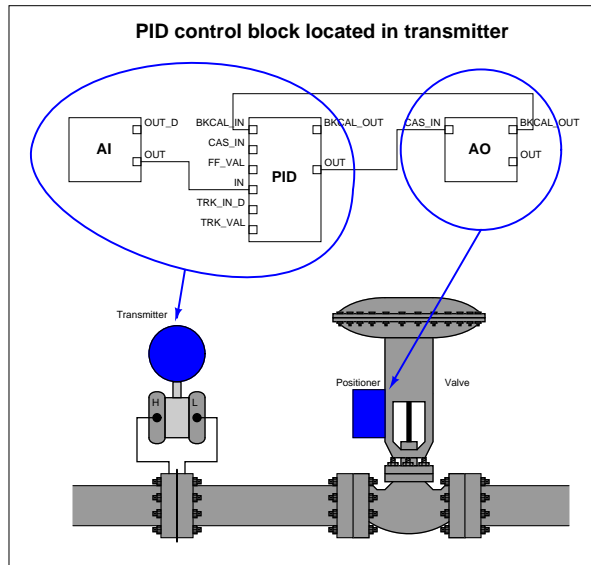
All function block inputs are on the left-hand sides of the blocks, and all outputs are on the right-hand sides. In this function block program, data from the analog input (AI) block flows into

the PID block. After calculating the proper output value, the PID block sends data to the analog output (AO) block where the final control element (e.g. valve, variable-speed motor) is adjusted. The AO block in turn sends a “back calculation” signal to the PID block to let it know the final control element has successfully reached the state commanded by the PID block’s output. This is important for the elimination of *reset windup*¹⁶ in the event the final control element fails to respond to the PID block’s output signal.

It should be obvious that the analog input (AI) block must reside in the transmitter, simply because only the transmitter is able to measure the process fluid flow rate. Likewise, it should be obvious that the analog output (AO) block must reside in the control valve positioner, simply because the valve is the only device capable of manipulating (exerting influence over) anything. However, given the lack of a separate controller device, the person configuring the Fieldbus loop may choose to locate the PID block in either the transmitter or the control valve positioner. So long as both the FF transmitter and the FF valve positioner possess PID function block capability, it is possible to locate the PID function block in either device.

¹⁶“Reset windup” which is also known as “integral windup” is what happens when any loop controller possessing reset (integral) action senses a difference between PV and SP that it cannot eliminate. The reset action over time will drive the controller’s output to saturation. If the source of the problem is a control valve that cannot attain the desired position, the controller will “wind up” or “wind down” in a futile attempt to drive the valve to a position it cannot go. In an FF system where the final control element provides “back calculation” feedback to the PID algorithm, the controller will not attempt to drive the valve farther than it is able to respond.


The following illustrations show the two possible locations of the PID function block in this system:



The only factor favoring one location over another for the PID function block is the number of communication broadcasts (“Compel Data” token distributions and replies) necessary per macrocycle. Note the lines connecting function blocks between the two instruments in the previous diagrams (lines crossing from one blue bubble to another). Each of these lines represents a VCR

(Virtual Communication Relationship) – an instance during each macrocycle where data must be transmitted over the network segment from one device to another. With the PID function block located in the flow transmitter, two lines connect blocks located in different physical devices. With the PID function block located in the valve positioner, only one line connects blocks in different physical devices. Thus, locating the PID function block in the valve positioner means only one CD message/reply is necessary per macrocycle, making the network communication more efficient.

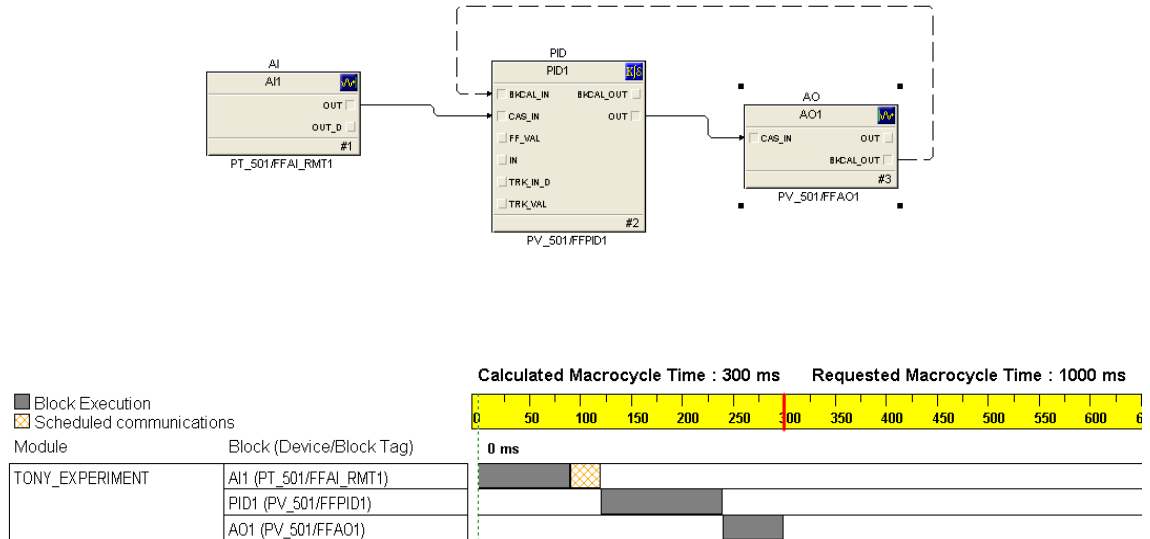
To illustrate the difference this re-location of the PID block makes, we will examine the function block diagram and macrocycle timing schedule on a simple pressure control FF loop, hosted on an Emerson DeltaV distributed control system. The first composite screenshot shows the function block diagram and schedule with the PID function block located in the transmitter (PT_501):



Note the two scheduled communication events (CD tokens and responses) necessary in the macrocycle schedule to enable communication between pressure transmitter PT_501's PID function block and valve positioner PV_501's analog output function block. The first CD token in this macrocycle schedule compels the PID block to publish its "output" signal (subscribed to by the analog output block), while the second token compels the analog output block to publish its "back calculation" signal (subscribed to by the PID block). The amount of time required for function block execution and their publisher/subscriber communications is 330 milliseconds, with a total macrocycle time of 1 second¹⁷.

¹⁷This is not an unreasonable loop execution time for a gas pressure control system. However, *liquid* pressure control is notoriously fast-acting, and will experience less than ideal response with a controller dead time of one second.

Now let's examine the same PID pressure control system with the PID function block moved to the valve. Here you see the function block diagram followed immediately by the updated macrocycle schedule:



In this macrocycle timing schedule, there is only one CD token needed: compelling the analog input block to publish its measurement signal (subscribed to by the PID block). This makes the block execution plus scheduled communication time 30 milliseconds shorter than before (300 milliseconds total as opposed to 330 milliseconds), since there is one less scheduled communications event happening. The total macrocycle time of 1 second remains unchanged, but now we have 30 milliseconds more unscheduled time during which other communication events may take place.

16.4.3 Standard function blocks

The FF standard specifies many different function blocks for the construction of control algorithms. Ten of them are considered “basic” FF function blocks:

- AI – Analog Input
- AO – Analog Output
- B – Bias
- CS – Control Selector
- DI – Discrete Input
- DO – Discrete Output
- ML – Manual Loader
- PD – Proportional/Derivative control
- PID – Proportional/Integral/Derivative control
- RA – Ratio

Nineteen more “Advanced” function blocks are incorporated in the FF standard:

- Pulse Input
- Complex Analog Output
- Complex Discrete Output
- Step Output PID
- Device Control
- Setpoint Ramp
- Splitter
- Input Selector
- Signal Characterizer
- Dead Time
- Calculate
- Lead/Lag
- Arithmetic
- Integrator

- Timer
- Analog Alarm
- Discrete Alarm
- Analog Human Interface
- Discrete Human Interface


Five more function blocks are specified as well:

- Multiple Analog Input
- Multiple Analog Output
- Multiple Digital Input
- Multiple Digital Output
- Flexible Function Block

The primary benefit of standardization is that the end-user may choose FF instruments manufactured by any standard-compliant vendor, and those function blocks should behave the same as the equivalent function blocks within any other manufacturer's model of FF device. There are, of course, examples where manufacturers have equipped their FF devices with "extended" capability function blocks going beyond the Fieldbus Foundation standard, and the user must beware of this.

16.4.4 Device-specific function blocks

In addition to the function blocks necessary to construct control schemes, all FF instruments contain one *Resource* block and usually one or more *Transducer* blocks describing details specific to that instrument. The following computer screenshot shows all function blocks within a Rosemount model 3095MV Fieldbus transmitter:



The Resource block appears first in this list, followed by three transducer blocks, then followed by the palette of general function blocks for use in constructing control algorithms. Information contained in the Resource block of an FF instrument includes the following:

- Identifier (the 32-byte code unique to every FF device)
- Type of device
- Device revision level
- Memory total and available (free) capacity
- Computation time
- Available features listing
- Current device state (Initializing, Standby, On-line, Failed, etc.)

Transducer blocks provide a means of organizing data relevant to the actual sensing inputs, outputs, calculated variables, and graphic displays of a FF device. There need not be a one-to-one correspondence between the number of transducer blocks in an FF device and the number of physical I/O channels it has. For example, in the Rosemount 3095MV multivariable transmitter, transducer block 1100 manages all physical measurement inputs (pressure and temperature sensors) while transducer block 1200 is reserved for inferred mass flow (based on calculations performed on the raw sensor measurements) and transducer block 1300 manages data for the liquid crystal display (LCD).

16.4.5 FF signal status

As mentioned earlier, function block programming bears a strong resemblance to analog function-block circuit design, where specific tasks are divided up into discrete elements, those elements connected together to form a larger system with more complex functionality. One of the important distinctions between legacy analog function block circuit design and FF function block programming is the data content of the lines connecting blocks together. In the analog world, each connecting line (wire) carries exactly one piece of information: a single variable represented in analog form by a voltage signal. In the world of Fieldbus, each connecting line carries not only the variable's numerical value, but also a *status* and in some cases an *engineering unit* (a unit of measurement). For example, a Fieldbus transmitter sensing temperature might output a digital process variable (PV) signal of “342 degrees Celsius, Good”, whereas a temperature transmitter with an analog (e.g. 4-20 mA) output is merely able to send a signal representing the temperature (no measurement unit or status information).

The inclusion of status along with data is a powerful concept, with roots in scientific practice. Scientists, as a rule, do their best to report the degree of *confidence* associated with the data they publish from experiments. Data is important, of course, but so is the degree of certainty with which that data was obtained. Obviously, data gathered with instruments of low quality (high uncertainty) will have different significance than data gathered with instruments of high precision and impeccable accuracy (low uncertainty). Any scientist basing research on a set of scientific data published by another scientist will have access to the data's certainty in addition to the data itself – a very valuable detail.

By the same token, data “published” by a FF device is only as good as the health of that device. A FF transmitter exhibiting noisy or wildly fluctuating measurements might very well be nearing complete failure, and therefore its published data should be treated with skepticism. Since FF devices are “smart” (meaning, among other things, they have self-diagnostic capability), they have the ability to flag their own data as “Bad” if an internal fault is detected. The data still gets published and sent to other FF function blocks, but the status sent along with that data warns all downstream blocks of its uncertainty.

The three major status conditions associated with every FF signal passed between function blocks are **Good**, **Bad**, and **Uncertain**. Sub-status states also exist¹⁸ to further delineate the nature of the uncertainty. “Sensor Failure” is an example of a sub-status value, describing the reason for a “Bad” status value from a process transmitter.

¹⁸For example, sub-statuses for a “Bad” status include *out of service*, *device failure*, *sensor failure*, and *non-specific*. Sub-statuses for an “Uncertain” status include *last usable value (LUV)*, *sensor conversion not accurate*, *engineering unit range violation*, *sub-normal*, and *non-specific*.

In computer science, there is a truism that “Garbage In equals Garbage Out,” sometimes abbreviated as *GIGO*. No algorithm, no matter how advanced, can guarantee an output of good data from an input of bad data¹⁹. This principle finds intelligent application in FF function block programming, as the blocks are programmed to switch mode when “Bad” or “Uncertain” input statuses are detected. For example, here are some of the possible actions a function block may be configured to take upon detection of a “Bad” input signal status:

- Set output signal to last “Good” value
- Fail high (set output signal to top-of-range value)
- Fail low (set output signal to bottom-of-range value)

Furthermore, status values are *propagated* in a FF system from the input to the output of every function block connected in series, reflecting the effect of an input signal’s uncertainty throughout the entire control loop. For example, an analog input (AI) block sending a “Bad” status signal to the process variable input of a PID control block will have its “Bad” status propagated to the output of the PID block as well. When that “Bad” PID output signal reaches the analog output (AO) function block, that final block knows the signal is not to be trusted, because its origin (the AI block) is untrustworthy. Any function blocks receiving the PID block’s output signal will likewise sense the “Bad” status and further propagate that status to their output signal(s). This “status propagation” ensures all function blocks in a Fieldbus control system are “aware” of the input data status, so that a “Bad” measurement does not result in “bad” control decisions made on that data.

¹⁹The great pioneer of mechanical computing technology, Charles Babbage, commented in his book *Passages from the Life of a Philosopher* in 1864 that not one but two members of the British parliament asked him whether his computer (which he called the Difference Engine) could output correct answers given incorrect data. His reaction was both frank and hilarious: “I am not able rightly to apprehend the kind of confusion of ideas that could provoke such a question.”

16.4.6 Function block modes

All FF function blocks must support multiple *modes* of operation, describing how the block should execute its intended function. Several different function block modes are commonly found for FF function blocks, though not all FF function blocks support all of these modes:

- **OOS** (Out Of Service) – *All function blocks are required to support this mode, where the block freezes its output at the last calculated value and attaches a “Bad” status value*
- **Man** (Manual) – *the output of the block is fixed at a value determined by the technician, with a “Good” status value attached*
- **Auto** (Automatic) – *the function block processes information normally*
- **Cas** (Cascade) – *the function block processes information normally*
- **Iman** (Initialization Manual) – *the output of the block is fixed at its last calculated value, due to the output signal path being incomplete*
- **LO** (Local Override) – *the output of the block is fixed at its last calculated value, due to a detected fault condition within the device*
- **RCas** (Remote Cascade) – *the function block processes information normally based on a setpoint sent from a remote source to the block’s RCas_In input*
- **ROut** (Remote Output) – *the function block passes data to its output sent from a remote source to the block’s ROut_In input*

Instrumentation technicians and professionals are already familiar with the concept of a controller having “Automatic,” “Manual,” and even “Cascade” operating modes, but Fieldbus function block programming extends this general concept to each and every function block. With FF, *each block* may be independently set into “Automatic” or “Manual” mode, which is a useful tool for testing FF algorithms and troubleshooting complex FF control schemes. The “Out of Service” mode, for instance, is commonly set when performing routine maintenance on an FF device (e.g. checking the calibration of an FF transmitter).

It is worth noting an important distinction here between Manual mode and OOS (Out Of Service) mode. In both cases, the function block’s output becomes fixed at some value, but a major difference between these two modes is their associated statuses. In Manual mode, the output value is fixed and the status is “Good,” allowing all function blocks downstream to remain operational. In OOS mode, the output value is fixed and the status is “Bad,” causing all downstream function blocks to react as they would when receiving any “Bad” signal status (usually by shedding to Manual mode themselves). Placing a function block in Manual mode is useful when performing tests on the control strategy because it allows the technician or engineer to simulate values that might come from transmitters and other “upstream” devices in the loop. All function blocks receiving a signal from a block in Manual mode will continue to operate as they are designed. However, placing a function block in OOS mode is quite different in that all function blocks receiving that signal will act as though there is a serious problem rather than acting normally.

In addition to these operating modes for FF function blocks (not all of which are supported by all FF blocks), FF function blocks also have four mode categories describing valid modes for the block to be in under various conditions:

- Target
- Actual
- Permitted
- Normal

A block’s “Target” mode is the mode it strives to be in if possible. The “Actual” mode is the mode the block is in at the present time. “Permitted” modes list all the different modes which may be used as “target” modes. “Normal” is a category describing to an operator interface what a block’s normal operation mode should be, but the block itself does not heed this setting.

16.5 H1 FF device configuration and commissioning

Fieldbus devices require far more attention in their initial setup and commissioning than their analog counterparts. Unlike an analog transmitter, for example, where the only “configuration” settings are its zero and span calibration adjustments, a FF transmitter has a substantial number of parameters describing its behavior. Some of these parameters must be set by the end-user, while others are configured automatically by the host system during the start-up process, which we generally refer to as *commissioning*.

16.5.1 Configuration files


In order for a FF device to work together with a host system (which may be manufactured by a different company), the device must have its capabilities explicitly described so the host system “knows what to do with it.” This is analogous to the need for *driver* files when interfacing a personal computer with a new peripheral device such as a printer, scanner, or modem.

A standardized language exists for digital instrumentation called the *Device Description Language*, or *DDL*. All FF instrument manufacturers are required to document their devices’ capabilities in this standard-format language, which is then compiled by a computer into a set of files known as the *Device Description* (DD) files for that instrument. DDL itself is a text-based language, much like C or Java, written by a human programmer. The DD files are generated from the DDL source file by a computer, output in a form intended for another computer’s read-only access. For FF instruments, the DD files end in the filename extensions `.sym` and `.ffo`, and may be obtained freely from the manufacturer or from the Fieldbus Foundation²⁰ website (<http://www.fieldbus.org>). The `.ffo` DD file is in a binary format readable only by a computer with the appropriate “DD services” software active. The `.sym` DD file is ASCII-encoded, making it viewable by a human by using a text editor program (although you should not attempt to edit the contents of a `.sym` file).

Other device-specific files maintained by the host system of a FF segment are the *Capability* and *Value* files, both referred to as *Common Format Files*, or `.cff` files. These are also text-readable (ASCII encoded) digital files describing device capability and specific configuration values for the device, respectively. The Capability file for a FF device is typically downloaded from either the manufacturer’s or the Fieldbus Foundation website along with the two DD files, as a three-file set (filename extensions being `.cff`, `.sym`, and `.ffo`, respectively). The Value file is generated by the host system during the device’s configuration, storing the specific configuration values for that specific device and system tag number. The data stored in a Value file may be used to duplicate the exact configuration of a failed FF device, ensuring the new device replacing it will contain all the same parameters.

²⁰One of the tasks of the Fieldbus Foundation is to maintain approved listings of FF devices in current manufacture. The concept is that whenever a manufacturer introduces a new FF device, it must be approved by the Fieldbus Foundation in order to receive the Fieldbus “badge” (a logo with a stylized letter “F”). Approved devices are cataloged by the Fieldbus Foundation, complete with their DD file sets. This process of approval is necessary for operational compatibility (called *interoperability*) between FF devices of different manufacture. Without some form of centralized standardization and approval, different manufacturers would invariably produce devices mutually incompatible with each other.

A screenshot of a .cff Capability file opened in a text editor program appears here, showing the first few lines of code describing the capabilities of a Yokogawa model DYF vortex flowmeter:



```

File Edit Tools Syntax Buffers Window Help
//*****
//  DYF-FF Phasel Capability File
//
// $Author:  Yoshida.Shinnosuke  $
// $Revision:  1.1  $
// $Date:  15 Nov 2006 22:03:14  $
//
//*****

//=====
//  Header
//=====
[File Header]
Description="This is a Capabilities File for DYF Vortex Flowmeter"
FileType=CapabilitiesFile
FileDate=2006,10,05
CffVersion=1,7

//=====
//  Device Header
//=====
[Device Header]
DeviceName="DYF"
DeviceClass=LINKMASTER
CommGroup=3
CommClass=Class31+Class32
CommSubClass=Class3Publisher+Class3Subscriber+Class3LinkMaster+Class3TimePublisher
Value

[Device VFD 1] //Management VFD
VendorName="Yokogawa Electric"
ModelName="DYF"


```

As with “driver” files needed to make a personal computer peripheral device function, it is important to have the correct versions of the Capability and DD files installed on the host system computer before attempting to commission the device. It is permissible to have Capability and DD files installed that are newer than the physical device, but not vice-versa (a newer physical device than the Capability and DD files). This requirement of proper configuration file management is a new task for the instrument technician and engineer to manage in their jobs. With every new FF device installed in a control system, the proper configuration files must be obtained, installed, and archived for safe keeping in the event of data loss (a “crash”) in the host system.

16.5.2 Device commissioning

This section illustrates the commissioning of a Fieldbus device on a real segment, showing screenshots of a host system's configuration menus. The particular device happens to be a Fisher DVC5000f valve positioner, and the host system is a *DeltaV* distributed control system manufactured by Emerson. All configuration files were updated in this system prior to the commissioning exercise. Keep in mind that the particular steps taken to commission any FF device will vary from one host system to another, and may not follow the sequence of steps shown here.

If an unconfigured FF device is connected to an H1 network, it appears as a “decommissioned” device. On the Emerson DeltaV host system, all decommissioned FF devices appear within a designated folder on the “container” hierarchy. Here, my Fisher DVC5000 device is shown highlighted in blue. A commissioned FF device appears just below it (PT_501), showing all available function blocks within that instrument:




Before any FF device may be recognized by the DeltaV host system, a “placeholder” and tag name must be created for it within the segment hierarchy. To do this, a “New Fieldbus Device” must be added to the H1 port. Once this option is selected²¹, a window opens up to allow naming of this new device:

The image shows a screenshot of the 'Fieldbus Device Properties' dialog box. The 'General' tab is selected. The 'Object type' is 'Fieldbus Device'. The 'Modified' and 'Modified by' fields are empty. The 'Device tag' is 'PV_501'. The 'Description' is 'Pressure control valve (positioner)'. The 'Device ID' field is empty. The 'Address' is '35'. The 'Use as backup link master' checkbox is unchecked. The 'Manufacturer' is 'Fisher Controls'. The 'Device type' is 'DVC5000f AO/PID/IS Digital Valve'. The 'Device revision' is '9'. The 'OK', 'Cancel', and 'Help' buttons are at the bottom.


Here, the tag name “PV_501” has been chosen for the Fisher valve positioner, since it will work in conjunction with the pressure transmitter PT_501 to form a complete pressure control loop. In addition to a tag name (PV_501), I have also added a text description (“Pressure control valve (positioner)”), and specified the device type (Fisher DVC5000f with AO, PID, and IS function block capability). The DeltaV host system chose a free address for this device (35), although it is possible to manually select the desired device address at this point. Note the “Backup Link Master” check box in this configuration window, which is grey in color (indicating the option is not available with this device).

²¹On the Emerson DeltaV system, most options are available as drop-down menu selections following a right-mouse-button click on the appropriate icon.

After the device information has been entered for the new tag name, a “placeholder” icon appears within the hierarchy for the H1 segment (connected to Port 1). You can see the new tag name (PV_501) below the last function block for the commissioned FF instrument (PT_501). The actual device is still decommissioned, and appears as such:



By right-clicking on the new tag name and selecting the “Commission” option, a new window opens to allow you to select which decommissioned device should be given the new tag name. Since there is only one decommissioned device on this particular H1 segment, only one option appears within the window:



After selecting the decommissioned device you wish to commission, the DeltaV host system prompts you to reconcile any differences between the newly created tag name placeholder and the decommissioned device. If you want to use the existing values stored within the physical (decommissioned) device, you skip the “reconcile” step. If you want to alter the values in the device from what they presently are, you choose the “reconcile” option which then opens up an editing window where you can set the device values however you wish.




After selecting (or not selecting) the “reconcile” option, the DeltaV system prompts you to confirm commissioning of the device, after which it goes through a series of animated²² display sequences as the device transitions from the “Standby” state to the “Commissioned” state:




As you can see, the commissioning process is not very fast. After nearly one full minute of waiting, the device is still “Initializing” and not yet “Commissioned.” The network speed of 31.25 kbps and the priority of scheduled communications are limiting factors when exchanging large quantities of configuration data over a FF H1 network segment. In order for device configuration to not interrupt or slow down process-critical data transfers, all configuration data exchanges must wait for unscheduled time periods, and then transmit at the relatively slow rate of 31.25 kbps when the allotted times arrive. Any technician accustomed to the fast data transfer rates of modern Ethernet devices will feel as though he or she has taken a step back in time when computers were *much* slower.

²²Animated graphics on the Emerson DeltaV control system prominently feature an anthropomorphized globe valve named Duncan. There’s nothing like a computer programmer with a sense of humor . . .

After commissioning this device on the DeltaV host system, several placeholders in the hierarchy appear with blue triangles next to them. In the DeltaV system, these blue triangle icons represent the need to download database changes to the distributed nodes of the system:



After “downloading” the data, the new FF valve positioner shows up directly below the existing pressure transmitter as a commissioned instrument, and is ready for service. The function blocks for pressure transmitter PT_501 have been “collapsed” back into the transmitter’s icon, and the function blocks for the new valve positioner (PV_501) have been “expanded” for view:




As you can see, the new instrument (PV_501) does not offer nearly as many function blocks as the original FF instrument (PT_501). The number of Fieldbus function blocks offered by any FF instrument is a function of that instrument’s computational ability, internal task loading, and the discretion of its designers. Obviously, this is an important factor to consider when designing a FF segment: being sure to include instruments that contain all the necessary function blocks to execute the desired control scheme. This may also become an issue if one of the FF instruments in a control scheme is replaced with one of a different manufacturer or model, having fewer available function blocks. If one or more mission-critical function blocks is not available in the replacement instrument, a different replacement must be sought.

16.5.3 Calibration and ranging

Calibration and ranging for a FF device is similar in principle to any other “smart” measurement instrument. Unlike analog instruments, where the “zero” and “span” adjustments completely define the instrument’s calibration *and* range, calibration and ranging are two completely different functions in a digital instrument.


To begin, we will examine a block diagram of an analog pressure transmitter showing the zero and span adjustments, with analog signaling between all functions inside the transmitter:



The “zero” and “span” adjustments together define the mathematical relationship between sensed pressure and current output. Calibration of an analog transmitter consists of applying known (reference standard) input stimuli to the instrument, and adjusting the “zero” and “span” settings until the desired current output values are achieved. The goal in doing this is to ensure accuracy of measurement.

The “range” of a transmitter is simply the input values associated with 0% and 100% output signals (e.g. 4 mA and 20 mA). Ranging an analog transmitter consists (also) of adjusting the “zero” and “span” settings until the output signal corresponds to the desired LRV and URV points of the measured variable. For an analog transmitter, the functions of ranging and calibration are always performed by the technician at the same time: to calibrate an analog transmitter is to range it, and vice-versa.

By contrast, a “smart” (digital) transmitter equipped with an analog 4-20 mA current output distinctly separates the calibration and range functions, each function determined by a different set of adjustments:



Calibration of a “smart” transmitter consists of applying known (reference standard) input stimuli to the instrument and engaging the “trim” functions until the instrument accurately registers the input stimuli. For a “smart” transmitter equipped with analog electronic (4-20 mA) output, there are *two* sets of calibration trim adjustments: one for the analog-to-digital converter and another for the digital-to-analog converter.

Ranging, by contrast, establishes the mathematical relationship between the measured input value and the output current value. To illustrate the difference between calibration and ranging, consider a case where a pressure transmitter is used to measure water pressure in a pipe. Suppose the transmitter’s pressure range of 0 to 100 PSI translates to a 4-20 mA output current. If we desired to re-range an analog transmitter to measure a greater span of pressures (say, 0 to 150 PSI), we would have to re-apply known pressures of 0 PSI and 150 PSI while adjusting the zero and span potentiometers so 0 PSI input gave a 4 mA output value and 150 PSI input gave a 20 mA output value. The only way to re-range an analog transmitter is to completely re-calibrate it.


In a “smart” (digital) measuring instrument, however, calibration against a known (standard) source need only be done at the specified intervals to ensure accuracy over long periods of time

given the instrument's inevitable drift. If our hypothetical transmitter were recently calibrated against a known pressure standard and trusted not to have drifted since the last calibration cycle, we could re-range it by simply changing the URV (upper range value) so that an applied pressure of 150 PSI now commands it to output 20 mA instead of an applied pressure of 100 PSI as was required before. Digital instrumentation allows us to re-range without re-calibrating, representing a tremendous savings in technician time and effort.

The distinction between *calibration* and *ranging* tends to confuse people, even some experienced technicians. When working with an analog transmitter, you cannot calibrate without setting the instrument's range as well: the two functions are merged in the same procedures of adjusting zero and span. When working with a digital transmitter, however, the function of calibration and the function of ranging are entirely separate.

For a detailed analogy explaining the distinction between calibration and ranging, refer to section 18.6 beginning on page 1275.

Fieldbus instruments, of course, are “smart” in the same way, and their internal block diagrams look much the same as the “smart” transmitters with analog current output, albeit with a far greater number of parameters within each block. The rectangle labeled “XD” in the following diagram is the Transducer block, while the rectangle labeled “AI” is the Analog Input block:



Calibration (trim) values are set in the transducer block along with the engineering unit, making the output of the transducer block a digital value scaled in real units of measurement (e.g. PSI, kPa, bar, mm Hg, etc.) rather than an abstract ADC “count” value. The analog input function block’s **Channel** parameter tells it which transducer output to receive²³ as the pre-scaled “Primary Value”, which it may then translate to another scaled value based on a proportionality between transducer scale values (**XD_Scale** high and low) and output scale values (**OUT_Scale** high and low).

To calibrate such a transmitter, the transducer block should first be placed in *Out Of Service* (OOS) mode using a handheld FF communicator or the Fieldbus host system. Next, a standard (calibration-grade) fluid pressure is applied to the transmitter’s sensor and the **Cal_Point_Lo**

²³Fieldbus transmitters often have multiple channels of measurement data to select from. For example, the multi-variable Rosemount 3095MV transmitter assigns channel 1 as differential pressure, channel 2 as static pressure, channel 3 as process temperature, channel 4 as sensor temperature, and channel 5 as calculated mass flow. Setting the **Channel** parameter properly in the AI block is therefore critical for linking it to the proper measurement variable.

parameter is set to equal this applied pressure. After that, a greater pressure is applied to the sensor and the `Cal.Point.Hi` parameter is set to equal this applied pressure. After setting the various calibration record-keeping parameters (e.g. `Sensor.Cal.Date`, `Sensor.Cal.Who`), the transducer block's mode may be returned to *Auto* and the transmitter used once again.

To range such a transmitter, a correspondence between sensed pressure and the process variable must be determined and entered into the analog input function block's `XD.Scale` and `OUT.Scale` parameters. If the pressure transmitter is being used to indirectly measure something other than pressure, these range parameters will become very useful, not only proportioning the numerical values of the measurement, but also casting the final digital output value into the desired “engineering units” (units of measurement).


Ranging in Fieldbus transmitters is a somewhat confusing topic due to the unfortunate names given to the different `L.Type` parameter options. Here is a list of the `L.Type` parameter options along with their meanings:

- Direct = the AI block will publish the signal output by the XD block, regardless of the specified `OUT.Scale` range
- Indirect = the AI block will mathematically scale the signal from the XD block into a range specified by `OUT.Scale` parameters using a linear equation (e.g. $y = mx + b$)
- Indirect square root = same as above, except that a square-root function is applied to the percentage of range (useful when characterizing flow transmitters based on differential pressure measurement)

The terms “direct” and “indirect” are unfortunate²⁴, because they often cause people to interpret them as “direct” and “reverse” (as though `L.Type` described the *direction of action* for the function block). This is *not* what these terms mean for the AI block! What a “direct” value for `L.Type` means is that the *raw* value of the XD block is what will be published onto the Fieldbus network by the AI block. What an “indirect” value for `L.Type` means is that the XD block's signal will be *scaled* to a different range (specified by the `OUT.Scale` parameter). In summary, the technician must set the `XD.Scale` range according to the primary signal sensed by the transmitter's sensing element, and set the `OUT.Scale` range according to what the rest of the control system needs to see proportional to that primary signal.

²⁴If I were king for a day, I would change the labels “direct” and “indirect” to “raw” and “scaled”, respectively. Alternatively, I would abandon the “direct” option altogether, because even when this option is chosen the `OUT.Scale` range still exists and may contain “scaled” values even though these are ignored in “direct” mode!

The concept of ranging a FF transmitter makes more sense when viewed in the context of a real application. Consider this example, where a pressure transmitter is being used to measure the level of ethanol (ethyl alcohol) stored in a 40 foot high tank. The transmitter connects to the bottom of the tank by a tube, and is situated 10 feet below the tank bottom:



Hydrostatic pressure exerted on the transmitter’s sensing element is the product of liquid density (γ) and vertical liquid column height (h). When the tank is empty, there will still be a vertical column of ethanol 10 feet high applying pressure to the transmitter’s “high” pressure port. Therefore, the pressure seen by the transmitter in an “empty” condition is equal to:

$$P_{empty} = \gamma h_{empty} = (49.3 \text{ lb/ft}^3)(10 \text{ ft})$$

$$P_{empty} = 493 \text{ lb/ft}^2 = 3.424 \text{ PSI}$$

When the tank is completely full (40 feet), the transmitter sees a vertical column of ethanol 50 feet high (the tank’s 40 foot height plus the suppression height of 10 feet created by the transmitter’s location below the tank bottom). Therefore, the pressure seen by the transmitter in a “full” condition is equal to:

$$P_{full} = \gamma h_{full} = (49.3 \text{ lb/ft}^3)(40 \text{ ft} + 10 \text{ ft})$$

$$P_{full} = 2465 \text{ lb/ft}^2 = 17.12 \text{ PSI}$$

Thus, the transducer (XD) block in this Fieldbus transmitter will sense a liquid pressure ranging from 3.424 PSI to 17.12 PSI over the full range of the tank’s storage capacity.

However, we do not want this transmitter to publish a signal to the Fieldbus network in units of PSI, because the operations personnel monitoring this control system want to see a measurement of ethanol *level* inside the tank, not hydrostatic *pressure* at the bottom of the tank. We may be exploiting the principle of hydrostatic pressure to sense ethanol level, but we do not wish to report this measurement as a pressure.

The proper solution for this application is to set the `L_Type` parameter to “indirect” which will instruct the AI function block to mathematically scale the XD block’s pressure signal into a different range. Then, we must specify²⁵ the expected pressure range and its corresponding level range as `XD_Scale` and `OUT_Scale`, respectively²⁶:

AI block parameter	Range values
<code>L_Type</code>	Indirect
<code>XD_Scale</code>	3.424 PSI to 17.12 PSI
<code>OUT_Scale</code>	0 feet to 40 feet

Now, the ethanol tank’s level will be accurately represented by the FF transmitter’s output, both in numeric value and measurement unit. An empty tank generating a pressure of 3.424 PSI causes the transmitter to output a “0 feet” digital signal value, while a full tank generating 17.12 PSI of pressure causes the transmitter to output a “40 feet” digital signal value. Any ethanol levels between 0 and 40 feet will likewise be represented proportionally by the transmitter.

If at some later time the decision is made to re-locate the transmitter so it no longer has a 10 foot “suppression” with regard to the tank bottom, the `XD_Scale` parameters may be adjusted to reflect the corresponding shift in pressure range, and the transmitter will still accurately represent ethanol level from 0 feet to 40 feet, without re-calibrating or re-configuring anything else in the transmitter.

If we wished, we could even mathematically determine the liquid *volume* stored inside this ethanol tank at different sensed pressures, and then scale the AI block’s `OUT_Scale` parameter to report a volume in units of gallons, liters, cubic feet, or any other appropriate volume unit. Using the “indirect” mode with appropriate `XD_Scale` and `OUT_Scale` parameter values gives us great flexibility in how the transmitter senses and represents process data.


In summary, we set the `XD_Scale` parameter to the physical range of measurement directly sensed by the transducer, we set the `OUT_Scale` parameter to the corresponding range of measurement we wish the transmitter to report to the rest of the control system, and we set `L_Type` to “indirect” to enable this translation from one range to another. We should only use the “direct” `L_Type` setting if the raw transducer range is appropriate to output to the rest of the control system (e.g. if the transmitter directly senses fluid pressure and we wish this very same pressure value to be published onto the Fieldbus network by the transmitter, with no scaling).

²⁵It is important to note that *you* must correctly calculate the corresponding `XD_Scale` and `OUT_Scale` parameter values in order for this to work. The Fieldbus instrument does not calculate the parameters for you, because it does not “know” how many PSI correspond to how many feet of liquid level in the tank. These values must be calculated by some knowledgeable human technician or engineer and then entered into the instrument’s AI block, after which the instrument will execute the specified scaling as a purely mathematical function.

²⁶When configuring the `XD_Scale` high and low range values, be sure to maintain consistency with the transducer block’s `Primary_Value_Range` parameter unit. Errors may result from mis-matched measurement units between the transducer block’s measurement channel and the analog input block’s `XD_Scale` parameter.

Here is another Fieldbus transmitter ranging application, this time a differential pressure transmitter sensing pressure dropped across an orifice plate in order to infer the rate of flow for fluid inside the pipe. The transmitter senses small amounts of pressure difference (expressed in a unit of pressure called *inches water column*), but what we want it to report to the Fieldbus network is an actual flow rate in gallons per minute.

If we happen to know that this orifice plate produces a pressure drop of 125 inches water column (125 "WC) at a flow rate of 350 gallons per minute (350 GPM), we could set up the scaling parameters as shown:



AI block parameter	Range values
L_Type	Indirect Square Root
XD_Scale	0 inches water to 125 inches water
OUT_Scale	0 GPM to 350 GPM

Note the use of the “indirect *square root*” L_Type parameter value instead of just “indirect” as we used in the ethanol tank example. The square root function is necessary in this application because the relationship between differential pressure (ΔP) and flow rate (Q) through an orifice is nonlinear, as described by the following formula:

$$Q = k\sqrt{\Delta P}$$

This particular nonlinearity is unique to pressure-based measurements of fluid flow, and does not find application in any other form of process measurement.

As before, though, we see a common theme with the XD_Scale and OUT_Scale parameter ranges: we set the XD_Scale parameter to the physical range of measurement directly sensed by the transducer, we set the OUT_Scale parameter to the corresponding range of measurement we wish the transmitter to report to the rest of the control system, and we set L_Type to “indirect” to enable this translation from one range to another.

16.6 H1 FF segment troubleshooting

Feedback obtained from industrial users of FF reveal a common pattern: Fieldbus is a powerful and reliable technology, but only if it is properly installed. Poor installations, usually driven by a desire to minimize capital expenses, will cause numerous problems during commissioning and operation.

One relatively easy way to avoid problems caused by short-circuits in FF wiring is to use coupling devices with built-in short-circuit protection. This feature does not add significant cost to the coupling device, and it will prevent the entire segment from failing due to a short-circuit on a single spur cable or within a device. Use coupling devices with indicator LEDs as well, since these give easy visual verification of network power which may greatly accelerate FF segment troubleshooting when the need arises.

16.6.1 Cable resistance

A simple check of an H1 segment's cabling consists of a series of resistance measurements performed with the segment unpowered (as is standard with any electrical resistance check), with all FF devices disconnected, and with the cable entirely disconnected (all three conductors) at the host end. The following table shows guidelines published by the Fieldbus Foundation for H1 segment cable resistance measurements:

Measurement points	Expected resistance
Between (+) and (−) conductors	> 50 k Ω , increasing over time
Between (+) conductor and shield (ground)	> 20 M Ω
Between (−) conductor and shield (ground)	> 20 M Ω
Between shield conductor and earth ground	> 20 M Ω

The last resistance check shown in the table checks for the presence of ground connections in the shield conductor *other than* the one ground connection at the host end (which has been disconnected for the purposes of the test). Since the shield should only be grounded at one point²⁷ (to avoid ground loops), and this one point has been disconnected, the shield conductor should register no continuity with earth ground during the test.

The necessity of disconnecting all FF devices and host system interfaces is essential so that the resistance measurements reflect the health of the cable and nothing else. The presence of any FF devices on the segment would substantially affect the resistance measurements, particularly resistance between the signal (+ and −) conductors.

²⁷An alternative method of shield grounding is to directly connect it to earth ground at one end, and then capacitively couple it to ground at other points along the segment length. The capacitor(s) provide an AC path to ground for “bleeding off” any induced AC noise without providing a DC path which would cause a ground loop.

16.6.2 Signal strength

The Fieldbus Foundation specifies a signal voltage (peak-to-peak) range of 350 mV to 700 mV for a healthy FF segment. Excessive signal voltage levels point to a lack of terminator resistor(s), while insufficient voltage levels point to an over-abundance of terminators (or perhaps even a device short):

Signal voltage (pk-pk)	Interpretation
800 mV or more	Possibly missing terminator resistor
350 mV to 700 mV	Good signal strength
150 mV to 350 mV	Marginally low signal – possible extra terminator resistor(s)
150 mV or less	Too little signal to function

16.6.3 Electrical noise

FF, like all digital networks, are unaffected by noise voltage below a certain threshold. If noise voltage is present in excessive quantity, though, it may cause bits to be misinterpreted, causing data errors. The Fieldbus Foundation gives the following recommendations²⁸ for noise voltage levels on a FF segment:

Noise voltage (pk-pk)	Interpretation
25 mV or less	Excellent
25 mV to 50 mV	Okay
50 mV to 100 mV	Marginal
100 mV or more	Poor

Fieldbus diagnostic tools measure noise on the network segment during times between message frames, when there should be purely DC voltage between the two conductors.

²⁸Bear in mind the tolerable level for noise will vary with signal voltage level as well. All other factors being equal, a strong signal is less affected by the presence of noise than a weak signal (i.e. the signal-to-noise ratio, or *SNR*, is crucial).

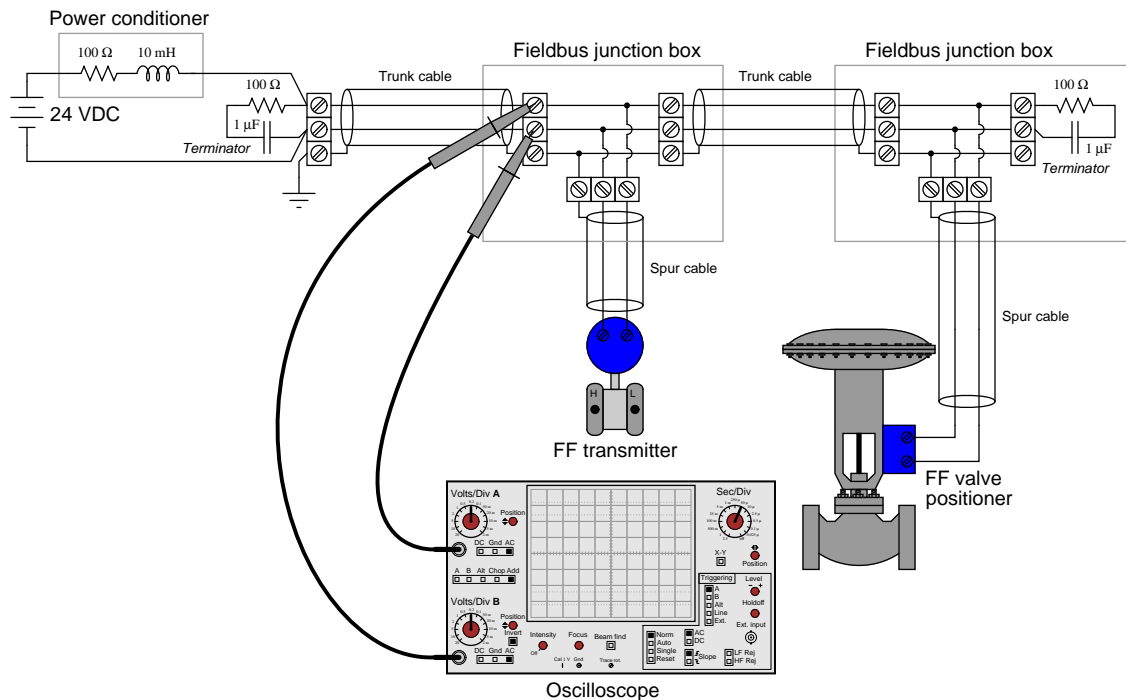
16.6.4 Using an oscilloscope on H1 segments

A tool available in most instrument shops is a digital-storage oscilloscope, which may be used to measure and display FF H1 signal waveforms for analysis of problems. Analog oscilloscopes are also useful for network troubleshooting, but to a lesser degree²⁹.

When using an oscilloscope to measure FF H1 signals, it is very important not to connect either of the FF segment conductors to earth ground through the oscilloscope. Introducing such a “ground fault” to the network segment will almost certainly *cause* communication problems, in addition to whatever problems already exist that compel you to diagnose with an oscilloscope. If a single channel of the oscilloscope is connected across the segment wires, the “ground” clip of the probe will force one of those conductors to earth ground potential via the metal chassis of the oscilloscope which is grounded through the third prong of the power plug for safety. An exception to this rule is if the oscilloscope itself is battery-powered and has an insulated case where no ground connection is made through the surface it sits on or the human hand that holds it. Otherwise, using a single channel on a line-powered oscilloscope to measure network signals is inviting trouble.

²⁹It is impossible to “lock in” (trigger) non-periodic waveforms on an analog oscilloscope, and so most network communications will appear as an incomprehensible blur when viewed on this kind of test instrument. Digital oscilloscopes have the ability to “capture” and display momentary pulse streams, making it possible to “freeze” any portion of a network signal for visual analysis.

If a line-powered oscilloscope must be used, the proper way to configure it is for *differential channel* measurement. In this mode, the oscilloscope will register the voltage *between* two probe tips, rather than register the voltage between a single probe tip and earth ground.



Configuring a dual-trace oscilloscope for differential mode is quite simple. On the front panel of the oscilloscope, you must set the multi-trace controls to the *Add* mode, where one trace on the screen represents the instantaneous sum of the two inputs (channels “A” and “B”). The volts per division “sensitivity” of both channels should be set to exactly the same value. Also, the *Invert* control must be engaged for the second input channel, forcing that channel’s signal to be inverted (register upside-down on the screen). The summation of channel “A” and an inverted channel “B” is equivalent to the mathematical difference (subtraction) between “A” and “B,” which means the single trace on the screen now represents the difference of potential between the two probe tips. The oscilloscope now behaves as an ungrounded voltmeter, where neither of the test leads is referenced to earth ground.

16.6.5 Message re-transmissions

Aside from voltage parameters (signal strength, noise amplitude), another good indicator of FF segment health is the number of message *re-transmissions* over time. Certain types of communication on an H1 segment require verification of a received signal (particularly client/server VCRs such as those used to communicate operator setpoint changes and diagnostic messages). If the signal received by the client FF device appears corrupted, the device will request a *re-transmission* of the message from the server device. Re-transmission events, therefore, are an indication of how often messages are getting corrupted, which is a direct function of signal integrity in a Fieldbus segment.

Most host systems provide re-transmission statistics in much the same way that computers communicating via TCP/IP protocol have the ability to display the number of “lost” data packets over time. Since nearly all FF segments function with a host system connected, this becomes a built-in diagnostic tool for technicians to troubleshoot FF network segments.

Hand-held diagnostic tools are also manufactured to detect signal voltage levels, noise voltage levels, and message re-transmissions. Relcom manufactures both the model FBT-3 and model FBT-6 hand-held Fieldbus testers at the time of this writing (2009), the FBT-6 being the more capable of the two test devices.

16.7 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader’s view of this chapter’s concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Analog vs. digital signals:** analog signals have infinite resolution but are susceptible to corruption by noise. Digital signals have limited resolution but are tolerant of any noise measuring less than the difference in thresholds between the high and low states.
- **Superposition theorem:** any linear, bilateral electrical network with multiple sources may be analyzed by taking one source at a time (while replacing all other sources with their internal impedance values) and analyzing all voltages and currents, then superimposing (summing) those voltage and current values to obtain the voltages and currents with all sources active. Relevant to analyzing FOUNDATION Fieldbus H1 networks carrying DC power plus AC signals simultaneously.
- **Transmission lines:** short-duration (pulsed) electrical signals travel along a cable at nearly the speed of light, reflecting off the end of that cable if not properly terminated. Relevant to signal cables carrying high-frequency signals.

References

ANSI/ISA-5.1-2009, Instrumentation Symbols and Identification, Research Triangle Park, NC, 2009.

“Fieldbus Book – A Tutorial” (TI 38K02A01-01E) 1st Edition , Yokogawa Electric Corporation, Tokyo, Japan, 2001.

“FOUNDATION Fieldbus Application Guide – 31.25 kbit/s Intrinsically Safe Systems” (AG 163) Revision 2.0, The Fieldbus Foundation, Austin, TX, 2004.

“FOUNDATION Fieldbus Blocks” (00809-0100-4783) Revision BA, Rosemount, Inc., Chanhassen, MN, 2000.

“FOUNDATION Fieldbus System Engineering Guidelines” (AG 181) Revision 2.0, The Fieldbus Foundation, Austin, TX, 2004.

“FOUNDATION Specification System Architecture” (FF 581) Revision FS 1.1, The Fieldbus Foundation, Austin, TX, 2000.

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Software and Digital Networks*, Third Edition, CRC Press, New York, NY, 2002.

“Model 3051 Transmitter with FOUNDATION Fieldbus” (00809-0100-4774) Revision AA, Rosemount, Inc., Chanhassen, MN, 1999.

“RSFieldbus – Configuring and Programming Foundation Fieldbus Devices Application Guide” (RSFBUS-AT001A-EN-E), Rockwell Software, Inc., Milwaukee, WI, 2004.

Smith, John I., *Modern Operational Circuit Design*, Wiley-Interscience, John Wiley & Sons, Inc., New York, NY, 1971.

Park, John; Mackay, Steve; Wright, Edwin; *Practical Data Communications for Instrumentation and Control*, IDC Technologies, published by Newnes (an imprint of Elsevier), Oxford, England, 2003.

“Rosemount 3095 MultiVariable Mass Flow Transmitter with HART or FOUNDATION Fieldbus Protocol” (00809-0100-4716) Revision JA, Rosemount, Inc., Chanhassen, MN, 2008.

“The FOUNDATION Fieldbus Primer” Revision 1.1, Fieldbus Inc., Austin, TX, 2001.

“Wiring and Installation 31.25 kbit/s, Voltage Mode, Wire Medium Application Guide” (AG-140) Revision 1.0, Fieldbus Foundation, Austin, TX, 2000.

Chapter 17

Wireless instrumentation

One of the most significant technological innovations in industrial instrumentation of late has been the introduction of *radio-based* or *wireless* field instrumentation. Although this technology is simply too immature at the time of this writing (2011) to displace many wired analog and digital systems, the day is undoubtedly coming when wireless instruments will be one of the major technologies of choice for industrial applications.

Wireless field instruments are naturally digital devices, and as such possess all the innate advantages of digital instrumentation: self-diagnostics, multivariable reporting, duplex communication, etc. Furthermore, wireless instruments (at least in theory) lack some of the major limitations of wired digital instruments: slow data rates, low node counts, and energy limitations for classified areas. The single most significant weakness of current wireless field instrument technology appears to be power. With chemical batteries being the only power source, data update times must be limited to a snail's pace in order to conserve battery life. With the ongoing development of "energy harvesting" devices to locally power wireless field instruments, we may very well see this technology leap ahead of fieldbus and wired-HART instruments.

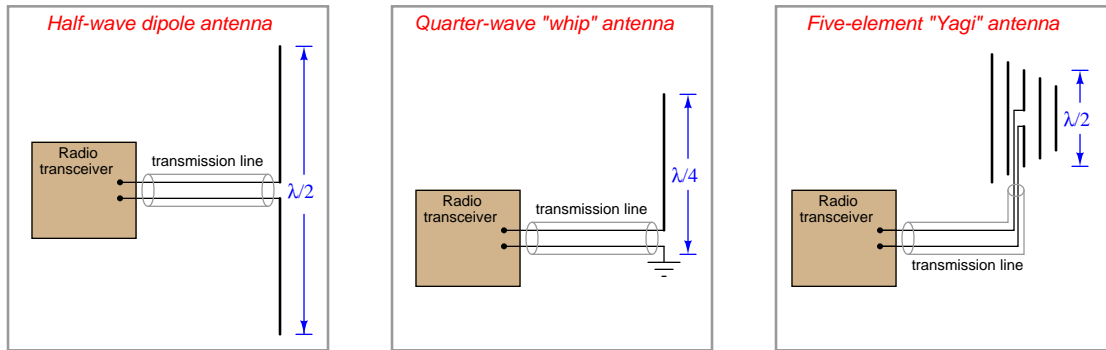
This chapter focuses on two strands of wireless technology: wireless field instruments (e.g. transmitters, control valves), and long-range wireless data links such as those used in SCADA systems. At the present time, *WirelessHART* (IEC standard 62591) is the dominant standard for radio-based field instruments, with multiple manufacturers already offering interoperable products. Exciting times, these are!

17.1 Radio systems

"Radio" systems use electromagnetic fields to communicate information over long distances through open space. This section explores some of the basic components common to all radio systems, as well as the mathematical analyses necessary to predict the performance of radio communication.

17.1.1 Antennas

A *radio wave* is a form of electromagnetic radiation, comprised of oscillating electric and magnetic fields. An *antenna*¹ is nothing more than a conductive structure designed to emit radio waves when energized by a high-frequency electrical power source, and/or generate high-frequency electrical signals when intercepting radio waves. Three common antenna designs appear here:



The Yagi antenna, with its “director” and “reflector” elements fore and aft of the dipole element, exhibits a high degree of *directionality*, whereas the dipole and “whip” antennas tend to emit and receive electromagnetic waves equally well in all directions perpendicular to their axes. Directional antennas are ideal for applications such as radar, and also in point-to-point communication applications. Omnidirectional antennas such as the dipole and whip are better suited for applications requiring equal sensitivity in multiple directions.

A photograph of an actual Yagi antenna used in a SCADA system appears here:



¹For a more detailed discussion of antennas and their electrical characteristics, refer to section 5.11 beginning on page 485.

The wavelength (λ) of any wave is its propagation velocity divided by its frequency. For radio waves, the propagation velocity is the speed of light (2.99792×10^8 meters per second, commonly represented as c), and the frequency is expressed in Hertz:

$$\lambda = \frac{c}{f}$$

Antenna dimensions are related to signal wavelength because antennas work most effectively in a condition of electrical *resonance*. In other words, the physical size of the antenna is such that it will electrically resonate at certain frequencies: a *fundamental* frequency as well as the *harmonics* (integer-multiples) of that fundamental frequency. For this reason, antenna size is inversely proportional to signal frequency: low-frequency antennas must be large, while high-frequency antennas may be small.

For example, a quarter-wave “whip” antenna designed for a 900 MHz industrial transceiver application will be approximately² 8.3 centimeters in length. The same antenna design applied to an AM broadcast radio transmitter operating at 550 kHz would be approximately *136 meters* in length!

The following photograph shows a *half-wave* “whip” antenna, located at the roofline of a building. The additional length of this design makes it more efficient than its quarter-wave cousin. This particular antenna stands approximately one meter in length from connector to tip, yielding a full wavelength value (λ) of 2 meters, equivalent to 150 MHz:




²Due to the “end effect” of lumped capacitance at the tip of the antenna, an actual quarter-wave antenna needs to be slightly shorter than an actual quarter of the wavelength. This holds true for dipoles and other antenna designs as well.

17.1.2 Decibels

One of the mathematical tools popularly used in radio-frequency (RF) work is the *common logarithm*, used to express power ratios in a unit called the *decibel*. The basic idea of decibels is to express a comparison of two electrical powers in logarithmic terms. Every time you see the unit of “decibel” you can think: *this is an expression of how much greater (or how much smaller) one power is to another*. The only question is which two powers are being compared.

Electronic amplifiers are a type of electrical system where comparisons of power are useful. Students of electronics learn to compare the output power of an amplifier against the input power as a unitless ratio, called a *gain*. Take for example an electronic amplifier with a signal input of 40 milliwatts and a signal output of 18.4 watts:



An alternative way to express the gain of this amplifier is to do so using the unit of the *Bel*, defined as the common logarithm of the gain ratio:

$$\log\left(\frac{P_{out}}{P_{in}}\right) = \log\left(\frac{18.4\text{ W}}{40\text{ mW}}\right) = 2.66276\text{ B}$$

When you see an amplifier gain expressed in the unit of “Bel”, it’s really just a way of saying “The output signal coming from this amplifier is x powers of ten greater than the input signal.” An amplifier exhibiting a gain of 1 Bel outputs 10 times as much power as the input signal. An amplifier with a gain of 2 Bels boosts the input signal by a factor of 100. The amplifier shown above, with a gain of 2.66276 Bels, boosts the input signal 460-fold.

At some point in technological history it was decided that the “Bel” (B) was too large and unwieldy of a unit, and so it became common to express powers in fractions of a Bel instead: the *deciBel* (1 dB = $\frac{1}{10}$ of a Bel). Therefore, this is the form of formula you will commonly see for expressing RF powers:

$$\text{dB} = 10 \log\left(\frac{P_{out}}{P_{in}}\right)$$

The gain of our hypothetical electronic amplifier, therefore, would be more commonly expressed as 26.6276 dB rather than 2.66276 B, although either expression is technically valid³.

³It is interesting to note that although the “Bel” is a metric unit, it is seldom if ever used without the metric prefix

An operation students often struggle with is converting a decibel figure back into a ratio, since the concept of logarithms seems to be universally perplexing. Here I will demonstrate how to algebraically manipulate the decibel formula to solve for the power ratio given a dB figure.

First, we will begin with the decibel formula as given, solving for a value in decibels given a power ratio:

$$\text{dB} = 10 \log(\text{Ratio})$$

If we wish to solve for the ratio, we must “undo” all the mathematical operations surrounding that variable. One way to determine how to do this is to reverse the order of operations we would follow if we knew the ratio and were solving for the dB value. After calculating the ratio, we would then take the logarithm of that value, and then multiply that logarithm by 10: start with the ratio, then take the logarithm, then multiply last. To un-do these operations and solve for the ratio, we must un-do each of these operations in reverse order. First, we must un-do the multiplication (by dividing by 10):

$$\frac{\text{dB}}{10} = \frac{10 \log(\text{Ratio})}{10}$$

$$\frac{\text{dB}}{10} = \log(\text{Ratio})$$

Next, we must un-do the logarithm function by applying its mathematical inverse to both sides of the formula – making each expression a power of 10:

$$10^{\frac{\text{dB}}{10}} = 10^{\log(\text{Ratio})}$$

$$10^{\frac{\text{dB}}{10}} = \text{Ratio}$$

To test our algebra, we can take the previous decibel value for our hypothetical RF amplifier and see if this new formula yields the original gain ratio:

$$\text{Ratio} = 10^{\frac{26.6276 \text{ dB}}{10}}$$


$$\text{Ratio} = 10^{2.66276} \text{ B}$$

$$\text{Ratio} = 460$$

Sure enough, we arrive at the correct gain ratio of 460, starting with the decibel gain figure of 26.6276 dB.

“deci” ($\frac{1}{10}$). One could express powers in microbels, megabels, or any other metric prefix desired, but it is never done in industry: only the decibel is used.

We may also use decibels to express power *losses* in addition to power *gains*. Here, we see an example of a radio-frequency (RF) signal cable losing power along its length⁴, such that the power out is less than the power in:



$$10 \log \left(\frac{P_{out}}{P_{in}} \right) = 10 \log \left(\frac{37\text{ mW}}{40\text{ mW}} \right) = -0.3386\text{ dB}$$


Contrasting this result against the previous result (with the amplifier) we see a very important property of decibel figures: any power *gain* is expressed as a *positive* decibel value, while any power *loss* is expressed as a *negative* decibel value. Any component outputting the exact same power as it takes in will exhibit a “gain” value of 0 dB (equivalent to a gain *ratio* of 1).

Remember that Bels and decibels are nothing more than logarithmic expressions of “greater than” and “less than”. Positive values represent powers that are *greater* while negative values represent powers that are *lesser*. Zero Bel or decibel values represent *no change* (neither gain nor loss) in power.

A couple of simple decibel values are useful to remember for approximations, where you need to quickly estimate decibel values from power ratios (or vice-versa). Each addition or subtraction of 10 dB exactly represents a 10-fold multiplication or division of power ratio: e.g. +20 dB represents a power ratio gain of $10 \times 10 = 100$, whereas -30 dB represents a power ratio reduction of $\frac{1}{10} \times \frac{1}{10} \times \frac{1}{10} = \frac{1}{1000}$. Each addition or subtraction of 3 dB approximately represents a 2-fold multiplication or division of power ratio: e.g. +6 dB is approximately equal to a power ratio gain of $2 \times 2 = 4$, whereas -12 dB is approximately equal to a power ratio reduction of $\frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} \times \frac{1}{2} = \frac{1}{16}$. We may combine $\pm 10\text{ dB}$ and $\pm 3\text{ dB}$ increments to come up with ratios that are products of 10 and 2: e.g. +26 dB is approximately equal to a power ratio gain of $10 \times 10 \times 2 \times 2 = 400$.

⁴The dominant mode of energy dissipation in an RF cable is *dielectric heating*, where the AC electric field between the cable conductors excites the molecules of the conductor insulation. This energy loss manifests as heat, which explains why there is less RF energy present at the load end of the cable than is input at the source end of the cable.

Observe what happens if we combine a “gain” component with a “loss” component and calculate the overall power out versus power in:



The overall gain of this RF amplifier and cable system expressed as a ratio is equal to the *product* of the individual component gain/loss ratios. That is, the gain ratio of the amplifier *multiplied* by the loss ratio of the cable yields the overall power ratio for the system:

$$\text{Overall gain} = \frac{17.02\text{ W}}{40\text{ mW}} = (460)(0.925) = 425.5$$

The overall gain may be alternatively expressed as a decibel figure, in which case it is equal to the *sum* of the individual component decibel values. That is, the decibel gain of the amplifier *added* to the decibel loss of the cable yields the overall decibel figure for the system:

$$\text{Overall gain} = 10 \log \left(\frac{17.02\text{ W}}{40\text{ mW}} \right) = 26.6276\text{ dB} + (-0.3386\text{ dB}) = 26.2890\text{ dB}$$

It is often useful to be able to estimate decibel values from power ratios and vice-versa. If we take the gain ratio of this amplifier and cable system (425.5) and round it down to 400, we may easily express this gain ratio as an expanded product of 10 and 2:

$$425.5 \approx 400 = (10) \times (10) \times (2) \times (2)$$

Knowing that every 10-fold multiplication of power ratio is an addition of +10 dB, and that every 2-fold multiplication of power is an addition of +3 dB, we may express the expanded product as a sum of decibel values:

$$(10) \times (10) \times (2) \times (2) = (10\text{ dB}) + (10\text{ dB}) + (3\text{ dB}) + (3\text{ dB}) = 26\text{ dB}$$

Therefore, our power ratio of 425.5 is approximately equal to +26 decibels.


Decibels always represent comparisons of power, but that comparison need not always be P_{out}/P_{in} for a system component. We may also use decibels to express an amount of power compared to some standard reference. If, for example, we wished to express the input power to our hypothetical RF amplifier (40 milliwatts) using decibels, we could do so by comparing 40 mW against a standard “reference” power of exactly 1 milliwatt. The resulting decibel figure would be written as “dBm” in honor of the 1 milliwatt reference:

$$P_{in} = 10 \log \left(\frac{40 \text{ mW}}{1 \text{ mW}} \right) = 16.0206 \text{ dBm}$$

The unit of “dBm” literally means the amount of dB “greater than” 1 milliwatt. In this case, our input signal of 40 milliwatts is 16.0206 dB greater than a standard reference power of exactly 1 milliwatt. The output power of that amplifier (18.4 watts) may be expressed in dBm as well:

$$P_{out} = 10 \log \left(\frac{18.4 \text{ W}}{1 \text{ mW}} \right) = 42.6482 \text{ dBm}$$

A signal power of 18.4 watts is 42.6482 dB greater than a standard reference power of exactly 1 milliwatt, and so it has a decibel value of 42.6482 dBm.



Notice how the output and input powers expressed in dBm relate to the power gain of the amplifier. Taking the input power and simply *adding* the amplifier’s gain factor yields the amplifier’s output power in dBm:

$$P_{in}(\text{dB}) + P_{gain}(\text{dB}) = P_{out}(\text{dB})$$


$$16.0206 \text{ dBm} + 26.6276 \text{ dB} = 42.6482 \text{ dBm}$$

An electronic signal that begins 16.0206 dB greater than 1 milliwatt, when boosted by an amplifier gain of 26.6276 dB, will become 42.6482 dB greater than the original reference power of 1 milliwatt.

We may alternatively express all powers in this hypothetical amplifier in reference to a 1-watt standard power, with the resulting power expressed in units of “dBW” (decibels greater than 1 watt):

$$P_{in} = 10 \log \left(\frac{40 \text{ mW}}{1 \text{ W}} \right) = -13.9794 \text{ dBW}$$

$$P_{out} = 10 \log \left(\frac{18.4 \text{ W}}{1 \text{ W}} \right) = 12.6482 \text{ dBW}$$



Note how the input power of 40 milliwatts equates to a negative dBW figure because 40 milliwatts is *less* than the 1 watt reference, and how the output power of 18.4 watts equates to a positive dBW figure because 18.4 watts is *more* than the 1 watt reference. A positive dB figure means “more than” while a negative dB figure means “less than.”

Note also how the output and input powers expressed in dBW still relate to the power gain of the amplifier by simple addition, just as they did when previously expressed in units of dBm. Taking the input power in units of dBW and simply *adding* the amplifier’s gain factor yields the amplifier’s output power in dBW:

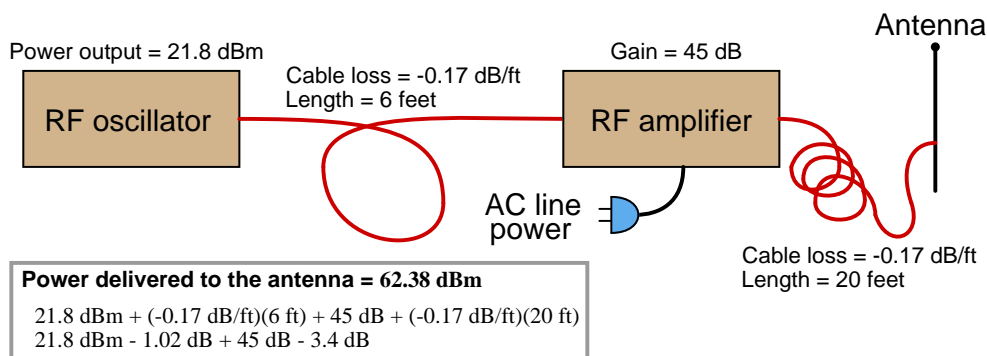
$$P_{in}(\text{dB}) + P_{gain}(\text{dB}) = P_{out}(\text{dB})$$

$$-13.9794 \text{ dBW} + 26.6276 \text{ dB} = 12.6482 \text{ dBW}$$

An electronic signal that begins 13.9794 dB less than 1 watt, when boosted by an amplifier gain of 26.6276 dB, will become 12.6482 dB greater than the original reference power of 1 watt.

This is one of the major benefits of using decibels to express powers: we may very easily calculate power gains and losses by summing a string of dB figures, each dB figure representing the power gain or power loss of a different system component. Normally, any conflation of *ratios* involves multiplication and/or division of those ratios, but with decibels we may simply add and subtract. One of the interesting mathematical properties of logarithms is that they “transform⁵” one type of problem into a simpler type: in this case, a problem of multiplying ratios into a (simpler) problem of adding decibel figures.

For example, we may express the power lost in an RF transmission line (two-conductor cable) in terms of decibels per foot. Most of this power loss is due to dielectric heating, as the high-frequency electric field of the RF signal causes polarized molecules in the cable insulation to vibrate and dissipate energy in the form of heat⁶. The longer the cable, of course, the more power will be lost this way, all other factors being equal. A type of cable having a loss figure of -0.15 decibels per foot at a signal frequency of 2.4 GHz will suffer -15 dB over 100 feet, and -150 dB over 1000 feet. To illustrate how decibels may be used to calculate power at the end of an RF system, accounting for various gains and losses along the way using decibel figures:




⁵In fact, logarithms are one of the simplest examples of a *transform function*, converting one type of mathematical problem into another type. Other examples of mathematical transform functions used in engineering include the *Fourier transform* (converting a time-domain function into a frequency-domain function) and the *Laplace transform* (converting a differential equation into an algebraic equation).

⁶This is precisely how a microwave oven works: water molecules are polar (that is to say, the electrical charges of the hydrogen and oxygen atoms are not symmetrical, and therefore each water molecule has one side that is more positive and an opposite side that is more negative), and therefore vibrate when subjected to electromagnetic fields. In a microwave oven, RF energy in the gigahertz frequency range is aimed at pieces of food, causing the water molecules within the food to heat up, thus indirectly heating the rest of the food. This is a practical example of an RF system where losses are not only expected, but are actually a design objective! The food represents a load to the RF energy, the goal being complete dissipation of all incident RF energy with no leakage outside the oven. In RF cable design, however, dissipative power losses are something to be avoided, the goal being complete delivery of RF power to the far end of the cable.


17.1.3 Antenna radiation patterns

Different antenna designs are unequal with regard to how well they radiate (and receive) electromagnetic energy. Every antenna design has a *pattern* of radiation and sensitivity: some directions in which it is maximally effective and other directions where it is minimally effective.

Some common antenna types and radiation patterns are shown in the following illustrations, the relative radii of the shaded areas representing the degree⁷ of effectiveness in those directions away from or toward the antenna:




⁷One should not think that the outer edges of the shaded radiation patterns represents some "hard" boundary beyond which no radiation is emitted (or detected). In reality, the radiation patterns extend out to infinity (assuming otherwise empty space surrounding the antenna). Instead, the size of each shaded area simply represents how effective the antenna is in that direction compared to other directions. In the case of the vertical whip and dipole antennas, for instance, the radiation patterns show us that these antennas have *zero* effectiveness along the vertical (Y) axis centerline. To express this in anthropomorphic terms, these antenna designs are "deaf and mute" in those directions where the radiation pattern is sketched having zero radius from the antenna center.




It should be noted that the radiation patterns shown here are approximate, and may modify their shapes if the antenna is operated harmonically rather than at its fundamental frequency.

A basic principle in antenna theory called *reciprocity* states that the efficiency of an antenna as a radiator of electromagnetic waves mirrors its efficiency as a collector of electromagnetic waves. In other words, a good transmitting antenna will be a good receiving antenna, and an antenna having a preferred direction of radiation will likewise be maximally sensitive to electromagnetic waves approaching from that same direction. To use a Yagi as an example:



Related to reciprocity is the concept of equivalent orientation between transmitting and receiving antennas for maximum effectiveness. The electromagnetic waves emitted by a transmitting antenna are *polarized* in a particular orientation, with the electric and magnetic fields perpendicular to each other. The same design of antenna will be maximally receptive to those waves if its elements are similarly oriented. A simple rule to follow is that antenna pairs should always be *parallel* to each other in order to maximize reception, in order that the electric and magnetic fields emanating from the wires of the transmitting antenna will "link" properly with the wires of the receiving antenna(s). If the goal is optimum communication in any direction (omnidirectionality), dipole and whip antennas should be arranged *vertically* so that all antenna conductors will be parallel to each other regardless of their geographic location.

Yagi antenna pairs may be horizontally or vertically oriented, so long as the transmitting and receiving Yagis are both mounted with the same polarization and face each other. In industrial SCADA radio applications, Yagi antennas are generally oriented with the dipole wires vertical, so that they may be used in conjunction with omnidirectional whip or dipole antennas. An illustration of such use is shown here, with multiple “Remote Terminal Unit” (RTU) transceivers communicating with a central “Master Terminal Unit” (MTU) transceiver:



Here, all the Yagi antennas on the RTUs are vertically oriented, so that they will match the polarization of the MTU’s whip antenna. The Yagi antennas all face in the direction of the MTU for optimum sensitivity. The MTU – which must broadcast to and receive from all the RTUs – really needs an omnidirectional antenna. The RTUs – which need only communicate with the one MTU and not with each other – work best with highly directional antennas.

If the MTU were equipped with a Yagi antenna instead of a whip, it would only communicate well with one of the RTUs, and perhaps not at all with some of the others. If all RTUs were equipped with whip antennas instead of Yagis, they would not be as receptive to the MTU’s broadcasts (lower gain), and each RTU would require greater transmitter power to transmit effectively to the MTU.


Another important principle to employ when locating any antenna is to keep it far away from any conductive surfaces or objects, including soil. Proximity to any conductive mass distorts an antenna’s radiation pattern, which in turn affects how well it can transmit and receive in certain directions. If there is any consistent rule to follow when setting up antennas for maximum performance, it is this: *position them as high as possible and as far away from interfering objects as possible!*

17.1.4 Antenna gain calculations

A common way to express the maximal effectiveness of any antenna design is as a ratio compared to some idealized form of antenna with a more uniform radiation pattern. As with most ratio measurements in radio technology, the standard unit for antenna gain is the *decibel* (dB), related to a ratio of powers as follows:

$$\text{Gain in dB} = 10 \log \left(\frac{P}{P_{ref}} \right)$$

The most common reference standard used to calculate antenna gain is a purely theoretical device called an *isotropic antenna*. This is an ideally omnidirectional antenna having a perfectly spherical radiation pattern:



If a directional antenna such as a Yagi radiates (and/or receives) 20 times as much power in its most sensitive direction as an isotropic antenna, it is said to have a power gain of 13.01 dBi (13.01 *decibels* more than an *isotropic*). An alternative “reference” for comparison is a half-wave dipole antenna. Decibel comparisons against a dipole are abbreviated dBd. The assumption of an isotropic antenna as the reference is so common in radio engineering, though, you often see antenna gains expressed simply in units of dB. The assumption of isotropic reference (dBi) for antenna gain expressions is analogous to the assumption of “RMS” measurements in AC circuits rather than “peak” or “peak-to-peak”: if you see an AC voltage expressed without any qualifier (e.g. “117 volts AC”), it is generally assumed to be an RMS measurement.

Whip antennas typically exhibit an optimal gain of 6 dBi (6 dB being approximately equal to a 4-fold magnification compared to an isotropic antenna), while Yagis may achieve up to 15 dBi (approximately equal to a 32-fold magnification). Parabolic “dish” antenna designs such as those used in microwave communication systems may achieve gains as high as 30 dBi (\approx 1000-fold magnification). Since antenna gain is not a real amplification of power – this being impossible according to the Law of Energy Conservation – greater antenna gain is achieved only by greater *focus* of the radiation pattern in a particular direction.

The concept of antenna gain is very easy to mis-comprehend, since it is tempting to think of any type of gain as being a true increase in power. Antenna gain is really nothing more than a way to express how *concentrated* the RF energy of a radiating⁸ antenna is in one direction compared to a truly omnidirectional antenna. An analogy to antenna gain is how a horn-style loudspeaker focuses its audio energy more than a loudspeaker lacking a horn. The horn-shaped speaker *sounds* louder than the horn-less speaker (in one direction only) because its audio energy is more focused. The two speakers may be receiving the exact same amount of electrical energy to produce sound, but the more directional of the two speakers will be more efficient transmitting sound in one direction than the other. Likewise, a horn-shaped microphone will have greater sensitivity in one direction than a comparable “omni” microphone designed to receive sound equally well from all directions. Connected to a recording device, the directional microphone seems to present a “gain” by sending a stronger signal to the recorder than the omnidirectional microphone is able to send, in that one direction.

The flip-side of high-gain antennas, loudspeakers, and microphones is how poorly they perform in directions other than their preferred direction. Any transmitting or receiving structure exhibiting a “gain” due to its focused radiation pattern must likewise exhibit a “loss” in performance when tested in directions other than its direction of focus. Referring back to directional radio antennas again, a Yagi with an advertised gain of 15 dBi (in its “preferred” direction) will exhibit a strong negative gain in the rearward direction where its ability to radiate and receive is almost non-existent.


If even more signal gain is necessary than what may be achieved by narrower radiation focus, an actual electronic amplifier may be added to an antenna assembly to boost the RF power sent to or received from the antenna. This is common for satellite antenna arrays, where the RF amplifier is often located right at the focal point of the parabolic dish. Satellite communication requires very high transmitter and receiver gains, due to inevitable signal weakening over the extremely long distances between a ground-based antenna and a satellite antenna in geosynchronous orbit around the Earth.

⁸Or – applying the principle of reciprocity – antenna gain is really nothing more than a way to express how sensitive a receiving antenna is compared to a truly omnidirectional antenna.

17.1.5 Effective radiated power

When determining the effectiveness of a radio system, one must include losses in cables, connectors, lightning arrestors, and other elements in the signal path in addition to the antenna itself. A commonly accepted way to quantify this effectiveness is to rate a radio system against a standard reference consisting of an ideal dipole antenna connected to a 1 milliwatt transmitter with no losses. Typically expressed in units of decibels, this is called the *Effective Radiated Power*, or *ERP*. If the ideal antenna model is isotropic instead of a dipole, then the calculation result is called the *Effective Isotropic Radiated Power*, or *EIRP*.

Let us consider the following example, where a 2.4 GHz radio transceiver outputs 250 milliwatts⁹ of radio-frequency (RF) power to a Yagi antenna through a type LMR 195 coaxial cable 12 feet in length. A lightning arrestor with 0.5 dB loss is also part of the cable system. We will assume an antenna gain of 9 dBi for the Yagi and a loss of 0.19 dB per foot for the LMR 195 cable:



The EIRP for this radio system is 30.2 dB: the sum of all gains and losses expressed in decibels. This means our Yagi antenna will radiate 30.2 dB (1047 times) more RF power in its most effective direction than an isotropic antenna would radiate in the same direction powered by a 1 milliwatt transmitter. Note that if our hypothetical radio system also included an RF amplifier between the transceiver and the antenna, its gain would have to be included in the EIRP calculation as well.

⁹Actual signal power is typically expressed as a decibel ratio to a reference power of either 1 milliwatt (dBm) or 1 watt (dBW). Thus, 250 mW of RF power may be expressed as $10 \log \frac{250}{1} = 23.98$ dBm or as $10 \log \frac{0.25}{1} = -6.02$ dBW. Power expressed in unit of dBm will always be 30 dB greater (1×10^3 greater) than power expressed in dBW.

A practical application of EIRP is how the Federal Communications Commission (FCC) sets limits on radio transmitters within the United States. Not only is gross transmitter power limited by law within certain frequency ranges, but also the EIRP of a transmitting station. This makes sense, since a more directional transmitting antenna (i.e. one having a greater gain value) will make it appear as though the transmitter is more powerful than it would be radiating from a less-directional antenna. If FCC limits were based strictly on transmitter power output rather than EIRP, it might still be possible for an otherwise power-compliant transmitting station to generate excessive interference through the use of a highly directional antenna. This is why it is illegal, for example, to connect a large antenna to a low-power transmitting device such as a two-way (“walkie-talkie”) radio unit: the two-way radio unit may be operated license-free only because its EIRP is low enough not to cause interference with other radio systems. If someone were to connect a more efficient antenna to this same two-way radio, its effective radiated power may increase to unacceptable levels (according to the FCC) even though the raw power output by the transmitter circuitry has not been boosted.

17.1.6 RF link budget

Electromagnetic radiation is used as a medium to convey information, to “link” data from one physical location to another. In order for this to work, the amount of signal loss between transmitter and receiver must be small enough that the signal does not become lost in radio-frequency “noise” originating from external sources and from within the radio receiver itself. We may express radio-frequency (RF) power in terms of its comparison to 1 milliwatt: 0 dBm being 1 milliwatt, 3.01 dBm being 2 milliwatts, 20 dBm being 100 milliwatts, etc. We may use dBm as an absolute measurement scale for transmitted and received signal strengths, as well as for expressing how much ambient RF noise is present (called the “noise floor” for its appearance at the bottom of a spectrum analyzer display). We may use plain dB to express relative gains and losses along the signal path.

The basic idea behind an “RF link budget” is to add *all* gains and losses in an RF system – from transmitter to receiver with all intermediate elements accounted for – to ensure there is a large enough difference between signal and noise to ensure good data communication integrity. If we account all gains as positive decibel values and all losses as negative decibel values, the signal power at the receiver will be the simple sum of all the gains and losses:

$$P_{rx} = P_{tx} + G_{total} + L_{total}$$

Where,

P_{rx} = Signal power delivered to receiver input (dBm)


P_{tx} = Transmitter output signal power (dBm)

G_{total} = Sum of all gains (amplifiers, antenna directionality, etc.), a positive dB value

L_{total} = Sum of all losses (cables, filters, path loss, fade, etc.), a negative dB value

This formula tells us how much signal power will be available at the radio receiver, but usually the purpose of calculating a link budget is to determine how much radio transmitter power will be *necessary* in order to have adequate signal strength at the receiver. More transmitter power adds expense, not only due to transmitter hardware cost but also to FCC licenses that are required if certain power limitations are exceeded. Excessive transmitter power may also create interference problems with other radio and electronic systems. Suffice it to say we wish to limit transmitter power to the minimum practical value.

In order for a radio receiver to reliably detect an incoming signal, that signal must be sufficiently greater than the ambient RF noise. All masses at temperatures above absolute zero radiate electromagnetic energy, with some of that energy falling within the RF spectrum. This *noise floor* value may be calculated¹⁰ or empirically measured using an RF spectrum analyzer as shown in this simulated illustration:



On top of the ambient noise, we also have the *noise figure* of the receiver itself (N_{rx}): noise created by the internal circuitry of the radio receiver. Thus, the minimum signal power necessary for the receiver to operate reliably ($P_{rx(min)}$) is equal to the decibel sum of the noise floor and the noise figure, by a margin called the minimum *signal-to-noise* ratio:

$$P_{rx(min)} = N_f + N_{rx} + S$$

Where,

$P_{rx(min)}$ = Minimum necessary signal power at the receiver input (dBm)

N_f = Noise floor value (dBm)

N_{rx} = Noise figure of radio receiver (dB)

S = Desired signal-to-noise ratio margin (dB)

Substituting this decibel sum into our original RF link budget formula and solving for the minimum necessary transmitter power output ($P_{tx(min)}$), we get the following result:

$$P_{tx(min)} = N_f + N_{rx} + S - (G_{total} + L_{total})$$

¹⁰Noise power may be calculated using the formula $P_n = kTB$, where P_n is the noise power in watts, k is Boltzmann's constant (1.38×10^{-23} J/K), T is the absolute temperature in Kelvin, and B is the bandwidth of the noise in Hertz. Noise power usually expressed in units of dBm rather than watts, because typical noise power values for ambient temperatures on Earth are so incredibly small.

It is common for radio receiver manufacturers to aggregate the noise floor, noise figure, and a reasonable signal-to-noise ratio into a single parameter called *receiver sensitivity*. The “sensitivity” of a radio receiver unit is the minimum amount of signal power (usually expressed in dBm) necessary at the input connector for reliable operation despite the inevitable presence of noise. If we simply express receiver sensitivity as $P_{rx(min)}$ and substitute this term for the sum of noise floor, noise figure, and signal-to-noise margin ($N_f + N_{rx} + S$) in the last formula, we see that the difference in receiver sensitivity (expressed in absolute decibels) and the sum of any gains and losses in the link (also expressed in decibels) tells us the minimum transmitter power required:

$$P_{tx(min)} = P_{rx(min)} - (G_{total} + L_{total})$$

Where,

$P_{tx(min)}$ = Minimum necessary transmitter output signal power, in dBm

$P_{rx(min)}$ = Receiver sensitivity (minimum necessary received signal power), in dBm

G_{total} = Sum of all gains (amplifiers, antenna directionality, etc.), a positive dB value

L_{total} = Sum of all losses (cables, filters, path loss, fade, etc.), a negative dB value

For digital radio receivers, sensitivity is a function of error rate: the fewer errors desired, the more signal power required. To give a practical example, one modern 900 MHz radio transceiver has a specified sensitivity of -110 dBm at a bit error rate (BER) of 10^{-4} bits (one error for every 10^4 bits received) and a sensitivity of -108 dBm at a BER of 10^{-6} bits. This relationship between signal power and error rate should make intuitive sense: the more powerful the signal compared to any background noise, the more reliably it will be received; the weaker the signal, the more it will become corrupted by noise and therefore the more errors we would expect to see over time.

Among the losses encompassed in L_{total} are *path loss* and *fade*. Path loss is the natural loss of signal strength with increasing distance from the radiation source. As electromagnetic waves propagate outward through space, they inevitably spread. The degradation in signal strength with increasing distance follows the *inverse square law*¹¹, where power decreases with the square of distance. Thus, doubling the distance from transmitting antenna to receiving antenna attenuates the signal by a factor of four ($\frac{1}{2^2}$, or -6.02 dB). Tripling the distance from transmitting antenna to receiving antenna attenuates the signal by a factor of nine ($\frac{1}{3^2}$, or -9.54 dB).

Path loss for free-space conditions is a rather simple function of distance and wavelength:

$$L_p = -20 \log \left(\frac{4\pi D}{\lambda} \right)$$

Where,

L_p = Path loss, a negative dB value

D = Distance between transmitting and receiving antennas

λ = Wavelength of transmitted RF field, in same physical unit as D

It should be emphasized that this simple path loss formula only applies to completely clear, empty space where the only mechanism of signal attenuation is the natural spreading of radio waves as they radiate away from the transmitting antenna. Path loss will be significantly greater if any objects or other obstructions lie between the transmitting and receiving antennas.

This same spreading effect also accounts for “fade,” where radio waves taking different paths destructively interfere (e.g. waves reflected off lateral objects reaching the receiving antenna out-of-phase with the straight-path waves), resulting in attenuated signal strengths in some places (but not in all). You may have personally experienced fade while driving a vehicle over long distances and listening to an analog (AM or FM) radio: sometimes a particular radio station’s signal will “fade out” as you drive and then “fade in” again while driving in the same direction, for no obvious reason (e.g. no immediate obstructions to the signal). This is due to radio waves from the station emanating in all directions, then reflecting off of large objects and/or ionized regions high in the earth’s atmosphere. There will inevitably be locations around that station where the incident wave from the transmitting antenna destructively interferes with those reflected waves, the result being regions of “dead” space where the signal is much weaker than one would expect from path loss alone.


Fade is a more difficult factor to predict than path loss, and so generally radio system designers include an adequate margin¹² to account for the effects of fade. This *fade margin* is typically 20

¹¹The inverse square law applies to any form of radiation that spreads from a point-source. In any such scenario, the intensity of the radiation received by an object from the point-source diminishes with the square of the distance from that source, simply because the rest of the radiated energy misses that target and goes elsewhere in space. This is why the path loss formula begins with a -20 multiplier rather than -10 as is customary for decibel calculations: given the fact that the inverse square law tells us path loss is proportional to the square of distance (D^2), there is a “hidden” second power in the formula. Following the logarithmic identity that exponents may be moved to the front of the logarithm function as multipliers, this means what would normally be a -10 multiplier turns into -20 and we are left with D rather than D^2 in the fraction.

¹²“Margin” is the professionally accepted term to express extra allowance provided to compensate for unknowns. A more colorful phrase often used in the field to describe the same thing is *fudge factor*.

dB to 30 dB, although it can be greater in scenarios where there are many signal paths due to reflections.

To illustrate, we will calculate the RF link budget for a 900 MHz radio transmitter/receiver pair directionally oriented toward each other with Yagi antennas. All sources of signal gain and loss will be accounted for, including the “path loss” of the RF energy as it travels through open air. The gains and losses of all elements are shown in the following illustration:



The path loss value shown in the illustration is a calculated function of the 900 MHz wavelength ($\lambda = \frac{c}{f} = 0.3331$ meters) and the distance between antennas (500 feet = 152.4 meters), assuming a completely obstruction-free path between antennas:

$$L_p = -20 \log \left(\frac{4\pi(152.4)}{0.3331} \right) = -75.19 \text{ dB}$$

According to the receiver manufacturer’s specifications, the receiver in this system has a sensitivity of -75 dBm, which means our transmitter must be powerful enough to deliver an RF signal at least as strong as -75 dBm at the receiver’s input connector in order to reliably communicate data. Inserting this receiver sensitivity figure into our RF link budget formula:

$$P_{tx(min)} = P_{rx(min)} - (G_{total} + L_{total})$$

$$P_{tx(min)} = -75 \text{ dBm} - (G_{total} + L_{total})$$

Now we need to tally all the gains and losses between the transmitter and the receiver. We will use a value of -20 dB for fade margin (i.e. our budget will leave room for up to 20 dB of power loss due to the effects of fade):

Gain or Loss	Decibel value
Transmitter cable loss	-5.64 dB
Transmitter arrestor loss	-0.5 dB
Transmitter antenna gain	$+12.1$ dBi
Path loss	-75.19 dB
Fade margin	-20 dB
Receiver antenna gain	$+12.1$ dBi
Receiver arrestor loss	-0.5 dB
Receiver cable loss	-4.23 dB
$G_{total} + L_{total}$	-81.86 dB

Inserting the decibel sum of all gains and losses into our RF link budget formula:

$$P_{tx(min)} = -75 \text{ dBm} - (-81.86 \text{ dB})$$

$$P_{tx(min)} = 6.86 \text{ dBm}$$

Converting a dBm value into milliwatts of RF power means we must manipulate the dBm power formula to solve for P_{mW} :

$$P_{dBm} = 10 \log \left(\frac{P_{mW}}{1 \text{ mW}} \right)$$

$$P_{mW} = 1 \text{ mW} \times 10^{\left(\frac{P_{dBm}}{10}\right)}$$

$$P_{tx} = 1 \text{ mW} \times 10^{\left(\frac{6.86}{10}\right)} = 4.85 \text{ milliwatts}$$


At this point we would do well to take stock of the assumptions intrinsic to this calculation. Power gains and losses inherent to the components (cables, arrestors, antennas) are quite certain because these are tested components, so we need not worry about these figures too much. What we know the least about are the environmental factors: noise floor can change, path loss *will* differ from our calculation if there is any obstruction near the signal path or under certain weather conditions (e.g. rain or snow dissipating RF energy), and fade loss is known to change dynamically as moving objects (people, vehicles) pass anywhere between the transmitter or receiver antennas. Our RF link budget calculation is really only an estimate of the transmitter power needed to get the job done.

How then do we improve our odds of building a reliable system? One way is to over-build it by equipping the transmitter with more power than the most pessimistic link budget predicts. However, this can cause other problems such as interference with nearby electronic systems if we are not careful. A preferable method is to conduct a *site test* where the real equipment is set up in the field and tested to ensure adequate received signal strength.


17.1.7 Fresnel zones

One of the many factors affecting RF link power transfer from transmitter to receiver is the openness of the signal path between the transmitting and receiving antennas. As shown in the previous subsection, *path loss* is relatively simple to calculate given the assumption of totally empty space between the two antennas. With no obstacles in between, path loss is simply a consequence of wave dispersion (spreading). Unless we are calculating an RF link budget between two airplanes or two spacecraft, though, there really is no such thing as totally empty space between the transmitting and receiving antennas. Buildings, trees, vehicles, and even the ground are all objects potentially disrupting what would otherwise be completely open space between transmitter and receiver.

A common expression in microwave radio communication is *line-of-sight*, or *LoS*. The very wording of this phrase evokes an image of being able to see a straight path between transmitting and receiving antennas. Obviously, if one does not have a clear line-of-sight path between antennas, the signal path loss will definitely be more severe than through open space. However, a common error is in thinking that the mere existence of an unobstructed straight line between antennas is all that is needed for unhindered communication, when in fact nothing could be further from the truth:



In fact, the free space necessary to convey energy in electromagnetic wave form takes on the form of football-shaped zones: the first one solid followed by annular (hollow) zones concentrically surrounding the first. These elliptical volumes are called *Fresnel zones*:



The precise shapes of these Fresnel zones are a function of wavelength and distance between antennas, *not* the size or configurations of the antennas themselves. In other words, you cannot change the Fresnel zone sizes or shapes by merely altering antenna types. This is because the

Fresnel zones do not actually map the distribution of electromagnetic fields, but rather map the free space we need to keep clear between antennas¹³.

If any object protrudes at all into any Fresnel zone, it diminishes the signal power communicated between the two antennas. In microwave communication (GHz frequency range), the inner-most Fresnel zone carries most of the power, and is therefore the most important from the perspective of interference. Keeping the inner-most Fresnel zone absolutely clear of interference is essential for maintaining ideal path-loss characteristics (equivalent to open-space). A common rule followed by microwave system designers is to try to maintain an inner Fresnel zone that is at least 60% free of obstruction.

In order to design a system with this goal in mind, we need to have some way of calculating the width of that Fresnel zone. Fortunately, this is easily done with the following formula:

$$r = \sqrt{\frac{n\lambda d_1 d_2}{D}}$$

Where,

r = Radius of Fresnel zone at the point of interest

n = Fresnel zone number (an integer value, with 1 representing the first zone)


d_1 = Distance between one antenna and the point of interest

d_2 = Distance between the other antenna and the point of interest

D = Distance between both antennas

λ = Wavelength of transmitted RF field, in same physical unit as D , d_1 , and d_2

Note: the units of measurement in this formula may be any unit of length, so long as they are all the same unit.



¹³The physics of Fresnel zones is highly non-intuitive, rooted in the wave-nature of electromagnetic radiation. It should be plain to see, though, that Fresnel zones cannot describe the actual electromagnetic field pattern between two antennas, because we know waves tend to spread out over space while Fresnel zones converge at each end. Likewise, Fresnel zones vary in size according to the distance between two antennas which we know radiation field patterns do not. It is more accurate to think of Fresnel zones as *keep-clear* areas necessary for reliable communication between two or more antennas rather than actual field patterns.

To illustrate by example, let us calculate the radius of the first Fresnel zone for two microwave antennas operating at 2.4 GHz, separated by one mile (1.609 km), at the widest portion of that zone. The wavelength of a 2.4 GHz signal is 0.1249 meters, from the formula $\lambda = \frac{c}{f}$. Distances d_1 and d_2 will both be equal to one-half of the total distance, since the widest portion of the Fresnel zone will be exactly at its mid-point between the two antennas ($d_1 = d_2 = 804.7$ meters). Here, we solve for r as follows:

$$r = \sqrt{\frac{n\lambda d_1 d_2}{D}}$$

$$r = \sqrt{\frac{(1)(0.1249)(804.7^2)}{1609}}$$

$$r = 7.089 \text{ meters}$$

Consider for a moment the significance of this dimension. At the very least, it means the antennas must be mounted this high off the ground in order to avoid having the most important Fresnel zone contact the earth itself (assuming level terrain), not to mention any objects above ground level such as buildings, vehicles, or trees between the two antennas. Consider also that the Fresnel zone is football-shaped, and therefore this 7.089 meter radius extends horizontally from the centerline connecting both antennas as well as vertically. This means that in order for this Fresnel zone to be untouched, there must be a clear path *14.18 meters wide* in addition to the antennas being at least 7.089 meters off the ground! If we were to consider a 900 MHz signal – another common frequency for industrial wireless devices – the minimum height above ground would be 11.58 meters, and the minimum clear path width 23.16 meters!


As you can see, “line of sight” is not as simple as it may first appear.

17.2 WirelessHART

An exciting development in industrial instrumentation is the *WirelessHART* radio communication standard, specifically designed for field instrument use (e.g. transmitters, valve positioners) as opposed to general data communication. The IEC (International Electrotechnical Commission) has codified the *WirelessHART* standard as IEC 62591.

17.2.1 Introduction to *WirelessHART*

WirelessHART is a subset of the HART industrial instrument communication standard as of version 7, communicating process data over 2.4 GHz radio waves. Individual instruments communicate with a common “gateway” device serving as an interface between the wireless network and a wired network or a host control system. In addition to this, though, individual *WirelessHART* devices also form links with one another, so that the network data routes look like a “mesh” with all nearby nodes interconnected in addition to connecting with the gateway:




In a mesh network, devices (nodes) perform double-duty as *repeaters* to relay data from other instruments to the gateway as needed. In other words, data transmitted from one *WirelessHART* instrument may not be directly received by the gateway device if that path is blocked or too far away. Instead, the data may “hop” from one device to another nearby, which then re-broadcasts that information to the gateway via a clearer path.

The purpose of a mesh network is to provide redundant data pathways in case of device failure or changes in the environment interrupting radio communication between devices. In this way, data

packets may be re-routed to the gateway if the shortest route fails, in a manner similar to how Terminal Control Protocol (TCP) and Internet Protocol (IP) work together to route data segments from source to destination over the “mesh” of the Internet. This feature is often referred to in *WirelessHART* technical literature as the *self-healing* property of the mesh network.

According to the HART Foundation, reliability for a well-designed *WirelessHART* mesh network is 99.7300204% minimum, and typically greater than 99.9999998%.

With each *WirelessHART* field instrument capable of functioning as a radio repeater, the potential exists to form wireless networks larger in size than the maximum broadcast/reception range of any one device. This illustration shows what is possible¹⁴:



¹⁴Some obvious connecting paths between field devices have been omitted from this illustration if the path length exceeds a certain maximum distance. As you can see, the instruments in the far-left cluster *must* rely on data packet relaying by instruments closer to the gateway, since they themselves are too far away from the gateway to directly communicate.

An important consideration when planning a *WirelessHART* network is battery life. With the main purpose of wireless field instruments being the elimination of wired connections to the host system, the field instruments cannot rely on a host system for their electrical power needs. Lithium-based batteries currently fulfill this role as primary¹⁵ power source, with life expectancies of several years. Interestingly, the amount of energy required by a *WirelessHART* device to transmit radio-frequency data is small compared to the energy required to power its essential instrument functions (e.g. pressure measurement, temperature measurement). This means a *WirelessHART* device operating as a radio repeater (in addition to being a measurement device) adds little burden to its battery.

Perhaps the greatest challenge in sustaining any wireless field instrument network is ensuring network integrity despite unending changes in the physical environment around the instruments. Remember that this is an *industrial*, field-instrument wireless network designed to be installed in less-than-ideal physical environments. These wireless devices must somehow reliably communicate with each other despite interference from high-power electrical devices (e.g. variable-frequency motor drive units), while mounted on or near metal objects such as girders, pipes, pipe racks, large vessels, motors, enclosures, shelters, and electrical conduits. Even the ground of an industrial environment can be an impediment to robust radio communication: steel-reinforced concrete and electrical grounding grids form what is essentially a solid “ground plane” that will interfere with *WirelessHART* devices mounted too close to ground level. Added to all this spatial complexity is the continual presence of large vehicles and other moving machines (cranes, forklifts, manlifts, etc.). It is not uncommon for scaffolding to be temporarily erected for maintenance work in industrial areas, presenting yet one more obstacle for RF signals.

In answer to these challenges is an integral and essential component of a *WirelessHART* network called the *Network Manager*: an advanced digital algorithm usually executed by the network gateway’s microprocessor. The purpose of the Network Manager is to manage the details of the network automatically, “tuning” various parameters for optimum reliability and data throughput. Among other tasks, the Network Manager assigns “timeslots” for individual devices to transmit, determines the frequency-hopping schedule, detects and authenticates new devices added to the network, dynamically adjusts device transmission power, and selects alternative routes between devices.

In a sense, the Network Manager in a *WirelessHART* network continually audits and tunes the RF system in an attempt to maximize reliability. The Network Manager’s functionality does not substitute for good planning during the design phase of the *WirelessHART* network, but it does eliminate the need for a human technician or engineer to continuously monitor the network’s performance and make the small adjustments necessary to compensate for changing conditions. When changes occur in a *WirelessHART* network that cannot be compensated by the Network Manager, the real-time statistics provided by the Network Manager are invaluable to the technician or engineer assigned to update the network.

¹⁵Another exciting technological development paralleling the implementation of *WirelessHART* in industry is that of *energy-harvesting* devices to generate DC electricity from nearby energy sources such as vibrating machines (mechanical motion), hot pipes (thermal differences), photovoltaic (solar) panels, and even small wind generators. Combined with rechargeable batteries to sustain instrument operation during times those energy sources are not producing, energy-harvesters promise great extension of battery life for wireless instruments of all types.

17.2.2 *Wireless*HART network protocol

The OSI reference model will be used here to identify and describe various features of the *Wireless*HART protocol.

Physical Layer

- 2.4 GHz to 2.5 GHz (“ISM” – Industrial, Scientific, Medical) signal band
- O-QPSK modulation (offset quadrature phase-shift keying)
- 250 kbps data rate
- Direct-sequence spread-spectrum (DSSS) with frequency-hopping between 15 channels within that band for security and interference reduction
- TDMA (Time-Division Multiple Access) bus arbitration, with 10-millisecond timeslots allocated for device transmission
- Variable transmit power, with 10 dBm (10 milliwatts) being default

*Wireless*HART uses 2.4 GHz (nominal) as its transmission frequency and low power levels (10 dBm nominal) because meeting these criteria allows *Wireless*HART devices to be unlicensed according to FCC (Federal Communications Commission) standards. If *Wireless*HART fell outside of these limits, the FCC would require end-users to obtain and maintain licenses for the use of these devices and licenses for maintenance personnel installing and maintaining the devices. Such requirements would make *Wireless*HART prohibitively expensive for all but the most challenging applications and thereby limit its marketability.

TDMA bus arbitration means that the Network Manager plans and schedules the transmission times of all field devices, giving each one its own dedicated time to “speak.” With these non-overlapping timeslots scheduled and broadcast to all the field devices, collisions are prevented while at the same time ensuring determinism (the guarantee that data packets *will* reach their destination within a certain specified time) barring any physical interruption of the data path.

The purpose of variable transmit power (as scheduled by the Network Manager) is to conserve battery life: an important priority for instruments whose main (or even sole) source of energy is a battery with a finite life. A secondary benefit of this power-limiting feature is that the interference potential of a *Wireless*HART network on other wireless devices sharing the same 2.4 GHz band is further minimized.

Data Link Layer

- Network ID number uniquely identifies each *Wireless*HART network, allowing multiple networks to overlap the same physical area
- Channel “blacklisting” – automatically avoids hopping to noisy channels

Network Layer

- “Mesh” networking – devices automatically establish links with any other nearby *WirelessHART* devices
- Signal repeating – devices may act as “repeaters” for other devices too far away from the master unit
- A *Network Manager* device determines communication routes between field devices, as well as timing schedules
- Four levels of data message priority (listed from highest to lowest):
 - Command:** network management messages
 - Process data:** PV values
 - Normal:** *all messages other than Command, Process, or Alarm*
 - Alarm:** messages reporting device alarms and events

The Network Manager in a *WirelessHART* network plays a role similar to the Link Active Scheduler (LAS) in a FOUNDATION Fieldbus network segment. The Network Manager assigns time-slots for individual devices to communicate, determines alternative communication routes (i.e. it designs and continually updates the mesh), and continually adjusts device transmit power in order to ensure optimal operation. This dynamic management of the wireless network is critically important in order to maintain low data latency times and high reliability in the face of changing environment variables such as objects coming into and out of the radio pathways (e.g. cranes, trucks, forklifts, man-lifts, scaffolding, and any other large metal structures which may temporarily alter the RF environment in an industrial setting.). Like FOUNDATION Fieldbus LAS devices, multiple (redundant) Network Managers are possible within a *WirelessHART* network with only one being active at any time.

Application Layer

- 128-bit encryption of data
- Backward-compatibility with wired-HART command structure and DDL (Device Description Language)

The backward compatibility of *WirelessHART* with wired-HART field instruments is an incredibly valuable feature of this standard, as it opens the door to wireless integration of legacy HART instruments. All that is needed to make a wired-HART instrument part of a functioning *WirelessHART* network is to attach the appropriate adapter, such as Emerson's *THUM*. Essentially, this step adds an antenna (and associated network interface electronics) on any legacy HART instrument, enabling it to communicate with native *WirelessHART* instruments and with the wireless gateway. This backward compatibility also improves integration of *WirelessHART* instruments, as they may communicate with legacy HART software application just as easily as wired-HART devices can. This means programs such as Emerson's AMS are able to interrogate *Wireless HART* instruments just as easily as they can wired-HART instruments, with no changes to the program code.

Other wireless networking protocols exist which are similar but not identical to *WirelessHART*. A few are listed here in contrast for better understanding.

***WirelessHART* versus Bluetooth**

Bluetooth is a popular wireless communication standard used in personal computing and other personal electronic devices such as cell phone headsets.

Like *WirelessHART*, Bluetooth supports channel-hopping and uses TDMA arbitration. However, Bluetooth uses a much simpler *star* network topology: up to seven Bluetooth slave devices may communicate with one Bluetooth master device. By contrast, *WirelessHART* allows for a greater number of field devices communicating with one Network Manager device, and the network topology is *mesh*, where any device may transmit data to any other device on the same network and have that other device "repeat" the data to the Network Manager.

***WirelessHART* versus ZigBee**

ZigBee is a mesh-networking wireless communication standard which has found application in building automation systems. It applies the IEEE 802.15.4-2006 standard for both Physical and Data Link layers, whereas *WirelessHART* employs its own unique Data Link layer including features such as channel "blacklisting" and time-slot synchronization to avoid collisions.

A major difference between ZigBee and *WirelessHART* is the methods of channel arbitration used: ZigBee uses CSMA/CA while *WirelessHART* uses TDMA. Time Division arbitration tends to be more time-efficient (and certainly more deterministic) when large numbers of devices are within range of each other.

WirelessHART versus Wi-Fi

Wi-Fi (IEEE 802.11) is a wireless communication standard that is extremely popular for personal computer Internet access. Unlike *WirelessHART*, Wi-Fi does not support channel-hopping for security and interference reduction. Wi-Fi, like ZigBee, also uses CSMA/CA channel arbitration, while *WirelessHART* uses TDMA channel arbitration to achieve determinism.

17.2.3 WirelessHART network gateway device


The *Network Gateway* is a critically important component in a *WirelessHART* system. It is the sole channel through which all field device data funnels to the host control system. Physically, a network gateway is nothing more than a box with an antenna on it, and connections within for electrical power and wired networks (e.g. Ethernet, EIA/TIA-485). Shown here is an Emerson model 1420¹⁶ “Smart Wireless Gateway”:



Electrically, these devices are quite complex. They are microprocessor-controlled, and often serve as the physical host for the Network Manager algorithm: orchestrating and tuning the wireless network communications.

¹⁶The model 1420 gateway has been superseded by the Smart Wireless Gateway, also manufactured by Emerson.

Since *WirelessHART* is a purely *digital* communication standard, all data points from the field devices are stored in the gateway in digital form, and must be accessed digitally. In the case of Emerson's Smart Wireless Gateway, the data may be accessed by any host system via Modbus query commands, communicated either serially (EIA/TIA-485, Modbus RTU format) or encapsulated in Ethernet packets (Modbus TCP). Screw terminal connections exist on the Emerson gateway for an EIA/TIA-485 (RS-485) cable to connect, as well as multiple RJ-45 Ethernet ports for connection to a hub or switch where other Ethernet-based computers and systems may connect as well:



Like so many other industrial Ethernet-ready devices, the Emerson Smart Wireless Gateway has a built-in web server, allowing password-protected access to configuration pages using nothing more than a personal computer with Ethernet connectivity and a web (Internet) browser program. Simply type the IP address of the gateway port into the browser's URL field, and the personal computer will connect to the gateway.

Individual device data points are custom-mapped by the user to specific Modbus registers inside the gateway's memory, as shown on this configuration page:

The screenshot displays the 'Modbus Register Map' configuration page for the Smart Wireless Gateway. The page is accessed via a web browser at the URL `http://169.254.5.3:8080/mb_regmap`. The interface features a navigation tree on the left with categories such as Diagnostics, Monitor, Explorer, Setup, Network, Internet protocol, Security, Time, System Backup, Page Options, Restart Apps, HART, Modbus, Communication, Mapping, Import/Export, OPC, and Trends. The main content area shows a table with the following data:

Register	Point Name	State	Invert
<input type="checkbox"/> 30001	TEMP #4.PV		<input type="checkbox"/>
<input type="checkbox"/> 30002	TEMP 2.PV		<input type="checkbox"/>


Below the table, there are navigation controls including '<< First', '<< Previous', a search input field, 'Search', 'Page 1 of 1', 'Next >>', and 'Last >>'. There are also buttons for 'New entry', 'Delete selected', 'Select All', 'None', 'Errors', and 'Submit'.

In this screenshot we see the primary variables¹⁷ (PV) of two Rosemount model 648 *WirelessHART* temperature transmitters mapped to Modbus registers 30001 and 30002. It should be noted that all *WirelessHART* field instruments are multi-variable devices, and as such are capable of reporting more than one variable to the gateway. If anyone were interested, it would have been possible in this example to assign battery voltage as a secondary variable (SV), tertiary variable (TV), or quaternary variable (QV) inside one or both temperature transmitters, then map those data points to their own Modbus registers in the gateway so that a host system could access and monitor battery voltage for the field instruments. Just as in wired-HART communication, multi-variable data communication from each transmitter is possible. This is not often done as a regular course of action with wired-HART instruments due to the very slow data rate of wired HART (1200 bps). However, with the much faster data rate of *WirelessHART* (250 kbps), the extra time required

¹⁷Device variables are addressed at the network gateway level by the device's HART tag (long tag, not short tag) and internal device variable name. Thus, the primary variable (PV) of temperature transmitter TEMP2 is specified as TEMP2.PV using a period symbol (.) as the delimiting character between the device name and the internal variable name.


for a field instrument to transmit three or four variables instead of just one variable is insignificant with respect to the needs of process measurement and control.

The next screenshot shows a portion of a simple PLC program written to query these two Modbus registers inside the Emerson gateway. The PLC in this example happens to be an Automation Direct “CLICK” model with a built-in EIA/TIA-485 data port, which connects directly to the gateway’s Modbus RTU network screw terminals.



Here, the “Receive” instruction in the PLC sends a Modbus function code 04 to read two analog input registers inside the slave device, that slave device being the Emerson Smart Wireless Gateway (Modbus address 10 on this particular EIA/TIA-485 network).

The result of this Modbus query is shown in the next screenshot, where the “Data View” window of the PLC is configured to display the two integer values obtained from the Modbus 04 command. These integer values (stored to registers DS1 and DS2 inside the PLC’s memory) happen to be 60 and 61, representing 60 degrees Fahrenheit and 61 degrees Fahrenheit, respectively. The two temperature transmitters happened to be measuring outdoor ambient temperature at the time this screenshot was taken:



Now that the temperature data resides in the PLC registers, the PLC may be programmed to take action on this data. For example, the PLC may be programmed to turn on cooling fans when the temperatures exceed pre-set limits.

Many modern HMI (Human-Machine Interface) display panels are also capable of serving as Modbus master devices, and may directly read from and write to the network gateway without the need of a PLC. For *WirelessHART* systems requiring no automatic control (i.e. monitoring and/or manual control functions only) interfacing an HMI panel to the gateway is a simple and practical solution.

17.2.4 WirelessHART device commissioning and configuration

WirelessHART field instruments look much like their wired counterparts, with the obvious addition of an antenna. A WirelessHART Rosemount model 648 temperature transmitter¹⁸ appears in this photograph:



Removing the large cover on this transmitter reveals the lithium battery:



A pair of metal terminals marked “Comm” on the transmitter where the battery plugs in provide a place to connect a standard HART communicator device, such as an Emerson model 475. Remember that *WirelessHART* instruments are fully HART-compliant devices, and may be configured identically to a wired-HART device using the same tools.

¹⁸This is an example of a first-generation Rosemount *WirelessHART* field instrument, back when the standard radio band was 900 MHz instead of 2.4 GHz. This explains why the antenna is longer than contemporary *WirelessHART* instruments.

Two parameters unique to *WirelessHART* devices, essential to specify in each field device (*WirelessHART* instrument) for establishing communication with the network gateway, are the *Network ID* and *Device Join Key*. These two parameters are contrasted in the following table:

Parameter	Format	Scope
Network ID	Integer between 0 and 36863	Shared by gateway and its field devices
Device Join Key	Four 4-byte fields (128 bits)	May be unique to each field device

The purpose of the Network ID is to simply associate each field device with one¹⁹ network gateway. Each *WirelessHART* gateway is programmed with one unique Network ID number, which is shared by all field devices communicating with that gateway. The purpose of the Device Join Key is altogether different: this is to provide *data security* by ensuring that only permitted devices can become a part of a particular gateway’s wireless mesh network. This essential difference explains why the Join Key is a much larger digital data field than the Network ID: the larger the “passcode” to join a network, the less likely any unauthorized agent will be able to randomly guess that passcode and establish a connection with that network.

An analogy to help understand the distinction between the Network ID and the Device Join Key is a street address versus a door key of a house, respectively. Each person living in a house must know where to find the house (thus the purpose for memorizing the street address), but access is granted only by possessing a key that unlocks the house door. In the simplest *WirelessHART* systems, all devices on a particular mesh network share the same Join Key, just as they (must) share the same Network ID. This is analogous to all residents of a house carrying identical keys to unlock the same door.

Although it is possible to configure a network gateway to have one “common” Join Key shared by all associated devices in that network, stronger security will be realized by assigning a unique Join Key to each device. In the latter case, the network gateway will maintain a list of all Join Keys and their associated devices, to ensure a device cannot become part of the wireless mesh network unless its programmed Join Key matches the one stored inside the gateway. Returning to our house analogy, this would be the equivalent of each resident having their own unique key to fit their own door on the house, with each door guarded by a security agent checking the name of the person trying to enter: in order to enter the house, your name would have to be on the resident list *and* you would have to be carrying the matching key for your name! For even stronger security, the gateway may be configured to generate random Join Keys (instead of the technician having to create their own 128-bit numbers), and may even be configurable to *rotate* the Device Join Keys on a periodic basis so that the Join Key for any particular device will not remain the same over time.

¹⁹Each gateway device can of course have backup gateways with the same Network ID, just waiting to take over if the primary gateway fails. The point of the Network ID is that it identifies a single *network* with only one active gateway.

Once a *WirelessHART* device has been powered, configured with the proper Network ID and Join Key parameters, and placed within range of a *WirelessHART* mesh network, it should be automatically detected by the Network Manager in time. Once detected, the device will appear in a list of network devices active in that *WirelessHART* network. Here are some tips to aid the commissioning process:

- Be sure to configure the device's *HART long tag* with the HART communicator prior to commissioning on the wireless network. This way the device will appear on the list of active devices with its proper tagname already configured, rather than as a cryptic MAC address. In the case of a *WirelessHART* adapter for a wired-HART device (e.g. an Emerson THUM connected to a legacy HART field instrument), you will need to place the instrument tagname in the wired HART device's "message" field. This tagname will become the leading portion of each variable name within the device: for example, the primary variable (PV) within a *WirelessHART* temperature transmitter with the tagname TT-35 will be addressed as TT-35.PV on the gateway's list of device variables once commissioned.
- Ensure a strong radio communication pathway between the *WirelessHART* field device and the gateway. This includes maintaining proper antenna orientation (either vertical up or down) and not too close to ground level, minimal obstructions between the device and the gateway, and not too far away from the gateway.
- Keep the field device powered down (i.e. its battery unplugged) until you have it in position and ready to commission. The default setting of *WirelessHART* devices is to request to join the network when powered up, so the act of plugging in the battery to a field device is the initiating event for commissioning on the wireless network.
- Turn the "Active Advertising" mode of the gateway *on*. This prompts the entire network (including all field devices) to actively search for uncommissioned devices and thereby expedites the joining process.
- Turn the "Rotate Network Key" feature of the gateway *off*. You do not want the Join Key randomly changing on you as you try to commission new devices!
- When commissioning several field devices in one area, begin with the device closest to the gateway antenna and proceed to the farthest device. This will exploit the ability of all *WirelessHART* field devices to act as *repeaters* for devices located far from the gateway.
- Refresh your web browser screen when checking device statuses on the gateway, because not all web browser software responds reliably to new data "pushed" from the gateway's HTTP server.
- If a field device is slow to join the wireless network, you may connect a HART communicator to the device's "COMM" terminals and monitor its join status directly. This will reveal any problems with the join process.
- Initially set the Update Rate to the fastest (i.e. shortest update time) possible in the field device. This does not affect the device's join time, but once joined it decreases the amount of time you must wait to monitor variables within the device. You may always re-set the update time to a slower value after commissioning, through the gateway.

- Be patient. Even when you have done everything correctly, the commissioning may take several minutes. Have other work ready to do (e.g. update instrument documentation, Modbus configuration in the gateway) while you are waiting for devices to join the wireless network. Having all field device tagnames pre-configured helps, because it allows you to populate the Modbus mapping table with proper variable names before the device has joined the wireless network.

Network gateways provide some basic statistical information on connected devices, which may be useful for diagnosing problems. Some of these statistics may be seen in the following computer screenshot taken from an Emerson model 1420 gateway:

The screenshot displays the 'Smart Wireless Gateway' interface. The main content area is titled 'Network Device Status' and contains a table with the following data:

HART Tag	State	Parents	Neighbors	Late updates	Path stability	RSSI	Joins	Join Time	Battery
TEMP #4	●	Unknown	2	2	100.0 %	-55 db	0	03/02/11 13:46:16	7.128 V
TEMP 2	●	Unknown	2	0	100.0 %	-49 db	0	03/02/11 13:43:10	7.161 V

“RSSI” refers to *Received Signal Strength Indication*, and is a measure of each device’s received RF signal strength, in units of dBm. Problems related to antennas, path loss, fade loss, and interference will result in decreased RSSI for that device. This same page shows the battery voltage for each field device. The “neighbors” parameter tells us how many *WirelessHART* devices are within range of each field device (the network gateway is counted among them). Thus, in this simple *WirelessHART* network consisting of two field devices and one gateway, all within range of each other, each field device reports having two neighbors.

17.3 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Conservation of energy:** energy cannot be created or destroyed, only converted between different forms. Relevant to antenna gain: high-gain antennas don't really amplify signals, merely focus them better, in the same way that a magnifying glass doesn't increase the amount of light but rather just focuses that light onto a smaller area.
- **Common logarithms:** used to express measurements spanning a tremendous range. Relevant to radio power and signal noise calculations.
- **Decibels:** used to express the ratio of one power to another in logarithmic form, such that the sum of component dB values is equivalent to the product of those same component gain/loss ratios. Decibels may also be used to express a power quantity relative to some reference power value such as 1 milliwatt (dBm) or 1 watt (dBW). Decibels are an example of a mathematical *transform function*, whereby one type of mathematical problem (multiplication/division) is transformed into an easier type of problem (addition/subtraction).
- **Analog vs. digital signals:** analog signals have infinite resolution but are susceptible to corruption by noise. Digital signals have limited resolution but are tolerant of any noise measuring less than the difference in thresholds between the high and low states.
- **Transmission lines:** short-duration (pulsed) electrical signals travel along a cable at nearly the speed of light, reflecting off the end of that cable if not properly terminated. Relevant to signal cables carrying high-frequency signals.
- **Electromagnetic induction:** occurs only when magnetic fields are perpendicular to the conductor. Relevant to optimal coupling between antennas, where antenna conductors should be parallel to each other so that their magnetic field polarizations will be perpendicular. Also relevant to minimizing coupling between antennas and interfering objects, where we try to distance the antenna from any parallel metal objects to reduce coupling and therefore reduce signal reflections and power loss.
- **Electrostatic coupling:** occurs when electric fields bridge between conductors, and cannot occur "behind" a grounded conductor. Relevant to minimizing signal degradation between antennas and interfering objects (especially grounded conductors), where we try to distance the antenna from any grounded metal objects to reduce coupling and therefore reduce signal reflections and power loss.
- **Inverse square law:** the strength of a field radiating away from a point-source diminishes proportionately to the square of the distance from the source. Relevant to path loss in radio power budget calculations, where path loss (L_p) is calculated based on the assumption of the radiator being a point-source isotropic (perfectly omnidirectional) antenna.

References

“Antenna Deployment Technical Brief”, Hewlett-Packard Development Company, 2006.

Code of Federal Regulations, Title 47, Volume 1, Chapter 1, Part 15, Subpart C, “Intentional Radiators”, document 47CFR15.247, pages 733-735, revised 2001.

Field Antenna Handbook, U.S. Marine Corps document MCRP 6-22D, 1999.

“FGR2 900 MHz Industrial Radio”, FreeWave Technologies, Inc., Boulder, CO, 2009.

“IEC 62591 WirelessHART System Engineering Guide”, Emerson Process Management, 2010.

IEEE Std 802.15.4-2006, The Institute of Electrical and Electronic Engineers, Inc., New York, NY, September 2006.

Lee, Jin-Shyan; Su, Yu-Wei; Shen, Chung-Chou, “A Comparative Study of Wireless Protocols: Bluetooth, UWB, ZigBee, and Wi-Fi”, Industrial Technology Research Institute, Hsinchu, Taiwan, November 2007.

Nixon, Mark, “A Comparison of WirelessHART and ISA100.11a”, Emerson Process Management, Round Rock, TX, 2012.

Song, Jianping, et. al., “WirelessHART: Applying Wireless Technology in Real-Time Industrial Process Control”, *IEEE Real-Time and Embedded Technology and Applications Symposium*, IEEE Computer Society, pages 377-386, Los Alamitos, CA, 2008.

The ARRL Antenna Book, Eleventh Edition, The American Radio Relay League, Inc., Newington, CT, 1968.

The ARRL Handbook For Radio Amateurs, 2001 Edition, ARRL – the national association for Amateur Radio, Newington, CT, 2001.

“WirelessHART Technical Data Sheet”, document HCF_LIT-89, HART Communication Foundation, Austin, TX, 2007.

Young, Michael F., “Planning a Microwave Radio Link”, YDI Wireless, Falls Church, VA, 2002.

Zyren, Jim and Petrick, Al, “Tutorial on Basic Link Budget Analysis”, Intersil Application Note AN9804.1, 1998.

Chapter 18

Instrument calibration

Every instrument has at least one *input* and one *output*. For a pressure sensor, the input would be some fluid pressure and the output would (most likely) be an electronic signal. For a loop indicator, the input would be a 4-20 mA current signal and the output would be a human-readable display. For a variable-speed motor drive, the input would be an electronic signal and the output would be electric power to the motor.

Calibration and *ranging* are two tasks associated with establishing an accurate correspondence between any instrument's input signal and its output signal. Simply defined, *calibration* assures the instrument accurately senses the real-world variable it is supposed to measure or control. Simply defined, *ranging* establishes the desired relationship between an instrument's input and its output.

18.1 Calibration versus re-ranging


To *calibrate* an instrument means to check and adjust (if necessary) its response so the output accurately corresponds to its input throughout a specified range. In order to do this, one must expose the instrument to an actual input stimulus of precisely known quantity. For a pressure gauge, indicator, or transmitter, this would mean subjecting the pressure instrument to known fluid pressures and comparing the instrument response against those known pressure quantities. One cannot perform a true calibration without comparing an instrument's response to known, physical stimuli.

To *range* an instrument means to set the lower and upper range values so it responds with the desired sensitivity to changes in input. For example, a pressure transmitter set to a range of 0 to 200 PSI (0 PSI = 4 mA output ; 200 PSI = 20 mA output) could be re-ranged to respond on a scale of 0 to 150 PSI (0 PSI = 4 mA ; 150 PSI = 20 mA).

In analog instruments, re-ranging could (usually) only be accomplished by re-calibration, since the same adjustments were used to achieve both purposes. In digital instruments, calibration and ranging are typically separate adjustments (i.e. it is possible to re-range a digital transmitter without having to perform a complete recalibration), so it is important to understand the difference.


18.2 Zero and span adjustments (analog instruments)

The purpose of *calibration* is to ensure the input and output of an instrument reliably correspond to one another throughout the entire range of operation. We may express this expectation in the form of a graph, showing how the input and output of an instrument should relate. For the vast majority of industrial instruments this graph will be linear:



This graph shows how any given percentage of input should correspond to the same percentage of output, all the way from 0% to 100%.

Things become more complicated when the input and output axes are represented by units of measurement other than “percent.” Take for instance a pressure *transmitter*, a device designed to sense a fluid pressure and output an electronic signal corresponding to that pressure. Here is a graph for a pressure transmitter with an input range of 0 to 100 pounds per square inch (PSI) and an electronic output signal range of 4 to 20 milliamps (mA) electric current:



Although the graph is still linear, zero pressure does not equate to zero current. This is called a *live zero*, because the 0% point of measurement (0 PSI fluid pressure) corresponds to a non-zero (“live”) electronic signal. 0 PSI pressure may be the LRV (Lower Range Value) of the transmitter’s input, but the LRV of the transmitter’s output is 4 mA, not 0 mA.

Any linear, mathematical function may be expressed in “slope-intercept” equation form:

$$y = mx + b$$

Where,

y = Vertical position on graph

x = Horizontal position on graph

m = Slope of line

b = Point of intersection between the line and the vertical (y) axis

This instrument’s calibration is no different. If we let x represent the input pressure in units of PSI and y represent the output current in units of milliamps, we may write an equation for this instrument as follows:

$$y = 0.16x + 4$$

On the actual instrument (the pressure transmitter), there are two adjustments which let us match the instrument’s behavior to the ideal equation. One adjustment is called the *zero* while the

other is called the *span*. These two adjustments correspond exactly to the b and m terms of the linear function, respectively: the “zero” adjustment shifts the instrument’s function vertically on the graph (b), while the “span” adjustment changes the slope of the function on the graph (m). By adjusting both zero and span, we may set the instrument for any range of measurement within the manufacturer’s limits.

The relation of the slope-intercept line equation to an instrument’s zero and span adjustments reveals something about how those adjustments are actually achieved in any instrument. A “zero” adjustment is always achieved by *adding* or *subtracting* some quantity, just like the y -intercept term b adds or subtracts to the product mx . A “span” adjustment is always achieved by *multiplying* or *dividing* some quantity, just like the slope m forms a product with our input variable x .

Zero adjustments typically take one or more of the following forms in an instrument:

- Bias force (spring or mass force applied to a mechanism)
- Mechanical offset (adding or subtracting a certain amount of motion)
- Bias voltage (adding or subtracting a certain amount of potential)

Span adjustments typically take one of these forms:

- Fulcrum position for a lever (changing the force or motion multiplication)
- Amplifier gain (multiplying or dividing a voltage signal)
- Spring rate (changing the force per unit distance of stretch)

It should be noted that for most analog instruments, zero and span adjustments are *interactive*. That is, adjusting one has an effect on the other. Specifically, changes made to the span adjustment almost always alter the instrument’s zero point¹. An instrument with interactive zero and span adjustments requires much more effort to accurately calibrate, as one must switch back and forth between the lower- and upper-range points repeatedly to adjust for accuracy.

18.3 Calibration errors and testing

The efficient identification and correction of instrument calibration errors is an important function for instrument technicians. For some technicians – particularly those working in industries where calibration accuracy is mandated by law – the task of routine calibration consumes most of their working time. For other technicians calibration may be an occasional task, but nevertheless these technicians must be able to quickly diagnose calibration errors when they cause problems in instrumented systems. This section describes common instrument calibration errors and the procedures by which those errors may be detected and corrected.

¹However, it is actually quite rare to find an instrument where a change to the zero adjustment affects the instrument’s span.

18.3.1 Typical calibration errors

Recall that the slope-intercept form of a linear equation describes the response of *any* linear instrument:

$$y = mx + b$$

Where,


y = Output

m = Span adjustment

x = Input


b = Zero adjustment

A *zero shift* calibration error shifts the function vertically on the graph, which is equivalent to altering the value of b in the slope-intercept equation. This error affects *all* calibration points equally, creating the same percentage of error across the entire range. Using the same example of a pressure transmitter with 0 to 100 PSI input range and 4 to 20 mA output range:




If a transmitter suffers from a zero calibration error, that error may be corrected by carefully moving the “zero” adjustment until the response is ideal, essentially altering the value of b in the linear equation.

A *span shift* calibration error shifts the slope of the function, which is equivalent to altering the value of m in the slope-intercept equation. This error's effect is unequal at different points throughout the range:




If a transmitter suffers from a span calibration error, that error may be corrected by carefully moving the “span” adjustment until the response is ideal, essentially altering the value of m in the linear equation.

A *linearity* calibration error causes the instrument's response function to no longer be a straight line. This type of error does not directly relate to a shift in either zero (b) or span (m) because the slope-intercept equation only describes straight lines:



Some instruments provide means to adjust the linearity of their response, in which case this adjustment needs to be *carefully* altered. The behavior of a linearity adjustment is unique to each model of instrument, and so you must consult the manufacturer's documentation for details on how and why the linearity adjustment works. If an instrument does not provide a linearity adjustment, the best you can do for this type of problem is "split the error" between high and low extremes, so the maximum absolute error at any point in the range is minimized.



A *hysteresis* calibration error occurs when the instrument responds differently to an increasing input compared to a decreasing input. The only way to detect this type of error is to do an *up-down* calibration test, checking for instrument response at the same calibration points going down as going up:



Hysteresis errors are almost always caused by mechanical friction on some moving element (and/or a loose coupling between mechanical elements) such as bourdon tubes, bellows, diaphragms, pivots, levers, or gear sets. Friction always acts in a direction opposite to that of relative motion, which is why the output of an instrument with hysteresis problems always *lags* behind the changing input, causing the instrument to register falsely low on a rising stimulus and falsely high on a falling stimulus. Flexible metal strips called *flexures* – which are designed to serve as frictionless pivot points in mechanical instruments – may also cause hysteresis errors if cracked or bent. Thus, hysteresis errors cannot be remedied by simply making calibration adjustments to the instrument – one must usually replace defective components or correct coupling problems within the instrument mechanism.

In practice, most calibration errors are some combination of zero, span, linearity, and hysteresis problems. An important point to remember is that with rare exceptions, zero errors always accompany other types of errors. In other words, it is extremely rare to find an instrument with a span, linearity, or hysteresis error that does not *also* exhibit a zero error. For this reason, technicians often perform a single-point calibration test of an instrument as a qualitative indication of its calibration health. If the instrument performs within specification at that one point, its calibration over the entire range is *probably* good. Conversely, if the instrument fails to meet specification at that one point, it definitely needs to be recalibrated.

A very common single-point test for instrument technicians to perform on differential pressure (“DP”) instruments is to close both block valves on the three-valve manifold assembly and then open the equalizing valve, to produce a known condition of 0 PSI differential pressure:

Normal operation**"Block and Equalize" test**

Most DP instrument ranges encompass 0 PSI, making this a very simple single-point check. If the technician “blocks and equalizes” a DP instrument and it properly reads zero, its calibration is *probably* good across the entire range. If the DP instrument fails to read zero during this test, it definitely needs to be recalibrated.

18.3.2 As-found and as-left documentation

An important principle in calibration practice is to document every instrument's calibration as it was found *and* as it was left after adjustments were made. The purpose for documenting both conditions is to make data available for calculating instrument *drift* over time. If only one of these conditions is documented during each calibration event, it will be difficult to determine how well an instrument is holding its calibration over long periods of time. Excessive drift is often an indicator of impending failure, which is vital for any program of predictive maintenance or quality control.

Typically, the format for documenting both As-Found and As-Left data is a simple table showing the points of calibration, the ideal instrument responses, the actual instrument responses, and the calculated error at each point. The following table is an example for a pressure transmitter with a range of 0 to 200 PSI over a five-point scale:

Percent of range	Input pressure	Output current (ideal)	Output current (measured)	Error (percent of span)
0%	0 PSI	4.00 mA		
25%	50 PSI	8.00 mA		
50%	100 PSI	12.00 mA		
75%	150 PSI	16.00 mA		
100%	200 PSI	20.00 mA		

The following photograph shows a single-point “As-Found” calibration report on a temperature indicating controller, showing the temperature of the calibration standard (−78.112 degrees Celsius), the display of the instrument under test (IUT, −79 degrees Celsius), and the error between the two (−0.888 degrees Celsius):

As found Readings		
Standard	IUT	Error
°C	°C	°C
-78.112	-79.0	-0.888

Note that the mathematical sign of the error is important. An instrument that registers −79 degrees when it should register −78.112 degrees exhibits a *negative error*, since its response is lower (i.e. more negative) than it should be. Expressed mathematically: Error = IUT − Standard. When the error must be expressed in percentage of span, the formula becomes:

$$\text{Error} = \frac{\text{IUT} - \text{Standard}}{\text{Span}} \times 100\%$$

18.3.3 Up-tests and Down-tests

It is not uncommon for calibration tables to show multiple calibration points going *up* as well as going *down*, for the purpose of documenting hysteresis and deadband errors. Note the following example, showing a transmitter with a maximum hysteresis of 0.313 % (the offending data points are shown in bold-faced type):

Percent of range	Input pressure	Output current (ideal)	Output current (measured)	Error (percent of span)
0%	0 PSI	4.00 mA	3.99 mA	-0.0625 %
25% ↑	50 PSI	8.00 mA	7.98 mA	-0.125 %
50% ↑	100 PSI	12.00 mA	11.99 mA	-0.0625 %
75% ↑	150 PSI	16.00 mA	15.99 mA	-0.0625 %
100% ↑	200 PSI	20.00 mA	20.00 mA	0 %
75% ↓	150 PSI	16.00 mA	16.01 mA	+0.0625 %
50% ↓	100 PSI	12.00 mA	12.02 mA	+0.125 %
25% ↓	50 PSI	8.00 mA	8.03 mA	+0.188 %
0% ↓	0 PSI	4.00 mA	4.01 mA	+0.0625 %

Note again how error is expressed as either a positive or a negative quantity depending on whether the instrument's measured response is above or below what it should be under each condition. The values of error appearing in this calibration table, expressed in percent of span, are all calculated by the following formula:


$$\text{Error} = \left(\frac{I_{\text{measured}} - I_{\text{ideal}}}{16 \text{ mA}} \right) (100\%)$$

In the course of performing such a directional calibration test, it is important not to overshoot any of the test points. If you do happen to overshoot a test point in setting up one of the input conditions for the instrument, simply “back up” the test stimulus and re-approach the test point from the same direction as before. Unless each test point's value is approached from the proper direction, the data cannot be used to determine hysteresis/deadband error.

18.3.4 Automated calibration

Maintaining the calibration of instruments at a large industrial facility is a daunting task. Aside from the actual labor of checking and adjusting calibration, records must be kept not only of instrument performance but also of test conditions and criteria (e.g. calibration tolerance, time interval between calibrations, number of points to check, specific procedures, etc.). Any practical method to minimize human error in this process is welcome. For this reason, automated and semi-automated calibration tools have been developed to help manage the data associated with calibration, and to make the instrument technician's job more manageable.

An example of a fully automated calibration system is a process chemical analyzer where a set of solenoid valves direct chemical samples of known composition to the analyzer at programmed time intervals, a computer inside the analyzer recording the analyzer's error (compared to the known standard) and auto-adjusting the analyzer in order to correct for whatever errors are detected. In the following illustration we see a schematic of a gas analyzer with two compressed-gas cylinders holding gases of 0% and 100% concentration of the compound(s) of interest, called "zero gas" and "span gas", connected through solenoid valves so that the chemical analyzer may be automatically tested against these standards:



The only time a human technician need attend to the analyzer is when parameters not monitored by the auto-calibration system must be checked, and when the auto-calibration system detects an

error too large to self-correct (thus indicating a fault).

An example of a semi-automated calibration system is an instrument such as Fluke's series of *Documenting Process Calibrators* (DPC). These devices function as standards for electrical measurements such as voltage, current, and resistance, with built-in database capability for storing calibration records and test conditions:



A technician using a documenting calibrator such as this is able to log As-Found and As-Left data in the device's memory and download the calibration results to a computer database at some later time. The calibrator may also be programmed with test conditions for each specific instrument on the technician's work schedule, eliminating the need for that technician to look up each instrument's test conditions in a book, and thereby reducing the potential for human error.

An example of database software used to schedule routine instrument calibrations and archive the results is Fluke's *DPCTrack2*, a point-and-click user interface serving as a front-end to an SQL database where the instrument data is maintained in digital format on the computer's hard drive:




Calibration management software allows managers to define calibration schedules, tolerances, and even technician work assignments, the software allowing for downloading of this setup information into a hand-held calibrator unit, as well as uploading and archival of calibration results following the procedure.

In some industries, this degree of rigor in calibration record-keeping is merely helpful; in other industries it is vital for business. Examples of the latter include pharmaceutical manufacturing, where regulatory agencies (such as the Food and Drug Administration in the United States) enforces rigorous standards for manufacturing quality including requirements for frequent testing and data archival of process instrument accuracy. Record-keeping in such industries is not limited to As-Found and As-Left calibration results, either; each and every action taken on that instrument by a human being must be recorded and archived so that a complete audit of causes may be conducted should there ever be an incident of product mis-manufacture.

18.4 Damping adjustments

The vast majority of modern process transmitters (both analog and digital) come equipped with a feature known as *damping*. This feature is essentially a low-pass filter function placed in-line with the signal, reducing the amount of process “noise” reported by the transmitter.


Imagine a pressure transmitter sensing water pressure at the outlet of a large pump. The flow of water exiting a pump tends to be extremely turbulent, and any pressure-sensing device connected to the immediate discharge port of a pump will interpret this turbulence as fluctuations in pressure. This means the pressure signal output by the transmitter will fluctuate as well, causing any indicator or control system connected to that transmitter to register a “noisy” water pressure:



Such “noise” wreaks havoc with most forms of feedback control, since the control system will interpret these rapid fluctuations as real pressure changes requiring corrective action. Although it is possible to configure some control systems to ignore such noise, the best solution is to correct the problem at the source either by relocating the pressure transmitter’s impulse line tap to a place where it will not be exposed to so much turbulence, or somehow prevent that sensed turbulence from being represented in the transmitter’s signal.


Since this noise is of a much greater frequency than the normal cycles of pressure in a process system, it is relatively easy to reduce the amount of noise in the transmitter signal simply by filtering that electronic signal using a low-pass filter circuit.

The simplest low-pass filter circuit is nothing more than a resistor and capacitor:



Low-frequency voltage signals applied to this circuit emerge at the output terminal relatively unattenuated, because the reactance of the capacitor is quite large at low frequencies. High-frequency signals applied to the same circuit become attenuated by the capacitor, which tends to “short” those signals to ground with its low reactance to high frequencies. The performance of such a filter circuit is primarily characterized by its *cutoff frequency*, mathematically defined as $f = \frac{1}{2\pi RC}$. The cutoff frequency is the point at which only 70.7% of the input signal appears at the output (i.e. a -3 dB attenuation in voltage).


If successfully applied to a process transmitter, such low-pass filtering has the effect of “quieting” an otherwise noisy signal so only the real process pressure changes are seen, while the effect of turbulence (or whatever else was causing the noise) becomes minimal. In the world of process control, the intentional low-pass filtering of process measurement signals is often referred to as *damping* because its effect is to “damp” (attenuate) the effects of process noise:



In order for damping to be a useful tool for the technician in mitigating measurement noise, it must be adjustable. In the case of the RC filter circuit, the degree of damping (cutoff frequency) may be adjusted by changing the value of either R or C , with R being the easier component to adjust.

In digital transmitters where the damping is performed by a digital algorithm², damping may be adjusted by setting a numerical value in the transmitter's configuration parameters. In pneumatic transmitters, damping could be implemented by installing viscous elements to the mechanism, or more simply by adding volume to the signal line (e.g. excess tubing length, larger tubing diameter, or even "capacity tanks" connected to the tube for increased volume).

The key question for the technician then becomes, "how much damping should I use?" Insufficient damping will allow too much noise to reach the control system (causing "noisy" trends, indications, and erratic control), while excessive damping will cause the transmitter to understate the significance of sudden (real) process changes. In my experience there is a bad tendency for instrument technicians to apply excessive damping in transmitters. A transmitter with too much damping (i.e. cutoff frequency set too low, or time constant value set too high) causes the trend graph to be very smooth, which at first appears to be a good thing. After all, the whole point of a control system is to hold the process variable tightly to setpoint, so the appearance of a "flat line" process variable trend is enticing indeed. However, the problem with excessive damping is that the transmitter gives a sluggish response to any sudden changes in the real process variable. A dual-trend graph of a pressure transmitter experiencing a sudden increase in process pressure shows this principle, where the undamped transmitter signal is shown in the upper portion and the over-damped signal in the lower portion (please note the vertical offset between these two trends is shown only for your convenience in comparing the two trend shapes):



Excessive damping causes the transmitter to "lie" to the control system by reporting a process variable that changes much slower than it actually does. The degree to which this "lie" adversely affects the control system (and/or the human operator's judgment in manually responding to the change in pressure) depends greatly on the nature of the control system and its importance to the overall plant operation.

²Various digital damping algorithms exist, but it may take as simple a form as successive averaging of buffered signal values coming out of a first-in-first-out ("FIFO") shift register.


One way damping may cause control problems is in systems where the loop controller is aggressively tuned. In such systems, even relatively small amounts of damping may cause the actual process variable to overshoot setpoint because the controller “thinks” the process variable is responding too slowly and takes action to speed its response. A common example of this is liquid flow control, where the process variable signal is typically “noisy” and the control action is typically aggressive. A technician may introduce damping to the transmitter with good intent, but unexpectedly causes the control system to wildly overshoot setpoint (or even oscillate) because the controller is trying to get a “sluggish” process variable to respond quicker than the transmitter filtering will allow the signal to change. In reality, the process variable (fluid flow rate) is not sluggish at all, but only appears that way because the transmitter is damped. What is worse, this instability will *not* appear on a trend of the process variable because the control system never sees the real process variable, but only the “lie” reported by the over-damped transmitter. If any rule may be given as to how much damping to use in any transmitter, it is this: use as *little* as necessary to achieve good control.

When calibrating a transmitter in a shop environment, the damping adjustment should be set to its absolute minimum, so the results of applying stimuli to the transmitter are immediately seen by the technician. Any amount of damping in a transmitter being calibrated serves only to slow down the calibration procedure without benefit.

18.5 LRV and URV settings, digital trim (digital transmitters)


The advent of “smart” field instruments containing microprocessors has been a great advance for industrial instrumentation. These devices have built-in diagnostic ability, greater accuracy (due to digital compensation of sensor nonlinearities), and the ability to communicate digitally³ with host devices for reporting of various parameters.

A simplified block diagram of a “smart” pressure transmitter looks something like this:



³Most popularly, using the HART digital-over-analog hybrid communication standard.

It is important to note all the adjustments within this device, and how this compares to the relative simplicity of an all-analog pressure transmitter:




Note how the only calibration adjustments available in the analog transmitter are the “zero” and “span” settings. This is clearly not the case with smart transmitters. Not only can we set lower- and upper-range values (LRV and URV) in a smart transmitter, but it is also possible to calibrate the analog-to-digital and digital-to-analog converter circuits independently of each other. What this means for the calibration technician is that a full calibration procedure on a smart transmitter potentially requires more work and a greater number of adjustments than an all-analog transmitter⁴.

A common mistake made among students and experienced technicians alike is to confuse the range settings (LRV and URV) for actual calibration adjustments. Just because you digitally set the LRV of a pressure transmitter to 0.00 PSI and the URV to 100.00 PSI does not necessarily mean it will register accurately at points within that range! The following example will illustrate this fallacy.

Suppose we have a smart pressure transmitter ranged for 0 to 100 PSI with an analog output range of 4 to 20 mA, but this transmitter’s pressure sensor is fatigued from years of use such that an

⁴Although those adjustments made on a digital transmitter tend to be easier to perform than repeated zero-and-span adjustments on analog transmitters due to the inevitable “interaction” between analog zero and span adjustments requiring repeated checking and re-adjustment during the calibration period.

actual applied pressure of 100 PSI generates a signal that the analog-to-digital converter interprets as only 96 PSI⁵. Assuming everything else in the transmitter is in perfect condition, with perfect calibration, the output signal will still be in error:

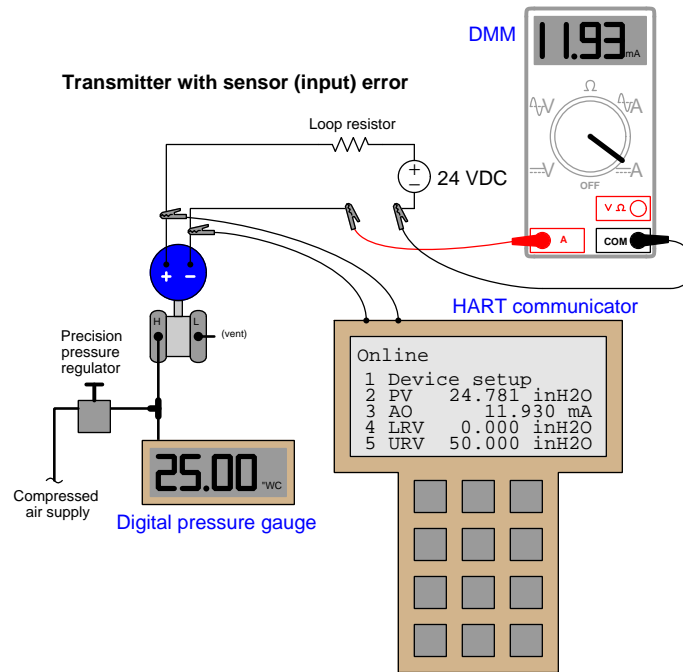


As the saying goes, “a chain is only as strong as its weakest link.” Here we see how the calibration of the most sophisticated pressure transmitter may be corrupted despite perfect calibration of both analog/digital converter circuits, and perfect range settings in the microprocessor. The microprocessor “thinks” the applied pressure is only 96 PSI, and it responds accordingly with a 19.36 mA output signal. *The only way anyone would ever know this transmitter was inaccurate at 100 PSI is to actually apply a known value of 100 PSI fluid pressure to the sensor and note the incorrect response.* The lesson here should be clear: digitally setting a smart instrument’s LRV and URV points does *not* constitute a legitimate calibration of the instrument.

For this reason, smart instruments always provide a means to calibrate both the ADC and DAC circuits, to ensure the microprocessor “sees” the correct representation of the applied stimulus and to ensure the microprocessor’s output signal gets accurately converted into a DC current, respectively. This calibration function is called *digital trim*.

⁵A 4% calibration error caused by sensor aging is enormous for any modern digital transmitter, and should be understood as an exaggeration presented only for the sake of illustrating how sensor error affects overall calibration in a smart transmitter. A more realistic amount of sensor error due to aging would be expressed in small fractions of a percent.

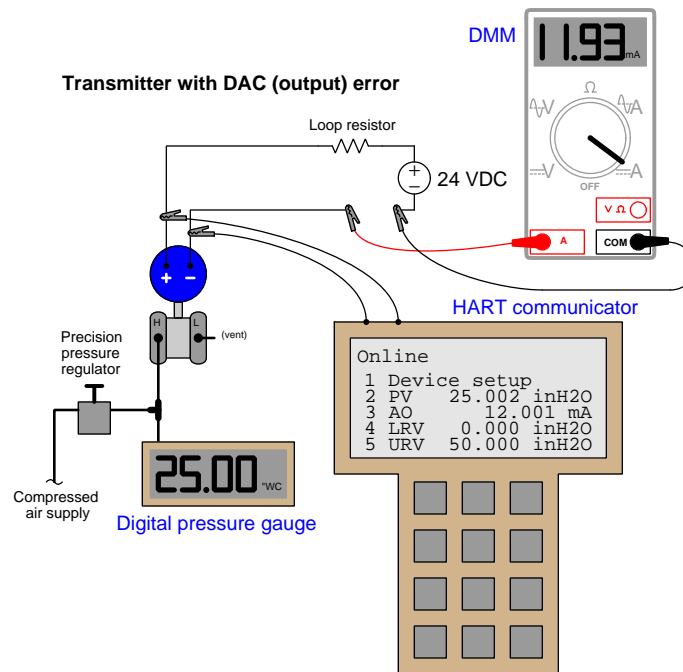
A convenient way to test a digital transmitter's analog/digital converters is to monitor the microprocessor's process variable (PV) and analog output (AO) registers while comparing the real input and output values against trusted calibration standards. A HART communicator⁶ provides this "internal view" of the registers so we may see what the microprocessor "sees." The following example shows a differential pressure transmitter with a sensor (analog-to-digital) calibration error:



Here, the calibration standard for pressure input to the transmitter is a digital pressure gauge, registering 25.00 inches of water column. The digital multimeter (DMM) is our calibration standard for the current output, and it registers 11.93 milliamps. Since we would expect an output of 12.00 milliamps at this pressure (given the transmitter's range values of 0 to 50 inches W.C.), we immediately know from the pressure gauge and multimeter readings that some sort of calibration error exists in this transmitter. Comparing the HART communicator's displays of PV and AO against our calibration standards reveals more information about the nature of this error: we see that the AO value (11.930 mA) agrees with the multimeter while the PV value (24.781 "W.C.) does *not* agree with the digital pressure gauge. This tells us the calibration error lies within the sensor (input) of the transmitter and not with the DAC (output). Thus, the correct calibration procedure to perform on this errant transmitter is a *sensor trim*.

⁶HART is a hybrid analog/digital communication protocol used by a great many field instruments, allowing maintenance personnel to access and edit digital parameters inside the instrument using a computer-based interface. Hand-held HART communicators exist for this purpose, as does HART software designed to run on a personal computer. HART modems also exist to connect personal computers to HART-compatible field instruments.

In this next example, we see what an output (DAC) error would look like with another differential pressure transmitter subjected to the same test:



Once again, the calibration standard for pressure input to the transmitter is a digital pressure gauge, registering 25.00 inches of water column. A digital multimeter (DMM) still serves as our calibration standard for the current output, and it registers 11.93 milliamperes. Since we expect 12.00 milliamperes output at this pressure (given the transmitter's range values of 0 to 50 inches W.C.), we immediately know from the pressure gauge and multimeter readings that some sort of calibration error exists in this transmitter (just as before). Comparing the HART communicator's displays of PV and AO against our calibration standards reveals more information about the nature of this error: we see that the PV value (25.002 inches W.C.) agrees with the digital pressure gauge while the AO value (12.001 mA) does *not* agree with the digital multimeter. This tells us the calibration error lies within the digital-to-analog converter (DAC) of the transmitter and not with the sensor (input). Thus, the correct calibration procedure to perform on this errant transmitter is an *output trim*.

Note how in both scenarios it was absolutely necessary to interrogate the transmitter's microprocessor registers with a HART communicator to determine where the error was located. Merely comparing the pressure and current standards' indications was not enough to tell us any more than the fact we had some sort of calibration error inside the transmitter. Not until we viewed the microprocessor's own values of PV and AO could we determine whether the calibration error was related to the ADC (input), the DAC (output), or perhaps even both.

Sadly, I have witnessed technicians attempt to use the LRV and URV settings in a manner not unlike the zero and span adjustments on an analog transmitter to correct errors such as these. While it may be possible to get an out-of-calibration transmitter to yield correct output current signal values over its calibrated range of input values by skewing the LRV and URV settings, it defeats the purpose of having separate “trim” and “range” settings inside the transmitter. Also, it causes confusion if ever the control system connected to the transmitter interrogates process variable values digitally rather than interpreting it via the 4-20 mA loop current signal. Finally, “calibrating” a transmitter by programming it with skewed LRV/URV settings corrupts the accuracy of any intentionally nonlinear functions such as square-root characterization (used for flow measurement applications) or strapping tables (used for liquid level measurement applications in vessels where the cross-sectional area varies with liquid height).

Once digital trims have been performed on both input and output converters, of course, the technician is free to re-range the microprocessor as many times as desired without re-calibration. This capability is particularly useful when re-ranging is desired for special conditions, such as process start-up and shut-down when certain process variables drift into uncommon regions. An instrument technician may use a hand-held HART communicator device to re-set the LRV and URV range values to whatever new values are desired by operations staff without having to re-check calibration by applying known physical stimuli to the instrument. So long as the ADC and DAC trims are both correct, the overall accuracy of the instrument will still be good with the new range. With analog instruments, the only way to switch to a different measurement range was to change the zero and span adjustments, which *necessitated* the re-application of physical stimuli to the device (a full re-calibration). Here and here alone we see where calibration is not necessary for a smart instrument. If overall measurement accuracy must be verified, however, there is no substitute for an actual physical calibration, and this entails both ADC and DAC “trim” procedures for a smart instrument.

Completely digital (“Fieldbus”) transmitters are similar to “smart” analog-output transmitters with respect to distinct trim and range adjustments. For an explanation of calibration and ranging on FOUNDATION Fieldbus transmitters, refer to section [16.5.3](#) beginning on page [1192](#).

18.6 An analogy for calibration versus ranging

The concepts of *calibration* (trimming) and *ranging* are often difficult for new students of instrumentation to immediately grasp. A simple analogy useful for understanding these topics is that of setting a digital alarm clock.

Suppose you purchase a digital alarm clock to wake you up at 7:00 AM in the morning so that you can get to school on time. It would be foolish to simply unpack your new clock from its box, power it up, and set the wake-up time to 7:00 AM expecting it will wake you at the correct time. Before *trusting* this alarm time of 7:00 AM, you would first have to synchronize your new clock to some standard time source (such as the time broadcast by your local telephone service, or better yet the shortwave radio broadcast of WWV or WWVH⁷) so that it accurately registers time for the zone in which you live. Otherwise, the wake-up setting of 7:00 AM will be hopelessly uncertain.

Once your clock is synchronized against a trusted time source, however, the wake-up (alarm) time may be set at will. If your class schedule changed, allowing one more hour of sleep, you could re-set the wake-up time from 7:00 AM to 8:00 AM without any need to re-synchronize (re-calibrate) the clock. The only reason for re-synchronizing your clock to the time standard is to compensate for inevitable *drift* due to imperfections in the clock circuitry.

Synchronizing the clock to a standard time source is analogous to “calibrating” or “trimming” a smart transmitter: you are establishing an accurate correspondence between what the device’s microprocessor *perceives* and what the actual (real-life) values *are*. This step need only be done at the very beginning of the device’s service, and every so often as warranted by the device’s calibration drift over time⁸.

Setting the wake-up (alarm) time on the clock is analogous to setting the LRV and URV parameters of a smart transmitter: you are defining the *action(s)* taken by the device at certain measured values. For the alarm clock, you are defining the hour and minute of day when the alarm sounds. For the transmitter, you are defining the measured variable values at which it will output 4 mA and 20 mA (for a 4-20 mA analog output range).

By contrast, an analog transmitter blends the functions of calibration and ranging into one. A useful analogy for this is to imagine using a simple wind-up mechanical timer to wake you at 7:00 AM. Such a crude timing device does not even register time in hours and minutes like a digital alarm clock: instead, it simply counts down time from its starting point and sounds an alarm when the descending count reaches zero. In order to set this device for a 7:00 AM wake-up alarm, you must first determine the current time and then calculate how many hours the timer must run before the time reaches 7:00 AM (e.g. if you are setting the wind-up alarm when you go to bed at 10:30 PM, this would equate to a timing period of 8.5 hours).

Every single time you set this wind-up alarm, you must consult a time standard to know how many hours and minutes of count-down time to set it for. If you decide to wake up at a different time in the morning, you must (once again) consult a standard time source, perform the necessary

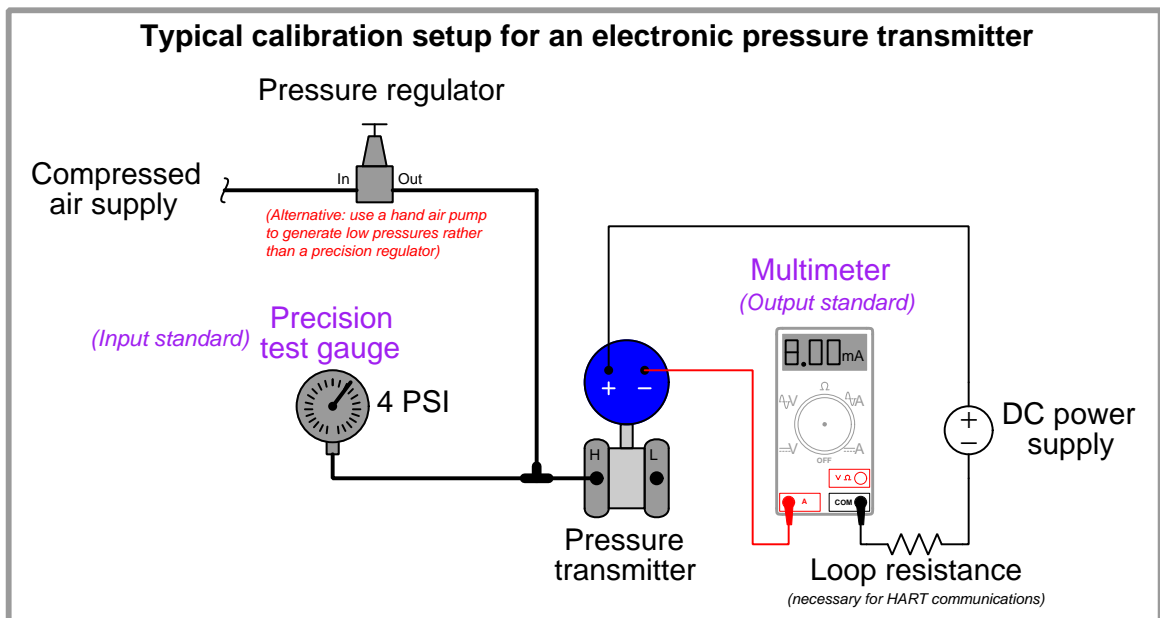
⁷The NIST broadcasts audio transmissions of “Coordinated Universal Time” (UTC) on the shortwave radio frequencies 5 MHz, 10 MHz, 15 MHz, 20 MHz, and 25 MHz. Announcements of time, in English, occur at the top of every minute.

⁸In the case of pressure transmitters, re-trimming may be necessary if the device is ever re-mounted in a different orientation. Changing the physical orientation of a pressure transmitter alters the direction in which gravity tugs on the sensing element, causing it to respond as though a constant bias pressure were applied to it. This bias is often on the order of an inch of water column (or less), and usually consequential only for low-pressure applications such as furnace draft pressure.

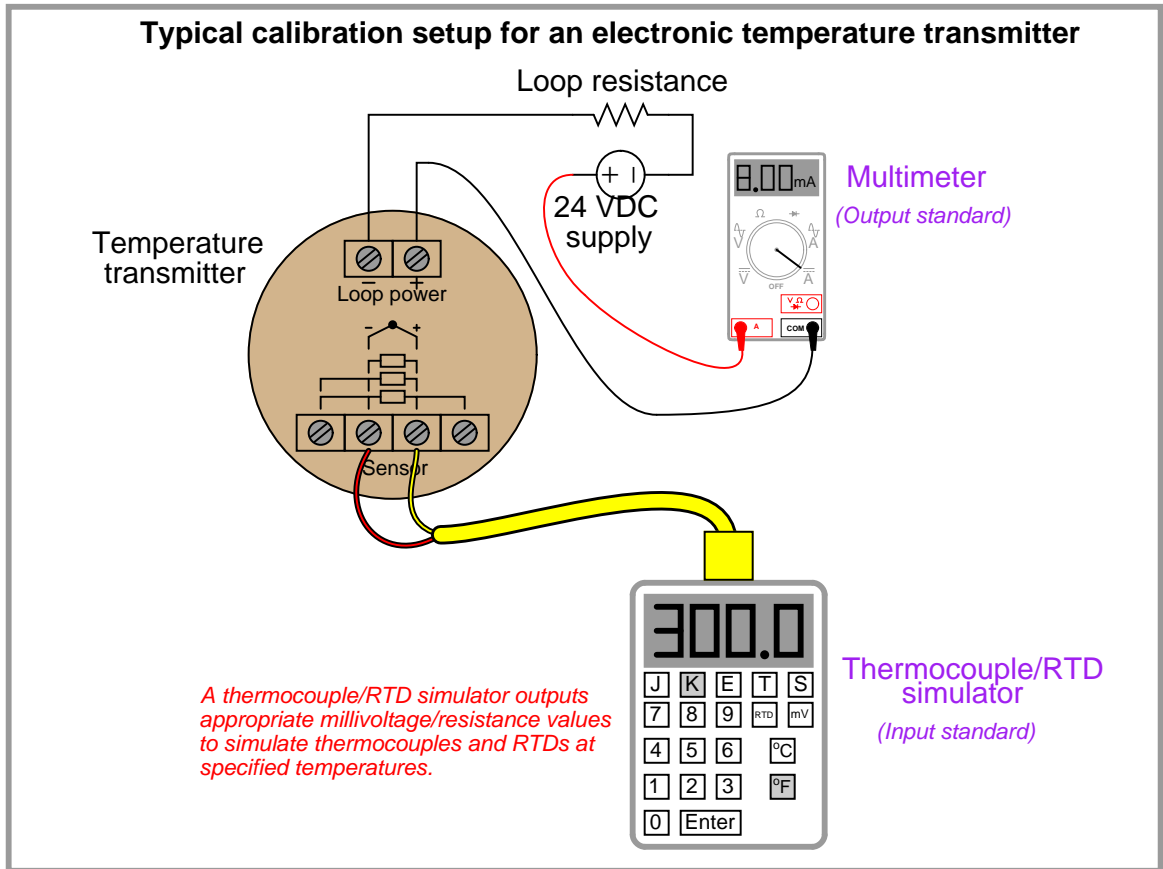
arithmetic, and set the timer accordingly. Setting the alarm time on this mechanism necessitates re-calibrating it to the local standard time without exception. Here, there is no distinction between synchronization and alarm setting; no distinction between calibration and ranging – to do one is to do the other.

18.7 Calibration procedures

As described earlier in this chapter, *calibration* refers to the adjustment of an instrument so its output accurately corresponds to its input throughout a specified range. The only way we can know that an instrument's output accurately corresponds to its input over a continuous range is to subject that instrument to known input values while measuring the corresponding output signal values. This means we must use trusted *standards* to establish known input conditions and to measure output⁹ signals. The following examples show both input and output standards used in the calibration of pressure and temperature transmitters:



⁹A noteworthy exception is the case of digital instruments, which output digital rather than analog signals. In this case, there is no need to compare the digital output signal against a standard, as digital numbers are not liable to calibration drift. However, the calibration of a digital instrument still requires comparison against a trusted standard in order to validate an analog quantity. For example, a digital pressure transmitter must still have its input calibration values validated by a pressure standard, even if the transmitter's digital output signal cannot drift or be misinterpreted.



It is the purpose of this section to describe procedures for efficiently calibrating different types of instruments.

18.7.1 Linear instruments

The simplest calibration procedure for an analog, linear instrument is the so-called *zero-and-span* method. The method is as follows:

1. Apply the lower-range value stimulus to the instrument, wait for it to stabilize
2. Move the “zero” adjustment until the instrument registers accurately at this point
3. Apply the upper-range value stimulus to the instrument, wait for it to stabilize
4. Move the “span” adjustment until the instrument registers accurately at this point
5. Repeat steps 1 through 4 as necessary to achieve good accuracy at both ends of the range

An improvement over this crude procedure is to check the instrument’s response at several points between the lower- and upper-range values. A common example of this is the so-called *five-point calibration* where the instrument is checked at 0% (LRV), 25%, 50%, 75%, and 100% (URV) of range. A variation on this theme is to check at the five points of 10%, 25%, 50%, 75%, and 90%, while still making zero and span adjustments at 0% and 100%. Regardless of the specific percentage points chosen for checking, the goal is to ensure that we achieve (at least) the minimum necessary accuracy at all points along the scale, so the instrument’s response may be trusted when placed into service.

Yet another improvement over the basic five-point test is to check the instrument’s response at five calibration points *decreasing* as well as *increasing*. Such tests are often referred to as *Up-down* calibrations. The purpose of such a test is to determine if the instrument has any significant *hysteresis*: a lack of responsiveness to a change in direction.

Some analog instruments provide a means to adjust linearity. This adjustment should be moved only if absolutely necessary! Quite often, these linearity adjustments are very sensitive, and prone to over-adjustment by zealous fingers. The linearity adjustment of an instrument should be changed only if the required accuracy cannot be achieved across the full range of the instrument. Otherwise, it is advisable to adjust the zero and span controls to “split” the error between the highest and lowest points on the scale, and leave linearity alone.

The procedure for calibrating a “smart” digital transmitter – also known as *trimming* – is a bit different. Unlike the zero and span adjustments of an analog instrument, the “low” and “high” trim functions of a digital instrument are typically non-interactive. This means you should only have to apply the low- and high-level stimuli *once* during a calibration procedure. Trimming the sensor of a “smart” instrument consists of these four general steps:

1. Apply the lower-range value stimulus to the instrument, wait for it to stabilize
2. Execute the “low” sensor trim function
3. Apply the upper-range value stimulus to the instrument, wait for it to stabilize
4. Execute the “high” sensor trim function

Likewise, trimming the output (Digital-to-Analog Converter, or DAC) of a “smart” instrument consists of these six general steps:

1. Execute the “low” output trim test function
2. Measure the output signal with a precision milliammeter, noting the value after it stabilizes
3. Enter this measured current value when prompted by the instrument
4. Execute the “high” output trim test function
5. Measure the output signal with a precision milliammeter, noting the value after it stabilizes
6. Enter this measured current value when prompted by the instrument

After both the input and output (ADC and DAC) of a smart transmitter have been trimmed (i.e. calibrated against standard references known to be accurate), the lower- and upper-range values may be set. In fact, once the trim procedures are complete, the transmitter may be ranged and ranged again as many times as desired. The only reason for re-trimming a smart transmitter is to ensure accuracy over long periods of time where the sensor and/or the converter circuitry may have drifted out of acceptable limits. This stands in stark contrast to analog transmitter technology, where re-ranging *necessitates* re-calibration every time.

18.7.2 Nonlinear instruments

The calibration of inherently nonlinear instruments is much more challenging than for linear instruments. No longer are two adjustments (zero and span) sufficient, because more than two points are necessary to define a curve.

Examples of nonlinear instruments include expanded-scale electrical meters, square root characterizers, and position-characterized control valves.

Every nonlinear instrument will have its own recommended calibration procedure, so I will defer you to the manufacturer’s literature for your specific instrument. I will, however, offer one piece of advice: when calibrating a nonlinear instrument, document all the adjustments you make (e.g. how many turns on each calibration screw) just in case you find the need to “re-set” the instrument back to its original condition. More than once I have struggled to calibrate a nonlinear instrument only to find myself further away from good calibration than where I originally started. In times like these, it is good to know you can always reverse your steps and start over!

18.7.3 Discrete instruments

The word “discrete” means *individual* or *distinct*. In engineering, a “discrete” variable or measurement refers to a true-or-false condition. Thus, a discrete sensor is one that is only able to indicate whether the measured variable is above or below a specified setpoint.

Examples of discrete instruments are *process switches* designed to turn on and off at certain values. A pressure switch, for example, used to turn an air compressor on if the air pressure ever falls below 85 PSI, is an example of a discrete instrument.

Discrete instruments require periodic calibration just like continuous instruments. Most discrete instruments have just one calibration adjustment: the *set-point* or *trip-point*. Some process switches have two adjustments: the set-point as well as a *deadband* adjustment. The purpose of a deadband adjustment is to provide an adjustable buffer range that must be traversed before the switch changes state. To use our 85 PSI low air pressure switch as an example, the set-point would be 85 PSI, but if the deadband were 5 PSI it would mean the switch would not change state until the pressure rose above 90 PSI (85 PSI + 5 PSI).

When calibrating a discrete instrument, you must be sure to check the accuracy of the set-point *in the proper direction of stimulus change*. For our air pressure switch example, this would mean checking to see that the switch changes states at 85 PSI *falling*, not 85 PSI *rising*. If it were not for the existence of deadband, it would not matter which way the applied pressure changed during the calibration test. However, deadband will always be present in a discrete instrument, whether that deadband is adjustable or not.

For example, a pressure switch with a deadband of 5 PSI set to trip at 85 PSI falling would re-set at 90 PSI rising. Conversely, a pressure switch (with the same deadband of 5 PSI) set to trip at 85 PSI rising would re-set at 80 PSI falling. In both cases, the switch “trips” at 85 PSI, but the direction of pressure change specified for that trip point defines which side of 85 PSI the re-set pressure will be found.

A procedure to efficiently calibrate a discrete instrument without too many trial-and-error attempts is to set the stimulus at the desired value (e.g. 85 PSI for our hypothetical low-pressure switch) and then move the set-point adjustment in the *opposite* direction as the intended direction of the stimulus (in this case, *increasing* the set-point value until the switch changes states). The basis for this technique is the realization that most comparison mechanisms cannot tell the difference between a rising process variable and a falling setpoint (or vice-versa). Thus, a falling pressure may be simulated by a rising set-point adjustment. You should still perform an actual changing-stimulus test to ensure the instrument responds properly under realistic circumstances, but this “trick” will help you achieve good calibration in less time.

18.8 Instrument turndown

An important performance parameter for transmitter instruments is something often referred to as *turndown* or *rangedown*. “Turndown” is defined as the ratio of maximum allowable span to the minimum allowable span for a particular instrument.

Suppose a pressure transmitter has a maximum calibration range of 0 to 300 pounds per square inch (PSI), and a turndown of 20:1. This means that a technician may adjust the span anywhere between 300 PSI (e.g. range = 0 to 300 PSI) and 15 PSI (e.g. range = 0 to 15 PSI). This is important to know in order to select the proper transmitter for any given measurement application. The odds of you finding a transmitter with just the perfect factory-calibrated range for your measurement application may be quite small, meaning you will have to adjust its range to fit your needs. The turndown ratio tells you how far you will be able to practically adjust your instrument’s range.

For example, suppose you were working at a facility where the operations personnel requested a pressure transmitter installed on a process vessel with a measurement range of 50 PSI to 90 PSI. You go to the warehouse where all the new instruments are stocked, and find a pressure transmitter with a (maximum) range of zero to 1000 PSI, and a turndown ratio of 20:1. Dividing the maximum span of 1000 PSI by 20, we arrive at a minimum span of 50 PSI. The span requested by operations for this pressure transmitter is 40 PSI (90 PSI – 50 PSI), which means the transmitter you found in the warehouse will *not* be able to “turn down” that far. At best, we could range it for 50 PSI to 100 PSI, or perhaps for 40 PSI to 90 PSI, but not the 50 PSI to 90 PSI requested by operations. At this point, you could return to the operations personnel to ask if a 50 PSI span would be acceptable – if not, you will have to order a different pressure transmitter with a smaller span (or with a greater turndown ratio¹⁰).

Another important consideration with turndown is the *accuracy* of the instrument at the stated turndown. The further an instrument is “turned down” from its maximum span, generally the worse its accuracy becomes at that reduced span. For example, the Micro Motion “ELITE” series of Coriolis mass flowmeters¹¹ are advertised to perform within an accuracy envelope of $\pm 0.05\%$ at turndown ratios up to 20:1, but that measurement uncertainty increases to $\pm 0.25\%$ at a turndown of 100:1, and to $\pm 1.25\%$ at a turndown of 500:1. It should be noted that the degradation of measurement accuracy at large turndown ratios is not some defect of Micro Motion flowmeters (far from it!), but rather an inescapable consequence of pushing an instrument’s turndown to its limit.

¹⁰Modern “smart” electronic pressure transmitters typically boast turndown ratios exceeding 100:1, with some having turndown ratios of 200:1 or more! Large turndown ratios are good because they allow users of instrumentation to maintain a smaller quantity of new transmitters in stock, since transmitters with large turndown ratios are more versatile (i.e. applicable to a wider variety of spans) than transmitters with small turndown ratios.

¹¹According to Emerson product datasheet PS-00374, revision L, June 2009.

18.9 NIST traceability

As defined previously, *calibration* means the comparison and adjustment (if necessary) of an instrument's response to a stimulus of precisely known quantity, to ensure operational accuracy. In order to perform a calibration, one must be reasonably sure that the physical quantity used to stimulate the instrument is accurate in itself. For example, if I try calibrating a pressure gauge to read accurately at an applied pressure of 200 PSI, I must be reasonably sure that the pressure I am using to stimulate the gauge is actually 200 PSI. If it is not 200 PSI, then all I am doing is adjusting the pressure gauge to register 200 PSI when in fact it is sensing something different.

Ultimately, this is a philosophical question of epistemology: *how do we know what is true?* There are no easy answers here, but teams of scientists and engineers known as *metrologists* devote their professional lives to the study of calibration standards to ensure we have access to the best approximation of "truth" for our calibration purposes. *Metrology* is the science of measurement, and the central repository of expertise on this science within the United States of America is the *National Institute of Standards and Technology*, or the *NIST* (formerly known as the *National Bureau of Standards*, or *NBS*).

Experts at the NIST work to ensure we have means of tracing measurement accuracy back to *intrinsic standards*, which are quantities inherently fixed (as far as anyone knows). The vibrational frequency of an isolated cesium atom when stimulated by radio energy, for example, is an intrinsic standard used for the measurement of time (forming the basis of the so-called *atomic clock*). So far as anyone knows, this frequency is fixed in nature and cannot vary: each and every isolated cesium atom has the exact same resonant frequency. The distance traveled in a vacuum by 1650763.73 wavelengths of light emitted by an excited krypton-86 (^{86}Kr) atom is the intrinsic standard for one meter of length. Again, so far as anyone knows, this distance is fixed in nature and cannot vary. This means any suitably equipped laboratory in the world should be able to build their own intrinsic standards to reproduce the *exact* same quantities based on the same (universal) physical constants. The accuracy of an intrinsic standard is ultimately a function of nature rather than a characteristic of the device. Intrinsic standards therefore serve as absolute references which we may calibrate certain instruments against.

The machinery necessary to replicate intrinsic standards for practical use is quite expensive and usually delicate. This means the average metrologist (let alone the average industrial instrument technician) simply will never have access to one. While the concept of an intrinsic standard is tantalizing in its promise of ultimate accuracy and repeatability, it is simply beyond the reach of most laboratories to maintain. An example of an intrinsic standard is this Josephson Junction array in the primary metrology lab at the Fluke corporation's headquarters in Everett, Washington:



A Josephson junction functions as an intrinsic standard for *voltage*, generating extremely precise DC voltages in response to a DC excitation current and a microwave radiation flux. Josephson junctions are superconducting devices, and as such must be operated in an extremely cold environment, hence the dewar vessel filled with liquid helium in the right-hand side of the photograph. The microwave radiation flux itself must be of a precisely known frequency, as the Josephson voltage varies in direct proportion to this frequency. Thus, the microwave frequency source is synchronized with the NIST's atomic clock (another intrinsic standard).

While theoretically capable of generating voltages with uncertainties in the low *parts per billion* range, a Josephson Array such as this one maintained by Fluke is quite an expensive¹² beast, being too impractical for most working labs and shops to justify owning. In order for these intrinsic

¹²According to the book *Philosophy in Practice* (second edition) published by Fluke, the initial expense of their Josephson Array in 1992 was \$85000, with another \$25000 budgeted for start-up costs. The annual operating cost of the array is approximately \$10000, mostly due to the cost of the liquid helium refrigerant necessary to keep the Josephson junction array at a superconducting temperature. This consumable cost does not include the salary of the personnel needed to maintain the system, either. Presumably, a metrology lab of this caliber would employ several engineers and scientists to maintain all standards in top condition and to perform continuing metrological research.

standards to be useful within the industrial world, we use them to calibrate other instruments, which are then used to calibrate other instruments, and so on until we arrive at the instrument we intend to calibrate for field service in a process. So long as this “chain” of instruments is calibrated against each other regularly enough to ensure good accuracy at the end-point, we may calibrate our field instruments with confidence. The documented confidence is known as *NIST traceability*: that the accuracy of the field instrument we calibrate is ultimately ensured by a trail of documentation leading to intrinsic standards maintained by the NIST. This “paper trail” proves to anyone interested that the accuracy of our calibrated field instruments is of the highest pedigree.

18.10 Practical calibration standards

As previously defined, *calibration* refers to the checking and adjustment of an instrument so that its output faithfully corresponds to its input throughout a specified range. In order to calibrate an instrument, we must have some means of knowing the input and/or output quantities associated with the instrument under test. A substance or device used as a reference to compare against an instrument's response is called a *calibration standard*. Simply put, a calibration standard is something we may *compare* the calibrated instrument to. Thus, any calibration can only be as good as the standard used¹³.

Calibration standards fall into two broad categories: standards used to *produce* accurate physical quantities (e.g. pressure, temperature, voltage, current, etc.), and standards used to simply *measure* physical quantities to a high degree of accuracy. An example of the former would be the use of boiling water (at sea level) to *produce* a temperature of 100 degrees Celsius (212 degrees Fahrenheit) in order to calibrate a temperature gauge, whereas an example of the latter would be the use of a laboratory-quality precision thermometer to measure some arbitrary source of temperature in comparison to the temperature gauge being calibrated.

In metrology labs, the ultimate standards are based on fundamental constants of nature, and are called *intrinsic standards*. A modern example of an intrinsic standard for time is the so-called *atomic clock*, using isolated atoms of Cesium to produce frequencies which are inherently fixed and reproduceable world-wide. Instrument shops located in industrial facilities cannot afford the capital and consumable costs associated with intrinsic standards, and so must rely on other devices for their calibration purposes. Ideally, there should be a "chain" of calibration from any device used as a shop standard traceable all the way back to some intrinsic standard in a national-level or primary metrology lab.

Calibration standards used in instrument shops for industrial calibration work should therefore be periodically sent to a local metrology lab for re-standardization, where their accuracy may be checked against other (higher-level) standards which themselves are checked against even higher-level calibration standards, ultimately traceable all the way to intrinsic standards. In each step of the calibration "chain," there is a progressive degree of measurement uncertainty. Intrinsic standards possess the least amount of uncertainty, while field instruments (e.g. pressure transmitters, temperature gauges, etc.) exhibit the greatest uncertainties.

It is important that the degree of uncertainty in the accuracy of a test instrument is *significantly less* than the degree of uncertainty we hope to achieve in the instruments we calibrate. Otherwise, calibration becomes a pointless exercise. This ratio of uncertainties is called the *Test Uncertainty Ratio*, or *TUR*. A good rule-of-thumb is to maintain a TUR of at least 4:1 (ideally 10:1 or better), the test equipment being many times more accurate (less uncertain) than the field instruments we calibrate with them.

I have personally witnessed the confusion and wasted time that results from trying to calibrate a

¹³This brings to mind a good joke. Once there was a man who walked by an antique store every day on his way to work and noticed all the wall clocks on display at this store always perfectly matched in time. One day he happened to see the store owner and complimented him on the consistent accuracy of his display clocks, noting how he used the owner's clocks as a standard to set his own wristwatch on his way to work. He then asked the owner how he kept all the clocks so perfectly set. The owner explained he set the clocks to the sound of the steam whistle at the local factory, which always blew precisely at noon. The store owner then asked the man what he did for a living. The man replied, "I operate the steam whistle at the factory."

field instrument to a tighter tolerance than what the calibration standard is capable of. In one case, an instrument technician attempted to calibrate a pneumatic pressure transmitter to a tolerance of $\pm 0.25\%$ of span using a test gauge that was only good for $\pm 1\%$ of the same span. This poor technician kept going back and forth, adjusting the transmitter's zero and span screws over and over again in a futile attempt to reign in the transmitter's response within the stated specification of $\pm 0.25\%$. After giving up, he tested the test gauges by comparing three of them at once, tied together on a common air pressure tube. When he did this, it became clear that no two test gauges would consistently agree with each other within the specified tolerance over the 3 to 15 PSI range. As he raised and lowered the pressure, the gauges' indications would deviate from one another far more than $\pm 0.25\%$ across the measurement range. Simply put, the inherent uncertainty of the gauges exceeded the uncertainty he was trying to calibrate the transmitter to. As a result, his calibration "standard" was in fact shifting on him as he performed the calibration. His actions were analogous to trying to set up a fixed-position cannon to repeatedly hit a moving target.

The lesson to be learned here is to always ensure the standards used to calibrate industrial instruments are reliably accurate (enough). No calibration standard is *perfect*, but perfection is not what we need. Our goal is to be *accurate enough* that the final calibration will be reliable within specified boundaries.

The next few subsections describe various standards used in instrument shops to calibrate industrial instruments.

18.10.1 Electrical standards


Electrical calibration equipment – used to calibrate instruments measuring voltage, current, and resistance – must be periodically calibrated against higher-tier standards maintained by outside laboratories. In years past, instrument shops would often maintain their own *standard cell* batteries (often called *Weston* cells) as a primary voltage reference. These special-purpose batteries produced 1.0183 volts DC at room temperature with low uncertainty and drift, but were sensitive to vibration and non-trivial to actually use. Now, electronic voltage references have all but displaced standard cells in calibration shops and laboratories, but these references must be checked and adjusted for drift in order to maintain their NIST traceability.

One enormous benefit of electronic calibration references is that they are able to generate accurate currents and resistances in addition to voltage (and not just voltage at one fixed value, either!). Modern electronic references are digitally-controlled as well, which lends themselves well to automated testing in assembly-line environments, and/or programmed multi-point calibrations with automatic documentation of as-found and as-left calibration data. A photograph of some electronic calibration references appears here:



If a shop cannot afford one of these versatile references for benchtop calibration use, an acceptable alternative in some cases is to purchase a high-accuracy multimeter and equip the calibration bench with adjustable voltage, current, and resistance sources. These sources will be simultaneously connected to the high-accuracy multimeter and the instrument under test, and adjusted until the high-accuracy meter registers the desired value. The measurement shown by the instrument under test is then compared against the reference meter and adjusted until matching (to within the required tolerance).

The following illustration shows how a high-accuracy voltmeter could be used to calibrate a handheld voltmeter in this fashion:



It should be noted that the variable voltage source shown in this test arrangement need not be sophisticated. It simply needs to be *variable* (to allow precise adjustment until the high-accuracy voltmeter registers the desired voltage value) and *stable* (so the adjustment will not drift appreciably over time). The accuracy of your calibration in the previous circuit originates not from the variable voltage source, but rather from the high-accuracy multimeter used as the calibration standard. It is the high-accuracy multimeter that serves as the calibration reference here, not the voltage source – it is the high-accuracy multimeter that functions as the *standard*.

18.10.2 Temperature standards

The most common technologies for industrial temperature measurement are electrical in nature: RTDs and thermocouples. As such, the standards used to calibrate such devices are the same standards used to calibrate electrical instruments such as digital multimeters (DMMs). For RTDs, this means a precision resistance standard such as a *decade box* used to precisely set known quantities of electrical resistance. For thermocouples, this means a *precision potentiometer* used to generate precise quantities of low DC voltage (in the millivolt range, with microvolt resolution).

Photographs of antique potentiometers used to calibrate thermocouple-sensing temperature instruments appear here:




Modern, electronic calibrators are also available now for RTD and thermocouple instrument calibration, capable of sourcing accurate quantities of electrical resistance and DC millivoltage for the simulation of RTD and thermocouple elements, respectively. A photograph of a Fluke model 525A laboratory standard is shown here:



Both the antique potentiometers and modern laboratory calibrators such as the Fluke 525A are self-contained *sources* useful for simulating the electrical outputs of temperature sensors. If you closely observe the potentiometer photos, you can see numbers engraved around the circumference of the dials, showing the user how much voltage the device output at any given setting.

Given an accurate enough voltmeter, it is possible to construct your own calibration potentiometer for simulating the millivoltage output of a thermocouple. A simple voltage divider set up to reduce the DC voltage of an ordinary variable-voltage power supply will suffice, so long as it provides fine enough adjustment:



Unlike the potentiometers of old, providing direct read-out of millivoltage at the potentiometer dial(s), we rely here on the accuracy of the precision multimeter to tell us when we have reached the required millivoltage with our power supply and voltage divider circuit. This means the high-accuracy multimeter functions as the calibration *standard* in this set-up, permitting the use of non-precision components in the rest of the circuit. Since the multimeter's indication is the only variable being trusted as accurate when calibrating the thermocouple-input temperature transmitter, the multimeter is the only¹⁴ component in the circuit affecting the uncertainty of our calibration.

¹⁴This, of course, assumes the potentiometer has a sufficiently fine adjustment capability that we may adjust the millivoltage signal to any desired precision. If we were forced to use a coarse potentiometer – incapable of being adjusted to the precise amount of millivoltage we desired – then the accuracy of our calibration would also be limited by our inability to precisely control the applied voltage.

Electrically simulating the output of a thermocouple or RTD may suffice when the instrument we wish to calibrate uses a thermocouple or an RTD as its sensing element. However, there are some temperature-measuring instruments that are not electrical in nature: this category includes bimetallic thermometers, filled-bulb temperature systems, and optical pyrometers. In order to calibrate these types of instruments, we must accurately create the calibration temperatures in the instrument shop. In other words, the instrument to be calibrated must be subjected to an actual temperature of accurately known value.

Even with RTDs and thermocouples – where the sensor signal may be easily simulated using electronic test equipment – there is merit in using an actual source of precise temperature to calibrate the temperature instrument. Simulating the voltage produced by a thermocouple at a precise temperature, for example, is fine for calibrating the instrument normally receiving the millivoltage signal from the thermocouple, but this calibration test does nothing to validate the accuracy of the thermocouple element itself! The *best* type of calibration for any temperature-measuring instrument, from the perspective of overall integrity, is to actually subject the sensing element to a precisely known temperature. For this we need special calibration equipment designed to produce accurate temperature samples on demand.

A time-honored standard for low-temperature industrial calibrations is pure water, specifically the freezing and boiling points of water. Pure water at sea level (full atmospheric pressure) freezes at 32 degrees Fahrenheit (0 degrees Celsius) and boils at 212 degrees Fahrenheit (100 degrees Celsius). In fact, the Celsius temperature scale is *defined* by these two points of phase change for water at sea level¹⁵.

To use water as a temperature calibration standard, simply prepare a vessel for one of two conditions: thermal equilibrium at freezing or thermal equilibrium at boiling. “Thermal equilibrium” in this context simply means equal temperature throughout the mixed-phase sample. In the case of freezing, this means a well-mixed sample of solid ice and liquid water. In the case of boiling, this means a pot of water at a steady boil (vaporous steam and liquid water in direct contact). What you are trying to achieve here is ample contact between the two phases (either solid and liquid; or liquid and vapor) to eliminate hot or cold spots. When the entire water sample is homogeneous in temperature and heterogeneous in phase (i.e. a mix of different phases), the sample will have only one degree of thermodynamic freedom: its temperature is an exclusive function of atmospheric pressure. Since atmospheric pressure is relatively stable and well-known, this fixes the temperature at a constant value. For ultra-precise temperature calibrations in laboratories, the *triple point* of water is used as the reference. When water is brought to its triple point (i.e. all three phases of solid, liquid, and gas co-existing in direct contact with each other), the sample will have *zero degrees* of thermodynamic freedom, which means both its temperature and its pressure will become locked at stable values: pressure at 0.006 atmospheres, and temperature at 0.01 degrees Celsius.

¹⁵The Celsius scale used to be called the *Centigrade* scale, which literally means “100 steps.” I personally prefer the name “Centigrade” to the name “Celsius” because the former actually describes something about the unit of measurement while the latter is a surname. In the same vein, I also prefer the older label “Cycles Per Second” (cps) to “Hertz” as the unit of measurement for frequency. You may have noticed by now that the instrumentation world does not yield to my opinions, much to my chagrin.

The following photograph shows a *triple-point cell* for water, used as a laboratory calibration standard:



The major limitation of water as a temperature calibration standard is it only provides two¹⁶ points of calibration: 0 °C and 100 °C, with the latter¹⁷ being strongly pressure-dependent. If other reference temperatures are required for a calibration, some substance other than water must be used.

¹⁶Three, if you count the triple point, but this requires more sophisticated testing apparatus to establish than either the freezing or boiling points.

¹⁷Pressure does have some influence on the freezing point of most substances as well, but not nearly to the degree it has on the boiling point. For a comparison between the pressure-dependence of freezing versus boiling points, consult a phase diagram for the substance in question, and observe the slopes of the solid-liquid phase line and liquid-vapor phase line. A nearly-vertical solid-liquid phase line shows a weak pressure dependence, while the liquid-vapor phase lines are typically much closer to horizontal.

A variety of substances with known phase-change points have been standardized as fixed points on the International Practical Temperature Scale (ITS-90). The following list is a sample of some of these substances and their respective phase states and temperatures¹⁸:

- Neon (triple point) = $-248.6\text{ }^{\circ}\text{C}$
- Oxygen (triple point) = $-218.8\text{ }^{\circ}\text{C}$
- Mercury (triple point) = $-38.83\text{ }^{\circ}\text{C}$
- Tin (freezing point) = $231.93\text{ }^{\circ}\text{C}$
- Zinc (freezing point) = $419.53\text{ }^{\circ}\text{C}$
- Aluminum (freezing point) = $660.32\text{ }^{\circ}\text{C}$
- Copper (freezing point) = $1084.62\text{ }^{\circ}\text{C}$

Substances at the triple point must be in thermal equilibrium with solid, liquid, and vaporous phases co-existing. Substances at the freezing point must be a two-phase mixture of solid and liquid (i.e. a liquid *in the process of freezing*, neither a completely liquid nor a completely solid sample). The physical principle at work in all of these examples is that of *latent heat*: the thermal energy exchange required to change the phase of a substance. So long as the minimum heat exchange requirement for complete phase change is not met, a substance in the midst of phase transition will exhibit a fixed temperature, and therefore behave as a temperature *standard*. Small amounts of heat gain or loss to such a sample will merely change the proportion of one phase to another (e.g. how much solid versus how much liquid), but the temperature will remain locked at a constant value until the sample becomes a single phase.

One major disadvantage of using phase changes to produce accurate temperatures in the shop is the limited availability of temperatures. If you need to create some other temperature for calibration purposes, you either need to find a suitable material with a phase change happening at that exact same temperature (good luck!) or you need to find a finely adjustable temperature source and use an accurate thermometer to compare your instrument under test against. The latter scenario is analogous to the use of a high-accuracy voltmeter and an adjustable voltage source to calibrate a voltage instrument: comparing one instrument (trusted to be accurate) against another (under test).

Laboratory-grade thermometers are relatively easy to secure. Variable temperature sources suitable for calibration use include *oil bath* and *sand bath* calibrators. These devices are exactly what they sound like: small pots filled with either oil or sand, containing an electric heating element and a temperature control system using a laboratory-grade (NIST-traceable) thermal sensor. In the case of sand baths, a small amount of compressed air is introduced at the bottom of the vessel to “fluidize” the sand so the grains move around much like the molecules of a liquid, helping the system reach thermal equilibrium. To use a bath-type calibrator, place the temperature instrument to be calibrated such the sensing element dips into the bath, then wait for the bath to reach the desired temperature.

¹⁸For each of these examples, the assumptions of a 100% pure sample and an airless testing environment are made. Impurities in the initial sample and/or resulting from chemical reactions with air at elevated temperatures, may introduce serious errors.


An oil bath temperature calibrator is shown in the following photograph, with sockets to accept seven temperature probes into the heated oil reservoir:



This particular oil bath unit has no built-in indication of temperature suitable for use as the calibration standard. A standard-grade thermometer or other temperature-sensing element must be inserted into the oil bath along with the sensor under test in order to provide a reference indication useful for calibration.

Dry-block temperature calibrators also exist for creating accurate calibration temperatures in the instrument shop environment. Instead of a fluid (or fluidized powder) bath as the thermal medium, these devices use metal blocks with blind (dead-end) holes drilled for the insertion of temperature-sensing instruments.

An inexpensive dry-block temperature calibrator intended for bench-top service is shown in this photograph:



This particular dry-block temperature calibrator does provide direct visual indication of the block temperature by means of a digital display on its front panel. If greater accuracy is desired, a laboratory reference-grade temperature sensor may be inserted into the same block along with the sensor being tested, and that reference-grade sensor relied upon as the temperature standard rather than the dry-block calibrator's digital display.


Optical temperature instruments require a different sort of calibration tool: one that emits radiation equivalent to that of the process object at certain specified temperatures. This type of calibration tool is called a *blackbody calibrator*,¹⁹ having a target area where the optical instrument may be aimed. Like oil and sand bath calibrators, a blackbody calibrator relies on an internal temperature sensing element as a reference, to control the optical emissions of the blackbody target at any specified temperature within a practical range.

¹⁹A "black body" is an idealized object having an emissivity value of exactly one (1). In other words, a black body is a perfect radiator of thermal energy. Interestingly, a blind hole drilled into any object at sufficient depth acts as a black body, and is sometimes referred to as a *cavity radiator*.

18.10.3 Pressure standards

In order to accurately calibrate a pressure instrument in a shop environment, we must create fluid pressures of known magnitude against which we compare the instrument being calibrated. As with other types of physical calibrations, our choice of standards falls into two broad categories: devices that inherently *produce* known pressures versus devices that accurately *measure* pressures created by some (other) adjustable source.

A *deadweight tester* (sometimes referred to as a *dead-test* calibrator) is an example in the former category. These devices *create* accurately known pressures by means of precise masses and pistons of precise area:



After connecting the gauge (or other pressure instrument) to be calibrated, the technician adjusts the secondary piston to cause the primary piston to lift off its resting position and be suspended by oil pressure alone. So long as the mass placed on the primary piston is precisely known, Earth's gravitational field is constant, and the piston is perfectly vertical, the fluid pressure applied to the instrument under test *must* be equal to the value described by the following equation:

$$P = \frac{F}{A}$$

Where,

P = Fluid pressure

F = Force exerted by the action of gravity on the mass ($F_{weight} = mg$)

A = Area of piston

The primary piston area, of course, is precisely set at the time of the deadweight tester's manufacture and does not change appreciably throughout the life of the device.

A very simple deadweight tester unit appears in the next photograph, mounted to a yellow wooden base:




When sufficient pressure has been accumulated inside the tester to overcome the weight on the piston, the piston rises off its rest and “floats” on the pressurized oil, as shown in this close-up photograph:




A common operating practice for any deadweight tester is to gently spin the mass during testing so the primary piston continually rotates within its cylinder. Any motion will prevent static friction from taking hold, helping to ensure the only force on the primary piston is the force of the fluid within the deadweight tester.

Most modern deadweight testers include extra features such as hand pumps and bleed valves in addition to secondary pistons, to facilitate both rapid and precise operation. The next photograph shows a newer deadweight tester, with these extra features:



There is also such a thing as a *pneumatic* deadweight tester. In these devices, a constant flow of gas such as compressed air or bottled nitrogen vents through a bleed port operated by the primary piston. The piston moves as necessary to maintain just enough gas pressure inside the unit to suspend the mass(es) against gravity. This gas pressure passes on to the instrument under test, just as liquid pressure in a hydraulic deadweight tester passes to the test instrument for comparison:

Pneumatic deadweight tester




In fact, the construction and operation of a pneumatic deadweight tester is quite similar to a self-balancing (force-balance) pneumatic instrument mechanism with a baffle/nozzle assembly. A

moving element opens or closes a variable restriction downstream of a fixed restriction to generate a varying pressure. In this case, that pressure directly operates the bleed vent to self-regulate gas pressure at whatever value is necessary to suspend the mass against gravity.

Deadweight testers (both hydraulic and pneumatic) lend themselves well to relatively high pressures, owing to the practical limitations of mass and piston area. You could use a deadweight tester to calibrate a 100 PSI pressure gauge used for measuring water mains pressure, for example, but you could not use a deadweight tester to calibrate a 0 to 1 "W.C. (zero to one inch water column) pressure gauge used to measure draft pressure in a furnace flue.

For low-pressure calibrations, the simple *manometer* is a much more practical standard. Manometers, of course, do not generate pressure on their own. In order to use a manometer to calibrate a pressure instrument, you must connect both devices to a source of variable fluid pressure, typically instrument air through a precision pressure regulator:



The difference in liquid column heights (h) within the manometer shows the pressure applied to the gauge. As with the deadweight tester, the accuracy of this pressure measurement is bound by just a few physical constants, none of which are liable to sudden change. So long as the manometer's liquid density is precisely known, Earth's gravitational field is constant, and the manometer tubes are perfectly vertical, the fluid pressure indicated by the manometer *must* be equal to the value described by the following equation (two different forms given):

$$P = \rho gh \quad (\text{or}) \quad P = \gamma h$$

Where,

P = Fluid pressure


ρ = Mass density of fluid

γ = Weight density of fluid

g = Acceleration of gravity

h = Height difference between manometer liquid columns

Of course, with pressure-measuring test instruments of suitable accuracy (preferably NIST-traceable), the same sort of calibration jig may be used for virtually any desired range of pressures:



When the electronic test gauge is designed for very low pressures (inches of water column), they are sometimes referred to as *electronic manometers*.

Instrument calibrations performed in the field (i.e. in locations near or at the intended point of use rather than in a professionally-equipped shop) are almost always done this way: a pressure-generating source is connected to both the instrument under test and a trusted calibration gauge (“test gauge”), and the two indications are compared at several points along the calibrated range. Test equipment suitable for field pressure calibrations include *slack-tube manometers* made from flexible plastic tubing hung from any available anchor point near eye level, and *test gauges* typically of the helical bourdon tube variety. Portable electronic test gauges are also available for field use, many with built-in hand pumps for generating precise air pressures.

A noteworthy example of a pneumatic pressure calibrator for field use was a device manufactured by the Wallace & Tiernan corporation, affectionately called a *Wally box* by at least one generation of instrument technicians. A “Wally box” consisted of a large dial pressure gauge (several inches in diameter) with a multi-turn needle and a very fine scale, connected to a network of valves and regulators which were used to set different air pressures from any common compressed air source. The entire mechanism was housed in an impact-resistance case for ruggedness. One of the many nice features of this calibration instrument was a selector valve allowing the technician to switch between two different pressures output by independent pressure regulators. Once the two pressure regulator values were set to the instrument’s lower- and upper-range values (LRV and URV), it was possible to switch back and forth between those two pressures at will, making the task of adjusting an analog instrument with interactive zero and span adjustments much easier than it would have been to precisely adjust a single pressure regulator again and again.

18.10.4 Flow standards

Most forms of continuous flow measurement are inferential; that is, we measure flow indirectly by measuring some other variable (such as pressure, voltage, or frequency) directly. With this in mind, we may usually achieve reasonable calibration accuracy simply by calibrating the primary sensor and replacing the flow element (if inspection proves necessary). In the case of an orifice plate used to measure fluid flow rate, this would mean calibrating the differential pressure transmitter to measure pressure accurately and replacing the orifice plate if it shows signs of wear.

In some cases, though, direct validation of flow measurement accuracy is needed. Most techniques of flow rate validation take the form of measuring accumulated fluid volume over time. This may prove to be complicated, especially if the fluids in question are hazardous in any way, and/or the flow rates are large, and/or the fluid is a gas or vapor.

For simple validation of liquid flow rates, the flow may be diverted from its normal path in the process and into a container where either accumulated volume or accumulated weight may be measured over time. If the rate of flow into this container is constant, the accumulated volume (or weight) should increase linearly over time. The actual flow rate may then be calculated by dividing the change in volume (ΔV) by the time period over which the change in volume was measured (Δt). The resulting quotient is the average flow rate between those two points in time, which is an approximation of instantaneous flow rate:

$$\frac{\Delta V}{\Delta t} = \text{Average flow}$$

$$\frac{\Delta V}{\Delta t} \approx \frac{dV}{dt} = \text{Instantaneous flow}$$

If a suitable vessel exists in the process with level-measuring capability (e.g. a liquid storage vessel equipped with a level transmitter), you may apply the same mathematical technique: use that vessel as an accumulator for the flow in question, tracking the accumulated (or lost) volume over time and then calculating $\frac{\Delta V}{\Delta t}$. The accuracy of this technique rests on some additional factors, though:

- The accuracy of the level transmitter (as a *volume* measuring instrument!)
- The ability to ensure only *one* flow path in or out of that vessel

The first condition listed here places significant limitations on the flow calibration accuracy one can achieve with this method. In essence, you are using the level instrument as the “test gauge” for the flow instrument, so it needs to be high-accuracy in order to achieve even reasonable accuracy for the flowmeter being calibrated.

A more sophisticated approach for direct flow validation is the use of a device called a *flow prover*. A “flow prover” is a precision piston-and-cylinder mechanism used to precisely measure a quantity of liquid over time. Process flow is diverted through the prover, moving the piston over time. Sensors on the prover mechanism detect when the piston has reached certain positions, and time measurements taken at those different positions enable the calculation of average flow ($\frac{\Delta V}{\Delta t}$).

18.10.5 Analytical standards

An *analyzer* measures intrinsic properties of a substance sample such as its density, chemical content, or purity. Whereas the other types of instruments discussed in this chapter measure quantities incidental to the composition of a substance (pressure, level, temperature, and flow rate), an analyzer measures something related to the *nature* of substance being processed.

As previously defined, to *calibrate* an instrument means to check and adjust (if necessary) its response so the output accurately corresponds to its input throughout a specified range. In order to do this, one must expose the instrument to an actual input stimulus of precisely known quantity. This is no different for an analytical instrument. In order to calibrate an analyzer, we must expose it to known quantities of substances with the desired physical and/or chemical properties (density, chemical composition, etc.). In other words, we need to use *chemical standards*.

A classic example of this is the calibration of a pH analyzer. pH is the measurement of hydrogen ion activity in an aqueous solution. The standard range of measurement is 0 pH to 14 pH, the number representing a negative power of 10 approximately describing the hydrogen ion molarity of the solution (how many moles of active hydrogen ions per liter of solution)²⁰.

The pH of a solution is typically measured with a pair of special electrodes immersed in the solution, which generate a voltage proportional to the pH of the solution. In order to calibrate a pH instrument, you must have a sample of liquid solution with a known pH value. For pH instrumentation, such calibration solutions are called *buffers*, because they are specially formulated to maintain stable pH values even in the face of (slight levels of) contamination.

pH buffers may be purchased in liquid form or in powder form. Liquid buffer solutions may be used directly out of the bottle, while powdered buffers must be dissolved in appropriate quantities of de-ionized water to generate a solution ready for calibration use. Pre-mixed liquid buffers are convenient to use, but have a fairly limited shelf life. Powdered buffer capsules are generally superior for long-term storage, and also enjoy the advantage of occupying less storage space in their dry state than a liquid buffer solution.

²⁰For example, a solution with a pH value of 4.7 has a concentration of $10^{-4.7}$ moles of active hydrogen ions per liter. For more information on “moles” and solution concentration, see section 3.7 beginning on page 262.

The following photograph shows a few 7.00 pH (± 0.02 pH) buffer capsules ready to be mixed with water to form a usable buffer solution:



After preparing the buffer solution in a cup, the pH probe is inserted into the buffer solution and given time to stabilize²¹. Once stabilized, the pH instrument may be adjusted to register the proper pH value. Buffer solutions should not be exposed to ambient air for any longer than necessary (especially alkaline buffers such as 10.0 pH) due to contamination²². Pre-mixed liquid buffer storage containers should be capped immediately after pouring into working cups. Used buffer solution should be discarded rather than re-used at a later date.

Analyzers designed to measure the concentration of certain gases in air must be calibrated in a similar manner. Oxygen analyzers, for example, used to measure the concentration of free oxygen in the exhaust gases of furnaces, engines, and other combustion processes must be calibrated against known standards of oxygen concentration. An oxygen analyzer designed to measure oxygen concentration over a range of ambient (20.9% oxygen) to 0% oxygen may be calibrated with ambient air as one of the standard values²³, and a sample of pure nitrogen gas (containing 0% oxygen) as the other standard value. An oxygen analyzer intended for the measurement of oxygen concentrations in excess of ambient air would require a different standard, most likely a sample of 100% pure oxygen, as a calibration reference.

²¹A clean and healthy pH probe should stabilize within about 30 seconds of being inserted in a buffer solution.


²²Carbon dioxide gas in ambient air will cause carbonic acid to form in an aqueous solution. This has an especially rapid effect on high-pH (alkaline) buffers.

²³It is assumed that the concentration of oxygen in ambient air is a stable enough quantity to serve as a calibration standard for most industrial applications. It is certainly an *accessible* standard!

An analyzer designed to measure the concentration of hydrogen sulfide (H_2S), a toxic gas produced by anaerobic bacterial decomposition of organic matter, will require a sample of gas with a precisely known concentration of hydrogen sulfide mixed in it as a calibration reference. A typical reference gas concentration might be 25 or 50 parts per million (ppm). Gas mixtures with such precise concentration values as this may be purchased from chemical laboratories for the purpose of calibrating concentration analyzers, and are often referred to as *span gases* because they are used to set the span of analyzer instruments.

Analytical instruments are generally subject to greater drifting over time than instruments that measure incidental quantities such as pressure, level, temperature, or flow rate. It is not uncommon for instrument technicians to be tasked with *daily* calibration checks of certain instruments responsible for monitoring atmospheric or water emissions at industrial facilities. For this reason, it is often practical to equip such critical analyzers with *self-calibration* systems. A self-calibration system is a system of solenoid (electrically controlled on-off) valves and reference gas bottles set up in such a way that a computer is able to switch the analyzer off-line and subject it to standard reference gases on a regular schedule to check calibration. Many analyzers are programmed to automatically calibrate themselves against these reference gases, thus eliminating tedious work for the instrument technician.

A typical self-calibration system for a gas analyzer might look like this:



The gas analyzer is equipped with its own auto-calibration controls and programming, allowing it to periodically shut off the process sample and switch to known reference gases for “zero” and “span” calibration checks. If these checks indicate excessive drift or any other questionable results, the analyzer has the ability to flag a maintenance alarm to alert an instrument technician to a potential problem that may require servicing. This sort of self-calibration and self-diagnostic capability saves the instrument technician from having to spend substantial time running manual calibration checks, yet alerts the technician if anything is in need of actual repair. Barring any component failures within this system, the only maintenance this system will need is periodic replacement of the calibration gas bottles.

18.11 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Linear equations:** any function represented by a straight line on a graph may be represented symbolically by the slope-intercept formula $y = mx + b$. Relevant to instrument input/output scaling.
- **Zero shift:** any shift in the offset of an instrument is fundamentally additive, being represented by the “intercept” (b) variable of the slope-intercept linear formula $y = mx + b$. Relevant to instrument calibration: adjusting the “zero” of an instrument always adds to or subtracts from its response.
- **Span shift:** any shift in the gain of an instrument is fundamentally multiplicative, being represented by the “slope” (m) variable of the slope-intercept linear formula $y = mx + b$. Relevant to instrument calibration: adjusting the “span” of an instrument always multiplies or divides its response.
- **Deadband and hysteresis:** the difference in response with the independent variable increasing versus decreasing. Usually caused by friction in a mechanism. Relevant to the calibration testing of instruments, both analog and discrete. For continuous measurement devices, the response of a sensor at some stimulus value (increasing) will not be the exactly the same as the response of that same sensor at that same value when decreasing. For process switches, the “trip” the value at which a switch changes state when its stimulus increases is not the same value it changes state when its stimulus decreases.

References

Agy, D. et al., *Calibration: Philosophy In Practice*, Second Edition, Fluke Corporation, Everett, WA, 1994.

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.

“Micro Motion ELITE Coriolis Flow and Density Meters”, product data sheet DS-00374 revision L, Micro Motion, Inc., June 2009.

Chapter 19

Continuous pressure measurement


In many ways, pressure is the primary variable for a wide range of process measurements. Many types of industrial measurements are actually inferred from pressure, such as:

- Flow (measuring the pressure dropped across a restriction)
- Liquid level (measuring the pressure created by a vertical liquid column)
- Liquid density (measuring the pressure difference across a fixed-height liquid column)
- Weight (hydraulic load cell)

Even temperature may be inferred from pressure measurement, as in the case of a fluid-filled chamber where fluid pressure and fluid temperature are directly related. As such, pressure is a very important quantity to measure, and measure accurately. This section describes different technologies for the measurement of pressure.


19.1 Manometers

A very simple device used to measure pressure is the *manometer*: a fluid-filled tube where an applied gas pressure causes the fluid height to shift proportionately. This is why pressure is often measured in units of liquid height (e.g. inches of water, inches of mercury). As you can see, a manometer is fundamentally an instrument of *differential* pressure measurement, indicating the difference between two pressures by a shift in liquid column height:




Of course, it is entirely acceptable to simply vent one tube of a manometer and use it as a *gauge* pressure instrument, comparing the applied pressure at one tube against atmospheric pressure in the other.

Liquid column height in a manometer should always be interpreted at the centerline of the liquid column, regardless of the shape of the liquid's meniscus (the curved air/liquid interface):




Manometers come in a variety of forms, the most common being the *U-tube*, *well* (sometimes called a *cistern*), *raised well*, and *inclined*:


U-tube manometer




"Well" manometer




"Raised well" manometer



"Raised-well" inclined manometer



U-tube manometers are very inexpensive, and are generally made from clear plastic (see the left-hand photo). Cistern-style manometers are the norm for calibration bench work, and are typically constructed from metal cisterns and glass tubes (see the right-hand photo):



Inclined manometers are used to measure very low pressures, owing to their exceptional sensitivity (note the fractional scale for inches of water column in the following photograph, extending from 0 to 1.5 inches on the scale, reading left to right):



Note that venting one side of a manometer is standard practice when using it as a *gauge pressure* indicator (responding to pressure in excess of atmospheric). Both pressure ports will be used if the manometer is applied to the measurement of differential pressure, just as in the case of the U-tube manometer first shown in this section. Absolute pressure may also be measured by a manometer, if one of the pressure ports connects to a sealed vacuum chamber. This is how a *mercury barometer* is constructed for the measurement of absolute ambient air pressure: by sealing off one side of a manometer and removing all the air in that side, such that the applied (atmospheric) pressure is always compared against a vacuum.

Manometers incorporating a “well” have the advantage of single-point reading: one need only compare the height of *one* liquid column, not the difference in height between *two* liquid columns. The cross-sectional area of the liquid column in the well is so much greater than that within the transparent manometer tube that the change in height within the well is usually negligible. In cases where the difference is significant, the spacing between divisions on the manometer scale may be skewed to compensate¹.


Inclined manometers enjoy the advantage of increased sensitivity. Since manometers fundamentally operate on the principle of pressure balanced by liquid height, and this liquid height is always measured parallel to the line of gravitational pull (perfectly vertical), inclining the manometer tube means that liquid must travel farther along the tube to generate the same change in (purely)

¹If you are having difficulty understanding this concept, imagine a simple U-tube manometer where one of the tubes is opaque, and therefore one of the two liquid columns cannot be seen. In order to be able to measure pressure just by looking at one liquid column height, we would have to make a custom scale where every inch of height registered as *two* inches of water column pressure, because for each inch of height change in the liquid column we can see, the liquid column we can't see also changes by an inch. A scale custom-made for a well-type manometer is just the same concept, only without such dramatic skewing of scales.

vertical height than it would in a vertical manometer tube. Thus, an inclined manometer tube causes an amplification in liquid motion for a given amount of pressure change, allowing measurements of greater resolution.

19.2 Mechanical pressure elements

Mechanical pressure-sensing elements include the *bellows*, the *diaphragm*, and the *bourdon tube*. Each of these devices converts a fluid pressure into a force. If unrestrained, the natural elastic properties of the element will produce a motion proportional to the applied pressure.




Bellows resemble an accordion constructed from metal instead of fabric. Increasing pressure inside a bellows unit causes it to elongate. A photograph of a bellows is shown here:



A diaphragm is nothing more than a thin disk of material which bows outward under the influence of a fluid pressure. Many diaphragms are constructed from metal, which gives them spring-like qualities. Some diaphragms are intentionally constructed out of materials with little strength, such that there is negligible spring effect. These are called *slack diaphragms*, and they are used in

conjunction with external mechanisms (e.g. springs) producing the necessary restraining force to prevent damage from applied pressure.

The following photograph shows the mechanism of a small pressure gauge using a brass diaphragm as the sensing element:



As pressure is applied to the rear of the diaphragm, it distends upward (away from the table on which it rests as shown in the photograph), causing a small shaft to twist in response. This twisting motion is transferred to a lever which pulls on a tiny link chain wrapped around the pointer shaft, causing it to rotate and move the pointer needle around the gauge scale. Both the needle and scale on this gauge mechanism have been removed for easier viewing of diaphragm and mechanism.

Bourdon tubes are made of spring-like metal alloys bent into a circular shape. Under the influence of internal pressure, a bourdon tube “tries” to straighten out into its original shape before being bent at the time of manufacture.

Most pressure gauges use a bourdon tube as their pressure-sensing element. Most pressure transmitters use a diaphragm as their pressure-sensing element. Bourdon tubes may be made in *spiral* or *helical* forms for greater motion (and therefore greater gauge resolution).

The Bourdon tube pressure element is a very robust and time-tested design. An illustration taken from page 471 of volume 1 of *Cassier's Magazine* published in the year 1891 shows a typical C-shaped bourdon tube pressure gauge mechanism complete with gears and pointing needle:

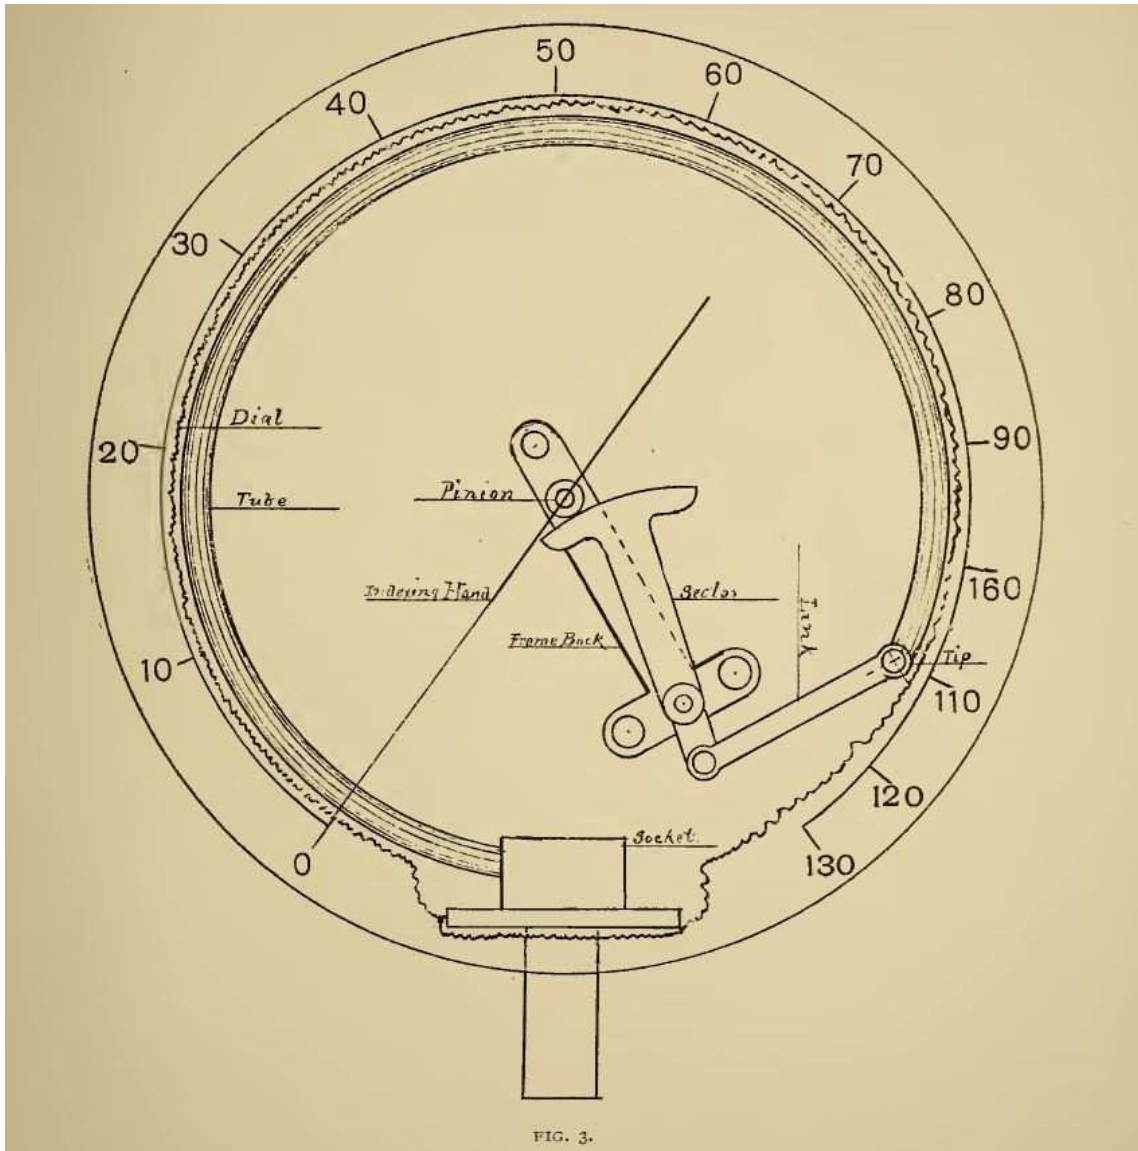


FIG. 3.

Looking closely at the labeled components of this mechanism, we see a circular “pinion” touching a curved “sector”. Both of these are *gears* meshing with one another, but as is typical with mechanical drawings the individual teeth of the meshing gears are not shown.

It is a useful mental exercise to imagine the components of this gauge moving under the influence

of a rising fluid pressure. The bourdon tube will straighten, resulting in its tip extending outward from center (up and right) as its socket remains stationary (anchored to the gauge body). This pulls on the link, which in turn causes the sector gear to rotate counter-clockwise on its bearing axis (that axis anchored on the backing plate of the gauge). This causes the pinion gear to rotate clockwise, driving the needle (the “indexing hand”) clockwise as well, so that the needle’s tip rises up the numerical scale printed on the gauge face.

A photograph of a C-tube pressure gauge mechanism (taken from the rear of the gauge, behind the pointer and scale) reveals its mechanical workings:



The dark, C-shaped tube is the bourdon tube sensing element, while the shiny metal parts are the linkage, lever, and gear assembly.


The next photograph shows a *spiral* bourdon tube, designed to produce a wider range of motion than a C-tube bourdon:



It should be noted that bellows, diaphragms, and bourdon tubes alike may all be used to measure differential and/or absolute pressure in addition to gauge pressure. All that is needed for these other functionalities is to subject the *other* side of each pressure-sensing element to either another applied pressure (in the case of differential measurement) or to a vacuum chamber (in the case of absolute pressure measurement).

This next set of illustrations shows how bellows, diaphragms, and bourdon tubes may be used as differential pressure-sensing elements:

Differential pressure sensing mechanisms



The challenge in doing this, of course, is how to extract the mechanical motion of the pressure-sensing element to an external mechanism (such as a pointer) while maintaining a good pressure seal. In gauge pressure mechanisms, this is no problem because one side of the pressure-sensing element must be exposed to atmospheric pressure anyway, and so that side is always available for mechanical connection.

A differential pressure gauge is shown in the next photograph. The two pressure ports are clearly evident on either side of the gauge:



19.3 Electrical pressure elements

Several different technologies exist for the conversion of fluid pressure into an electrical signal response. These technologies form the basis of electronic *pressure transmitters*: devices designed to measure fluid pressure and transmit that information via electrical signals such as the 4-20 mA analog standard, or in digital form such as HART or FOUNDATION Fieldbus.


A brief survey of electronic pressure transmitters in contemporary² use reveals a diverse representation of electrical pressure-sensing elements:

Manufacturer	Model	Pressure sensor technology
ABB/Bailey	PTSD	Differential reluctance
ABB/Bailey	PTSP	Piezoresistive (strain gauge)
Foxboro	IDP10	Piezoresistive (strain gauge)
Honeywell	ST3000	Piezoresistive (strain gauge)
Rosemount	1151	Differential capacitance
Rosemount	3051	Differential capacitance
Rosemount	3095	Differential capacitance
Yokogawa	EJX series	Mechanical resonance


²As of this writing, 2008.

19.3.1 Piezoresistive (strain gauge) sensors

Piezoresistive means “pressure-sensitive resistance,” or a resistance that changes value with applied pressure. The *strain gauge* is a classic example of a piezoresistive element, a typical strain gauge element shown here on the tip of my finger:



In order to be practical, a strain gauge must be glued (*bonded*) on to a larger specimen capable of withstanding an applied force (stress):



As the test specimen is stretched or compressed by the application of force, the conductors of the strain gauge are similarly deformed. Electrical resistance of any conductor is proportional to the ratio of length over cross-sectional area ($R \propto \frac{l}{A}$), which means that tensile deformation (stretching)

will increase electrical resistance by simultaneously increasing length and decreasing cross-sectional area while compressive deformation (squishing) will decrease electrical resistance by simultaneously decreasing length and increasing cross-sectional area.


Attaching a strain gauge to a diaphragm results in a device that changes resistance with applied pressure. Pressure forces the diaphragm to deform, which in turn causes the strain gauge to change resistance. By measuring this change in resistance, we can infer the amount of pressure applied to the diaphragm.

The classic strain gauge system represented in the previous illustration is made of metal (both the test specimen and the strain gauge itself). Within its elastic limits, many metals exhibit good spring characteristics. Metals, however, are subject to *fatigue* over repeated cycles of strain (tension and compression), and they will begin to “flow” if strained beyond their elastic limit. This is a common source of error in metallic piezoresistive pressure instruments: if overpressured, they tend to lose accuracy due to damage of the spring and strain gauge elements.³

Modern manufacturing techniques have made possible the construction of strain gauges made of silicon instead of metal. Silicon exhibits very linear spring characteristics over its narrow range of motion, and a high resistance to fatigue. When a silicon strain gauge is over-stressed, it fails completely rather than “flows” as is the case with metal strain gauges. This is generally considered a better result, as it clearly indicates the need for sensor replacement (whereas a metallic strain sensor may give the false impression of continued function following an over-stress event).


³For a simple demonstration of metal fatigue and metal “flow,” simply take a metal paper clip and repeatedly bend it back and forth until you feel the metal wire weaken. Gentle force applied to the paper clip will cause it to deform in such a way that it returns to its original shape when the force is removed. Greater force, however, will exceed the paper clip’s elastic limit, causing permanent deformation and also altering the spring characteristics of the clip.

Thus, most modern piezoresistive-based pressure instruments use silicon strain gauge elements to sense deformation of a diaphragm due to applied fluid pressure. A simplified illustration of a diaphragm / strain gauge pressure sensor is shown here:



As the diaphragm bows outward with applied fluid pressure, the strain gauge stretches to a greater length, causing its resistance to increase. This change in resistance imbalances the bridge circuit, causing a voltage (V_{out}) proportional to the amount of applied pressure. Thus, the strain gauge works to convert an applied pressure into a measurable voltage signal which may be amplified and converted into a 4-20 mA loop current signal (or into a digital “fieldbus” signal).

In some designs, a single silicon wafer serves as both the diaphragm and the strain gauge so as to fully exploit the excellent mechanical properties of silicon (high linearity and low fatigue). However, silicon is not chemically compatible with many process fluids, and so pressure must be transferred to the silicon diaphragm/sensor via a non-reactive *fill fluid* (commonly a silicone-based or fluorocarbon-based liquid). A metal *isolating diaphragm* transfers process fluid pressure to the fill fluid, which in turn transfers pressure to the silicon wafer. Another simplified illustration shows how this works:



The isolating diaphragm is designed to be much more flexible (less rigid) than the silicon diaphragm, because its purpose is to seamlessly transfer fluid pressure from the process fluid to the fill fluid, not to act as a spring element. In this way, the silicon sensor experiences the same pressure that it would if it were directly exposed to the process fluid, without having to contact the process fluid. The flexibility of the metal isolating diaphragm also means it experiences much less stress than the silicon sensing diaphragm, which avoiding the problems of metal fatigue experienced by transmitter designs using metal as the sensing (spring) element.


This use of a fill fluid to transfer pressure from an isolating diaphragm to a sensing diaphragm inside the transmitter is used in most if not all modern pressure transmitter designs, even those that are not piezoresistive.

An example of a pressure instrument utilizing a silicon strain gauge element is the Foxboro model IDP10 differential pressure transmitter, shown in the following photograph:



19.3.2 Differential capacitance sensors


Another common electrical pressure sensor design works on the principle of *differential capacitance*. In this design, the sensing element is a taut metal diaphragm located equidistant between two stationary metal surfaces⁴, comprising three plates for a complementary pair of capacitors. An electrically insulating fill fluid (usually a liquid silicone compound) transfers motion from the isolating diaphragms to the sensing diaphragm, and also doubles as an effective dielectric for the two capacitors:



Any difference of pressure across the cell causes the diaphragm to flex in the direction of least pressure. The sensing diaphragm is a precision-manufactured spring element, meaning that its displacement is a predictable function of applied force. The applied force in this case can only be a function of differential pressure acting against the surface area of the diaphragm in accordance with the standard force-pressure-area equation $F = PA$. In this case, we have two forces caused by two fluid pressures working against each other, so our force-pressure-area equation may be rewritten to describe *resultant* force as a function of differential pressure ($P_1 - P_2$) and diaphragm area: $F = (P_1 - P_2)A$. Since diaphragm area is constant, and force is predictably related to diaphragm displacement, all we need now in order to infer differential pressure is to accurately measure displacement of the diaphragm.

⁴In the following diagram, both the sensing diaphragm and the stationary metal surfaces are shown colored blue, to distinguish these electrical elements from the other structural components of the device.

The diaphragm's secondary function as one plate of two capacitors provides a convenient method for measuring displacement. Since capacitance between conductors is inversely proportional to the distance separating them, capacitance on the low-pressure side will increase while capacitance on the high-pressure side will decrease:



A capacitance detector circuit connected to this cell uses a high-frequency AC excitation signal to measure the different in capacitance between the two halves, translating that into a DC signal which ultimately becomes the signal output by the instrument representing pressure.

These pressure sensors are highly accurate, stable, and rugged. An interesting feature of this design – using two isolating diaphragms to transfer process fluid pressure to a single sensing diaphragm through an internal “fill fluid” – is that the solid frame bounds the motion of the two isolating diaphragms such that neither one is able to force the sensing diaphragm past its elastic limit. As the illustration shows, the higher-pressure isolating diaphragm gets pushed toward the metal frame, transferring its motion to the sensing diaphragm via the fill fluid. If too much pressure is applied to that side, the isolating diaphragm will merely “flatten” against the solid frame of the capsule and stop moving. This positively limits the isolating diaphragm’s motion so that it cannot possibly exert any more force on the sensing diaphragm, even if additional process fluid pressure is applied. This use of isolating diaphragms and fill fluid to transfer motion to the sensing diaphragm, employed in other styles of differential pressure sensor as well, gives modern differential pressure instruments excellent resistance to overpressure damage.


It should be noted that the use of a liquid fill fluid is key to this overpressure-resistant design. In order for the sensing diaphragm to accurately translate applied pressure into a proportional capacitance, it must not contact the conductive metal frame surrounding it. In order for any diaphragm to be protected against overpressure, however, it must contact a solid backstop to limit further travel. Thus, the need for non-contact (capacitance) and for contact (overpressure protection) are mutually exclusive, making it nearly impossible to perform both functions with a single sensing diaphragm. Using fill fluid to transfer pressure from isolating diaphragms to the sensing diaphragm allows us to separate the function of capacitive measurement (sensing diaphragm) from the function

of overpressure protection (isolation diaphragms) so that each diaphragm may be optimized for a separate purpose.


A classic example of a pressure instrument based on the differential capacitance sensor is the Rosemount model 1151 differential pressure transmitter, shown in assembled form in the following photograph:



By removing four bolts from the transmitter, we are able to remove two flanges from the pressure capsule, exposing the isolating diaphragms to plain view:



A close-up photograph shows the construction of one of the isolating diaphragms, which unlike the sensing diaphragm is designed to be very flexible. The concentric corrugations in the metal of the diaphragm allow it to easily flex with applied pressure, transmitting process fluid pressure through the silicone fill fluid to the taut sensing diaphragm inside the differential capacitance cell:



The interior of the same differential capacitance sensor (revealed by cutting a Rosemount model 1151 sensor in half with a chop saw⁵) shows the isolating diaphragms, the sensing diaphragm, and the ports connecting them together:



Here, the left-side isolating diaphragm is clearer to see than the right-side isolating diaphragm. A feature clearly evident in this photograph is the small clearance between the left-side isolating diaphragm and the internal metal frame, versus the spacious chamber in which the sensing diaphragm resides. Recall that these internal spaces are normally occupied by *fill fluid*, the purpose of which is to transfer pressure from the isolating diaphragms to the sensing diaphragm. As mentioned before, the solid metal frame limits the travel of each isolating diaphragm in such a way that the higher-pressure isolating diaphragm “bottoms out” on the metal frame before the sensing diaphragm can be pushed past its elastic limit. In this way, the sensing diaphragm is protected against damage from overpressure because the isolating diaphragms are simply not allowed to move any farther.

⁵A chop saw is admittedly not a tool of finesse, and it did a fair job of mangling this unfortunate differential capacitance cell. A bandsaw was tried at first, but made virtually no progress in cutting the hard stainless steel of the capsule assembly. The chop saw’s abrasive wheel created a lot of heat, discoloring the metal and turning the silicone fill fluid into a crystalline mass which had to be carefully chipped out by hand using an ice pick so as to not damage the thin metal sensing diaphragm. Keep these labors in mind, dear reader, as you enjoy this textbook!


The differential capacitance sensor inherently measures *differences* in pressure applied between its two sides. In keeping with this functionality, this pressure instrument has two threaded ports into which fluid pressure may be applied. A later section in this chapter will elaborate on the utility of differential pressure transmitters (section 19.5 beginning on page 1341). All the electronic circuitry necessary for converting the sensor's differential capacitance into an electronic signal representing pressure is housed in the blue-colored structure above the capsule and flanges.

A more modern realization of the differential capacitance pressure-sensing principle is the Rosemount model 3051 differential pressure transmitter:



As is the case with all differential pressure devices, this instrument has *two* ports through which fluid pressure may be applied to the sensor. The sensor, in turn, responds only to the *difference* in pressure between the ports.


The differential capacitance sensor construction is more complex in this particular pressure instrument, with the plane of the sensing diaphragm perpendicular to the plane of the two isolating diaphragms. This “coplanar” design is more compact than the older style of sensor, and more importantly it isolates the sensing diaphragm from flange bolt stress – one of the main sources of error in the previous design⁶.



Take particular note of how the sensor assembly is not embedded in the solid metal frame as was the case with the original Rosemount design. Instead, the sensor assembly is relatively isolated from the frame, connected only by two capillary tubes joining it to the isolating diaphragms. This way, stresses inside the metal frame imparted by flange bolts have virtually no effect on the sensor.

⁶Not only did applied torque of the four capsule bolts affect measurement accuracy in the older 1151 model design, but changes in temperature resulting in changing bolt tension also had a detrimental impact on accuracy. Most modern differential pressure transmitter designs strive to isolate the sensing diaphragm assembly from flange bolt stress for these reasons.

A cutaway model of a Rosemount model 3051S (“supermodule”) DP transmitter shows how this all looks in real life:



Process fluid pressure applied to the isolating diaphragm(s) transfers to fill fluid inside the capillary tubes, conveying pressure to the taut diaphragm inside the differential capacitance sensor. Like the classic Rosemount model 1151 design, we see the fill fluid performing multiple functions:

- The fill fluid protects the delicate sensing diaphragm from contact with unclean or corrosive process fluids
- The fill fluid allows the isolating diaphragms to provide overpressure protection for the sensing diaphragm
- The fill fluid provides a medium of constant permittivity for the differential capacitance circuit to function

The “supermodule” series of Rosemount pressure transmitters shares the same coplanar design as the earlier 3051 models, but adds a new design feature: inclusion of the electronics within the stainless-steel module rather than the blue-painted upper housing. This feature allows the transmitter size to be significantly reduced if needed for applications with limited space.

19.3.3 Resonant element sensors

As any guitarist, violinist, or other stringed-instrument musician can tell you, the natural frequency of a tensed string increases with tension. This, in fact, is how stringed instruments are tuned: the tension on each string is precisely adjusted to achieve the desired resonant frequency.

Mathematically, the resonant frequency of a string may be described by the following formula:

$$f = \frac{1}{2L} \sqrt{\frac{F_T}{\mu}}$$

Where,

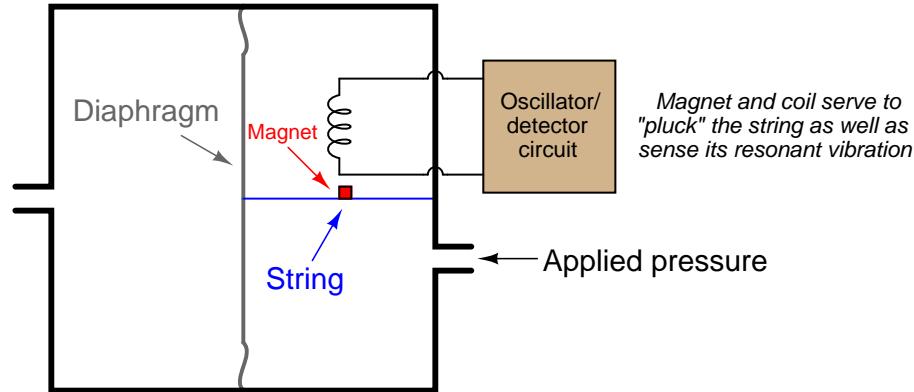
f = Fundamental resonant frequency of string (Hertz)

L = String length (meters)

F_T = String tension (newtons)

μ = Unit mass of string (kilograms per meter)

It stands to reason, then, that a string may serve as a force sensor. All that is needed to complete the sensor is an oscillator circuit to keep the string vibrating at its resonant frequency, and that frequency becomes an indication of tension (force). If the force originates from pressure applied to some sensing element such as a bellows or diaphragm, the string's resonant frequency will indicate fluid pressure. A proof-of-concept device based on this principle might look like this:



It should be noted that this principle of force measurement is nonlinear⁷, as indicated by the equation for resonant frequency (tension force F lies inside the radicand). This means the pressure transmitter must be designed with an electronic characterizing function to “linearize” the frequency measurement into a pressure measurement.

⁷For example, a doubling of force results in a frequency increase of 1.414 (precisely equal to $\sqrt{2}$). A four-fold increase in pressure would be necessary to double the string's resonant frequency. This particular form of nonlinearity, where diminishing returns are realized as the applied stimulus increases, yields excellent rangeability. In other words, the instrument is inherently more sensitive to changes in pressure at the low end of its sensing range, and “de-sensitizes” itself toward the high end of its sensing range.

The Foxboro company pioneered this concept in an early *resonant wire* design of pressure transmitter. Later, the Yokogawa corporation of Japan applied the concept using a pair of micro-machined⁸ silicon resonator structures bonded to a single sensing diaphragm, which became the basis for their successful line of “DPharp” pressure transmitters.

A photograph of a Yokogawa model EJA110 pressure transmitter with this technology is seen here:



Process pressure enters through ports in two flanges, presses against a pair of isolating diaphragms, transferring motion to a single sensing diaphragm via fill fluid where the resonant elements change frequency with diaphragm strain. Motion of the sensing diaphragm in either direction tenses one resonant element and compresses the other, causing their frequencies to deviate from each other. Electronic circuits within the upper housing measure the two resonant elements' frequencies and generate an output signal proportional to their frequency difference. This, of course, is a representation of applied differential pressure.

⁸This is an example of a micro-electro-mechanical system, or *MEMS*.

Even when disassembled, the transmitter does not look much different from the more common differential capacitance sensor design.



The important design differences are hidden from view, inside the sensing capsule. Functionally, though, this transmitter is much the same as its differential-capacitance and piezoresistive cousins. This design even uses fill fluid to protect the delicate silicon resonators from potentially destructive process fluids, just like differential capacitance sensors and most piezoresistive sensor designs.

An interesting advantage of the resonant element pressure sensor is that the sensor signal is easily digitized. The vibration of each resonant element is sensed by the electronics package as an AC frequency. This frequency signal is “counted” by a digital counter circuit over a given span of time and converted to a binary digital representation without any need for an analog-to-digital converter (ADC) circuit. Quartz crystal electronic oscillators are extremely precise, providing the stable frequency reference necessary for comparison in any frequency-based instrument.

In the Yokogawa “DPharp” design, the two resonant elements oscillate at a nominal frequency of approximately 90 kHz. As the sensing diaphragm deforms with applied differential pressure, one resonator experiences tension while the other experiences compression, causing the frequency of the former to shift up and the latter to shift down (as much as ± 20 kHz). The signal conditioning electronics inside the transmitter measures this difference in resonator frequency to infer applied pressure.

19.3.4 Mechanical adaptations

Most modern electronic pressure sensors convert very small diaphragm motions into electrical signals through the use of sensitive motion-sensing techniques (strain gauge sensors, differential capacitance cells, etc.). Diaphragms made from elastic materials behave as springs, but circular diaphragms exhibit very nonlinear behavior when significantly stretched unlike classic spring designs such as coil and leaf springs which exhibit linear behavior over a wide range of motion. Therefore, in order to yield a linear response to pressure, a diaphragm-based pressure sensor must be designed in such a way that the diaphragm stretches very little over the normal range of operation. Limiting the displacement of a diaphragm necessitates highly sensitive motion-detection techniques such as strain gauge sensors, differential capacitance cells, and mechanical resonance sensors to convert that diaphragm's very slight motion into an electronic signal.

An alternative approach to electronic pressure measurement is to use mechanical pressure-sensing elements with more linear pressure-displacement characteristics – such as bourdon tubes and spring-loaded bellows – and then detect the large-scale motion of the pressure element using a less-sophisticated electrical motion-sensing device such as a potentiometer, LVDT, or Hall Effect sensor. In other words, we take the sort of mechanism commonly found in a direct-reading pressure gauge and attach it to a potentiometer (or similar device) to derive an electrical signal from the pressure measurement.

The following photographs show front and rear views of an electronic pressure transmitter using a large C-shaped bourdon tube as the sensing element (seen in the left-hand photograph):




This alternative approach is undeniably simpler and less expensive to manufacture than the more sophisticated approaches used with diaphragm-based pressure instruments, but is prone to greater inaccuracies. Even bourdon tubes and bellows are not perfectly linear spring elements, and the substantial motions involved with using such pressure elements introduces the possibility of hysteresis errors (where the instrument does not respond accurately during reversals of pressure, where the mechanism changes direction of motion) due to mechanism friction, and deadband errors due to backlash (looseness) in mechanical connections.

You are likely to encounter this sort of pressure instrument design in direct-reading gauges equipped with electronic transmitting capability. An instrument manufacturer will take a proven product line of pressure gauge and add a motion-sensing device to it that generates an electric signal proportional to mechanical movement inside the gauge, resulting in an inexpensive pressure transmitter that happens to double as a direct-reading pressure gauge.

19.4 Force-balance pressure transmitters

An important legacy technology for all kinds of continuous measurement is the *self-balancing system*. A “self-balance” system continuously balances an adjustable quantity against a sensed quantity, the adjustable quantity becoming an indication of the sensed quantity once balance is achieved. A common manual-balance system is the type of scale used in laboratories to measure mass:




Here, the unknown mass is the sensed quantity, and the known masses are the adjustable quantity. A human lab technician applies as many masses to the left-hand side of the scale as needed to achieve balance, then counts up the sum total of those masses to determine the quantity of the unknown mass.

Such a system is perfectly linear, which is why these balance scales are popularly used for scientific work. The scale mechanism itself is the very model of simplicity, and the only thing the pointer needs to accurately sense is a condition of balance (equality between masses).

If the task of balancing is given to an automatic mechanism, the adjustable quantity will continuously change and adapt as needed to balance the sensed quantity, thereby becoming a representation of that sensed quantity. In the case of pressure instruments, pressure is easily converted into force by acting on the surface area of a sensing element such as a diaphragm or a bellows. A balancing force may be generated to exactly cancel the process pressure's force, making a *force-balance* pressure instrument. Like the laboratory balance scale, an industrial instrument built on the principle of balancing a sensed quantity with an adjustable quantity will be inherently linear, which is a tremendous advantage for measurement purposes.


Here, we see a diagram of a force-balance pneumatic pressure transmitter⁹, balancing a sensed differential pressure with an adjustable air pressure which becomes a pneumatic output signal:



Differential pressure is sensed by a liquid-filled diaphragm “capsule,” which transmits force to a “force bar.” If the force bar moves out of position due to this applied force, a highly sensitive “baffle” and “nozzle” mechanism senses it and causes a pneumatic amplifier (called a “relay”) to send a different amount of air pressure to a bellows unit. The bellows presses against the “range bar” which pivots to counter-act the initial motion of the force bar. When the system returns to equilibrium, the air pressure inside the bellows will be a direct, linear representation of the process fluid pressure applied to the diaphragm capsule.

⁹Based on the design of Foxboro’s popular model 13A pneumatic “DP cell” differential pressure transmitter.

With minor modifications to the design of this pressure transmitter¹⁰, we may convert it from pneumatic to electronic force-balancing:



Differential pressure is sensed by the same type of liquid-filled diaphragm capsule, which transmits force to the force bar. If the force bar moves out of position due to this applied force, a highly sensitive electromagnetic sensor detects it and causes an electronic amplifier to send a different amount of electric current to a force coil. The force coil presses against the range bar which pivots to counteract the initial motion of the force bar. When the system returns to equilibrium, the milliamperage current through the force coil will be a direct, linear representation of the process fluid pressure applied to the diaphragm capsule.

A distinct advantage of force-balance pressure instruments (besides their inherent linearity) is the constraining of sensing element motion. Unlike a modern diaphragm-based pressure transmitter which relies on the spring characteristics of the diaphragm to convert pressure into force and then

¹⁰Very loosely based on the design of Foxboro's now-obsolete E13 electronic "DP cell" differential pressure transmitter.

into motion (displacement) which is sensed and converted into an electronic signal, a force-balance transmitter works best when the diaphragm is slack and has no spring characteristics at all. Balance with the force of the process fluid pressure is achieved by the application of either an adjustable air pressure or an adjustable electric current, not by the natural tensing of a spring element. This makes a force-balance instrument far less susceptible to errors due to metal fatigue or any other degradation of spring characteristics.

Unfortunately, force-balance instruments have significant disadvantages as well. Force-balance mechanisms tend to be bulky¹¹, and they translate external vibration into inertial force which adds “noise” to the output signal. Also, the amount of electrical power necessary to provide adequate balancing force in an electronic force-balance transmitter is such that it is nearly impossible to limit below the level necessary to ensure intrinsic safety (protection against the accidental ignition of explosive atmospheres by limiting the amount of energy the instrument could possibly discharge into a spark).


19.5 Differential pressure transmitters

One of the most common, and most useful, pressure measuring instruments in industry is the *differential pressure transmitter*. This device senses the difference in pressure between two ports and outputs a signal representing that pressure in relation to a calibrated range. Differential pressure transmitters may be based on any of the previously discussed pressure-sensing technologies, so this section focuses on application rather than theory.


¹¹One instrument technician I know referred to the Foxboro E13 differential pressure transmitter as “pig iron” after having to hoist it by hand to the top of a distillation column.

19.5.1 DP transmitter construction and behavior


Differential pressure transmitters constructed for industrial measurement applications typically consist of a strong (forged metal) body housing the sensing element(s), topped by a compartment housing the mechanical and/or electronic components necessary to translate the sensed pressure into a standard instrumentation signal (e.g. 3-15 PSI, 4-20 mA, digital fieldbus codes):



Two models of electronic differential pressure transmitter appear in the following photographs, the Rosemount model 1151 (left) and model 3051 (right):




Two more models of electronic differential pressure transmitter are shown in the next photograph, the Yokogawa EJA110 (left) and the Foxboro IDP10 (right):




In each of these differential pressure transmitter examples, the pressure-sensing element is housed in the bottom half of the device (the forged-steel structure) while the electronics are housed in the top half (the colored, round, cast-aluminum structure).

Regardless of make or model, every differential pressure (“DP”, “d/p”, or ΔP)¹² transmitter has *two* pressure ports to sense different process fluid pressures. These ports typically have $\frac{1}{4}$ inch female NPT threads for convenient connection to the process. One of these ports is labeled “high” and the other is labeled “low”. This labeling does not necessarily mean that the “high” port must always be at a greater pressure than the “low” port. What these labels represent is the effect any increasing fluid pressure applied to that port will have on the *direction* of the output signal’s change.




¹²As far as I have been able to determine, the labels “D/P” and “DP cell” were originally trademarks of the Foxboro Company. Those particular transmitter models became so popular that the term “DP cell” came to be applied to nearly *all* makes and models of differential pressure transmitter, much like the trademark “Vise-Grip” is often used to describe *any* self-locking pliers, or “Band-Aid” is often used to describe *any* form of self-adhesive bandage.

The most common sensing element used by modern DP transmitters is the diaphragm. One side of this diaphragm receives process fluid pressure from the “high” port, while the other receives process fluid pressure from the “low” port. Any difference of pressure between the two ports causes the diaphragm to flex from its normal resting (center) position. This flexing is then translated into an output signal by any number of different technologies, depending on the manufacturer and model of the transmitter:




The concept of differential pressure instrument port labeling is very similar to the “inverting” and “noninverting” labels applied to operational amplifier input terminals:



The “+” and “-” symbols do not imply polarity of the input voltage(s); i.e. it is not as though the “+” input must be more positive than the “-” input. These symbols merely represent the different direction each input tends to drive the output signal. An increasing potential applied to the “+” input drives the opamp’s output positive, while an increasing potential applied to the “-” input drives the opamp’s output negative. Phrasing this in terms common to closed-loop control systems, we could say that the “+” input is *direct-acting* while the “-” input is *reverse-acting*.

Similarly, the “H” and “L” labels on a DP transmitter’s ports do not imply magnitude of input pressures; i.e. it is not as though the “H” port’s pressure must be greater than the “L” port’s pressure. These symbols merely represent the different effects on the output signal resulting from pressure applied to each port. An increasing pressure applied to the “high” port of a DP transmitter will drive the output signal to a greater level (up), while an increasing pressure applied to the “low” port of a DP transmitter will drive the output signal to a lesser level (down)¹³:



The ability to arbitrarily connect a DP transmitter to a process in such a way that it is either direct-acting or reverse-acting is a great advantage, as we will later see.


In the world of electronics, we refer to the ability of a differential voltage sensor (such as an operational amplifier) to sense small differences in voltage while ignoring large potentials measured with reference to ground by the phrase *common-mode rejection*. An ideal operational amplifier completely ignores the amount of voltage common to both input terminals, responding only to the *difference* in voltage *between* those terminals. This is precisely what a well-designed DP instrument does, except with fluid pressure instead of electrical voltage. A DP instrument ignores gauge pressure common to both ports, while responding only to *differences* in pressure *between* those two ports. Stated in other words, a differential pressure instrument (ideally¹⁴) responds only to differential pressure while ignoring common-mode pressure.

¹³One transmitter manufacturer I am aware of (ABB/Bailey) actually does use the “+” and “-” labels to denote high- and low-pressure ports rather than the more customary “H” and “L” labels found on other manufacturers’ DP products.

¹⁴Perfect common-mode rejection is impossible for differential pressure instruments just as it is impossible for electronic voltage-measuring instruments, but in either case the effect is usually minimal. For differential pressure transmitters, the effect of common-mode pressure on the instrument’s output signal is sometimes referred to as the *line pressure effect* or *static pressure effect*, typically stated as a percentage of the instrument’s upper range limit per unit of common-mode pressure.


To illustrate, we may connect the “high” and “low” ports of a differential pressure transmitter together using pipe or tube, then expose both ports simultaneously to a source of fluid pressure such as pressurized air from an air compressor. If the transmitter is in good working order, it should continue to register zero differential pressure even as we vary the amount of static pressure applied to both ports. So long as the applied pressures to each port are equal, the transmitter’s sensing diaphragm should experience zero net force pushing left or right. All force applied to the diaphragm from the “high” port’s fluid pressure should be precisely countered (canceled) by force applied to the diaphragm from the “low” port’s fluid pressure.

An electrical analogy to this would be connecting both red and black test leads of a voltmeter to a common point in an electrical circuit, then varying the amount of voltage between that point and earth ground. Since the voltmeter only registers *differences* of potential between its test leads, and those test leads are now electrically common to one another, the magnitude of common-mode voltage between that one point of the circuit and earth ground is irrelevant from the perspective of the voltmeter:



In each case the differential measurement device *rejects* the common-mode value, registering only the amount of difference (zero) between its sensing points.

The same common-mode rejection principle reveals itself in more complex fluid and electrical circuits. Consider the case of a DP transmitter and a voltmeter, both used to measure differential quantities in a “divider” circuit¹⁵:



In each case the differential measurement device responds only to the difference between the two measurement points, rejecting the common-mode value (97.5 PSI for the pressure transmitter, 97.5 volts for the voltmeter). Just to make things interesting in this example, the “high” side of each measuring instrument connects to the point of lesser value, such that the measured difference is a negative quantity. Like digital voltmeters, modern DP transmitters are equally capable of accurately measuring negative pressure differences as well as positive pressure differences.

¹⁵The electrical circuit shown on the right uses a pair of series-connected resistors to divide the source voltage into two parts, 5 volts and 95 volts. The pneumatic circuit shown on the left uses a pair of series-connected hand valves to divide the source pressure into two parts, 5 PSI and 95 PSI.

A vivid contrast between *differential* pressure and *common-mode* pressure for a DP instrument is seen in the pressure ratings shown on the nameplate of a Foxboro model 13A differential pressure transmitter:



This nameplate tells us that the transmitter has a calibrated differential pressure range of 50" H₂O (50 inches water column, which is only about 1.8 PSI). However, the nameplate also tells us that the transmitter has a *maximum working pressure* (MWP) of 1500 PSI. "Working pressure" refers to the amount of gauge pressure common to each port, not the differential pressure between ports. Taking these figures at face value means this transmitter will register zero (no differential pressure) even if the gauge pressure applied equally to both ports is a full 1500 PSI! In other words, this differential pressure transmitter will *reject* up to 1500 PSI of common-mode gauge pressure, and respond only to small differences in pressure between the ports (1.8 PSI differential being enough to stimulate the transmitter to full scale output).

19.5.2 DP transmitter applications


The combination of two differential pressure ports makes the DP transmitter very versatile as a pressure-measuring device. This one instrument may be used to measure pressure differences, positive (gauge) pressures, negative (vacuum) pressures, and even absolute pressures, just by connecting the “high” and “low” sensing ports differently.

In every DP transmitter application, there must be some means of connecting the transmitter’s pressure-sensing ports to the points in a process. Metal or plastic tubes (or pipes) work well for this purpose, and are commonly called *impulse lines*, or *gauge lines*, or *sensing lines*¹⁶. This is equivalent to the test wires used to connect a voltmeter to points in a circuit for measuring voltage. Typically, these tubes are connected to the transmitter and to the process by means of *compression fittings* which allow for relatively easy disconnection and reconnection of tubes. For more information on instrument tube fittings, refer to section 8.2.1 beginning on page 580.

¹⁶Also called *impulse tubes*, *gauge tubes*, or *sensing tubes*.

Measuring process vessel clogging


We may use the DP transmitter to measure an actual difference of pressure across a process vessel such as a filter, a heat exchanger, or a chemical reactor. The following illustration shows how a differential pressure transmitter may be used to measure clogging of a water filter:



Note how the high side of the DP transmitter connects to the upstream side of the filter, and the low side of the transmitter to the downstream side of the filter. This way, increased filter clogging will result in an increased transmitter output. Since the transmitter's internal pressure-sensing diaphragm only responds to *differences* in pressure between the "high" and "low" ports, the pressure in the filter and pipe relative to the atmosphere is completely irrelevant to the transmitter's output signal. The filter could be operating at a line pressure of 10 PSI or 10000 PSI – the only variable the DP transmitter measures is the pressure *drop* across the filter. If the upstream side is at 10 PSI and the downstream side is at 9 PSI, the differential pressure will be 1 PSI (sometimes labeled as PSID, "D" for *differential*). If the upstream pressure is 10000 PSI and the downstream pressure is 9999 PSI, the DP transmitter will still see a differential pressure of just 1 PSID. Likewise, the technician calibrating the DP transmitter on the workbench could use a precise air pressure of just 1 PSI (applied to the "high" port, with the "low" port vented to atmosphere) to simulate either of these real-world conditions. The DP transmitter simply cannot tell the difference between these three scenarios, nor should it be able to tell the difference if its purpose is to exclusively measure differential pressure.

Measuring positive gauge pressure

DP instruments may also serve as simple *gauge pressure* instruments if needed, responding to pressures in excess of atmosphere. If we simply connect the “high” side of a DP instrument to a process vessel using an impulse tube, while leaving the “low” side vented to atmosphere, the instrument will interpret any positive pressure in the vessel as a positive *difference* between the vessel and atmosphere:



Although this may seem like a waste of the transmitter’s abilities (why not just use a simpler gauge pressure transmitter with just one port?), it is actually a very common application for DP transmitters. This usage of a differential device may not actually be a “waste” if true-differential applications exist at the same facility for that pressure transmitter, which means only one spare transmitter need be stocked in the facility’s warehouse instead of two spare transmitters (one of each type).

Most DP instrument manufacturers offer “gauge pressure” versions of their differential instruments, with the “high” side port open for connection to an impulse line and the “low” side of the sensing element capped off with a special vented flange, effectively performing the same function we see in the above example at a slightly lesser cost. A close-up photograph of a Rosemount model 1151GP gauge pressure transmitter shows the port-less flange on the “low” side of the pressure-sensing module. Only the “high” side of the sensor has a place for an impulse line to connect:




A closer look at this flange reveals a vent near the bottom, ensuring the “low” side of the pressure-sensing capsule always senses ambient (atmospheric) pressure:



Measuring absolute pressure

Absolute pressure is defined as the difference between a given fluid pressure and a perfect vacuum, as opposed to gauge pressure which is the difference between a fluid's pressure and the atmospheric air pressure. We may build an absolute pressure sensing instrument by taking a DP instrument and sealing the "low" side of its pressure-sensing element in connection to a vacuum chamber. This way, any pressure greater than a perfect vacuum will register as a positive difference:

A vacuum transmitter




Most absolute pressure transmitters resemble "gauge pressure" adaptations of DP transmitters, with only one port available to connect an impulse line. Unlike gauge pressure transmitters, though, absolute pressure transmitters do *not* have vent holes on their "low" sides. The "low" side of an absolute pressure transmitter must be a sealed vacuum in order to accurately measure the "high" side fluid pressure in absolute terms.

Absolute pressure measurement is important for a variety of process applications, including boiling-point control and mass flow measurement of gases. The boiling temperature of any liquid is a function of the absolute pressure it experiences, and in applications where boiling temperature must be precisely controlled in order to achieve a certain outcome (e.g. vacuum distillation of crude oil, for example) the best type of pressure measurement to use absolute. When computing the mass flow rate of gases in a pipe, the relationship between volume and molecular count is a function of both temperature and pressure (both absolute), and so absolute pressure measurement is indispensable here as well.

Measuring vacuum

The same principle of connecting one port of a DP device to a process and venting the other works well as a means of measuring *vacuum* (pressures below that of atmosphere). All we need to do is connect the “low” side to the vacuum process and vent the “high” side to atmosphere:



Any pressure in the process vessel less than atmospheric will register to the DP transmitter as a *positive* difference (with P_{high} greater than P_{low}). Thus, the stronger the vacuum in the process vessel, the greater the signal output by the transmitter.

This last statement deserves some qualification. It used to be, the way analog pneumatic and electronic transmitters were designed many years ago, that the only way to obtain an increasing signal from a DP instrument was to ensure the “high” port pressure *rose* in relation to the “low” port pressure (or conversely stated, to ensure the “low” port pressure *dropped* in relation to the “high” side pressure). However, with the advent of digital electronic technology, it became rather easy to program a DP instrument with a *negative* range, for example 0 to -10 PSI. This way, a *decreasing* pressure as interpreted by the transmitter would yield an *increasing* output signal.

It is rare to find a pressure transmitter calibrated in such a way, but bear in mind that it is possible. This opens the possibility of using a regular “gauge” pressure transmitter (where the

“high” port connects to the process vessel and the “low” port is always vented to atmosphere by virtue of a special flange on the instrument) as a vacuum instrument. If a gauge pressure transmitter is given a negative calibration span, any decreasing pressure seen at the “high” port will yield an increasing output signal.

19.5.3 Inferential measurement applications

A very common technique in industrial instrumentation is to calculate the value of a process variable from the values of related variables which are easier to measure¹⁷. As it so happens, there are a host of variables which one may infer from readings of differential pressure. This makes DP transmitters very versatile devices, not just limited to measuring process variables of pressure and vacuum. This portion of the book will explore some of the more common inferred measurements possible with DP instruments.


¹⁷Truth be told, *most* process variables are inferred rather than directly measured. Even pressure, which is being used here to infer measurements such as liquid level and fluid flow, is itself inferred from some other variable inside the DP instrument (e.g. capacitance, strain gauge resistance, resonant frequency)!

Inferring liquid level

Liquids generate pressure proportional to height (depth) due to their weight. The pressure generated by a vertical column of liquid is proportional to the column height (h), and liquid's mass density (ρ), and the acceleration of gravity (g):

$$P = \rho gh$$

Knowing this, we may use a DP transmitter as a liquid level-sensing device if we know the density of the liquid remains fairly constant¹⁸:




As liquid level in the vessel increases, the amount of hydrostatic pressure applied to the transmitter's "high" port increases in direct proportion. The width of the vessel is irrelevant to the amount of pressure produced – only the liquid *height* (h), *density* (ρ), and Earth's gravity (g) are significant. Thus, the transmitter's increasing signal represents the height of liquid inside the vessel, no matter the size or shape of the vessel:

$$h = \frac{P}{\rho g}$$

¹⁸We simply assume Earth's gravitational acceleration (g) to be constant as well.

This simple technique works even if the vessel is under pressure from a gas or a vapor (rather than being vented as was the case in the previous example). All we need to do to compensate for this other pressure is to connect the DP transmitter's "low" port to the top of the vessel so it senses nothing but the gas pressure:




Since the transmitter responds only to differences of pressure between its two sensing ports, and the only cause for a difference of pressure in this application will be pressure generated by the height of a liquid column, the transmitter's signal becomes an exclusive representation of liquid level in the vessel, rejecting potential measurement errors caused by changes in gas pressure within the vessel. Any gas pressure within the vessel will be sensed equally by both ports on the transmitter as a "common-mode" pressure, thus canceling each other and having no effect on the differential pressure measurement. Only changes in liquid level within the vessel will cause the "high" port pressure to change independently of the "low" port pressure, changing the transmitter's output signal.

Inferring gas and liquid flow

Another common inferential measurement using DP transmitters is the measurement of fluid flow through a pipe. Pressure dropped across a constriction in the pipe varies in relation to flow rate (Q) and fluid density (ρ). So long as fluid density remains fairly constant, we may measure pressure drop across a piping constriction and use that measurement to infer flow rate.

The most common form of constriction used for this purpose is called an *orifice plate*, being nothing more than a metal plate with a precisely machined hole in the center. As fluid passes through this hole, its velocity changes, causing a pressure drop to form:



Once again, we see the common-mode rejection abilities of the pressure transmitter used for practical advantage. Since both ports of the transmitter connect to the same process line, static fluid pressure within that line has no effect on the measurement. Only *differences* of pressure between the upstream and downstream sides of the constriction (orifice plate) cause the transmitter to register flow.


19.6 Pressure sensor accessories

Multiple accessories exist for pressure-sensing devices to function optimally in challenging process environments. Sometimes, we must use special accessories to protect the pressure instrument against hazards of certain process fluids. One such hazard is pressure *pulsation*, for example at the discharge of a piston-type (positive-displacement) high-pressure pump. Pulsating pressure can quickly damage mechanical sensors such as bourdon tubes, either by wear of the mechanism transferring pressure element motion to an indicating needle, and/or fatigue of the metal element itself.

19.6.1 Valve manifolds


An important accessory to the DP transmitter is the *valve manifold*. This device incorporates manual valves to isolate and equalize pressure from the process to the transmitter, for maintenance and calibration purposes.

The following illustration shows the three valves comprising a three-valve manifold (within the dotted-line box), as well as a fourth valve called a “bleed” valve used to vent trapped fluid pressure to atmosphere:



While this illustration shows the three valves as separate devices, connected together and to the transmitter by tubing, three-valve manifolds are more commonly manufactured as monolithic devices: the three valves cast together into one block of metal, attaching to the pressure transmitter by way of a flanged face with O-ring seals. Bleed valves are most commonly found as separate devices threaded into one or more of the ports on the transmitter’s diaphragm chambers.

The following photograph shows a three-valve manifold bolted to a Honeywell model ST3000 differential pressure transmitter. A bleed valve fitting may be seen inserted into the upper port on the nearest diaphragm capsule flange:




In normal operation, the two block valves are left open to allow process fluid pressure to reach the transmitter. The equalizing valve is left tightly shut so no fluid can pass between the “high” and “low” pressure sides. To isolate the transmitter from the process for maintenance, one must close the block valves and open the equalizing valve. The best sequence to follow is to first close the high-pressure block valve, then open the equalizing valve, then close the low-pressure block valve. This sequence ensures the transmitter cannot be exposed to a high differential pressure during the isolation procedure, and that the trapped fluid pressure inside the transmitter will be as low as possible prior to “venting” to atmosphere. Finally, the “bleed” valve is opened at the very last step to relieve pent-up fluid pressure within the manifold and transmitter chambers¹⁹:


¹⁹To return the transmitter to live service, simply reverse these steps: close the bleed valve, open the low-pressure block valve, close the equalizing valve, and finally open the high-pressure block valve.

Final valve positions for both states are shown in the following illustrations:


Normal operation



Removed from service




A variation on this theme is the *five-valve manifold*, shown in this illustration:




The presence of a built-in bleed valve in the five-valve manifold allows the technician to vent trapped pressure through a tube to some remote location, rather than directly venting at the transmitter. Valve positions for normal operation and maintenance on this manifold are as follows:

Normal operation




Removed from service



It is critically important that the equalizing valve(s) never be open on any transmitter manifold while both block valves are open! Doing so will allow process fluid to flow through the equalizing valve(s) from the high-pressure side of the process to the low-pressure side of the process. If the impulse tubes connecting the manifold to the process are intentionally filled with a *fill fluid* (such as glycerin, to displace process water from entering the impulse tubes; or water in a steam system), this fill fluid will be lost. Also, if the process fluid is dangerously hot or radioactive, a combination of open equalizing and block valves will let that dangerous fluid reach the transmitter and manifold, possibly causing damage or creating a personal hazard. Speaking from personal experience, I once

made this mistake on a DP transmitter connected to a steam system, causing hot steam to flow through the manifold and overheat the equalizing valve so that it seized open and could not be shut again! The only way I was able to stop the flow of hot steam through the manifold was to locate and shut a sliding-gate hand valve between the impulse tube and the process pipe. Fortunately, this cast steel valve was not damaged by the heat and was still able to shut off the flow.

Pressure transmitter valve manifolds also come in single block-and-bleed configurations, for gauge pressure applications. Here, the “low” pressure port of the transmitter is vented to atmosphere, with only the “high” pressure port connected to the impulse line:



The following photograph shows a bank of eight pressure transmitters, seven out of the eight being equipped with a single block-and-bleed manifold. The eighth transmitter (bottom row, second-from left) sports a 5-valve manifold:





If you look closely at the photograph, you can see the bleed valve fittings installed on all the upper ports. Only the transmitter with the 5-valve manifold has *two* bleed valve fittings because it is the only DP transmitter of the group. The other seven transmitters are all *gauge pressure* units, and so only have one port to bleed.

A good habit to cultivate when operating valve handles on transmitter manifolds is to “back off” the open valves approximately one-quarter turn after opening. This discourages seizing in the full-open position, and also makes it possible for someone to more easily tell the states of the valves by feel: a closed valve will not easily turn (because it is tightened onto its seat) while an open valve is free to turn either direction a bit. Since there should be no flow going through the valves of a transmitter manifold, it is irrelevant whether an open manifold valve is 100% open or 90% open or 80% open, so there is no harm in “backing off” an open valve from the full-open position. It would of course be bad to do this with a closed valve, since any valve plug must be pressed tight into its seat in order to achieve positive shut-off.


19.6.2 Bleed (vent) fittings

Before removing a pressure transmitter from live service, the technician must “bleed” or “vent” accumulated fluid pressure to atmosphere in order to achieve a *zero energy state* prior to disconnecting the transmitter from the impulse lines. Some valve manifolds provide a bleed valve for doing just this, but many do not²⁰. An inexpensive and common accessory for pressure-sensing instruments (especially transmitters) is the *bleed valve fitting* or *vent valve fitting*, installed on the instrument as a discrete device. The most common bleed fitting is equipped with 1/4 inch male NPT pipe threads, for installation into one of the 1/4 inch female NPT pipe ports typically provided on pressure transmitter flanges. The bleed fitting is operated with a small wrench, loosening a ball-tipped plug off its seat to allow process fluid to escape through a small vent hole in the side of the fitting. The following photographs show close-up views of a bleed fitting both assembled (left) and with the plug fully extracted from the fitting (right). The bleed hole may be clearly seen in both photographs:



When installed directly on the flanges of a pressure instrument, these bleed valves may be used to bleed unwanted fluids from the pressure chambers, for example bleeding air bubbles from an instrument intended to sense water pressure, or bleeding condensed water out of an instrument intended to sense compressed air pressure.

The following photographs show bleed fittings installed two different ways on the side of a pressure transmitter flange, one way to bleed gas out of a liquid process (located on top) and the other way to bleed liquid out of a gas process (located on bottom):




²⁰The standard 3-valve manifold, for instance, does not provide a bleed valve – only block and equalizing valves.

With the bleed plug completely removed, the open bleed fitting provides a port through which one may apply air pressure for testing the response of the pressure transmitter. A special test fitting called a *bleed port adapter* or *DP transmitter calibration fitting* – colloquially known as a *stinger* – threads into the opened bleed fitting. A photograph of a bleed port adapter is shown here:



This special fitting allows a compression-style tube to be temporarily connected to the opened bleed port, which then allows the connection of an air pump and test pressure gauge to the transmitter. Thus, the bleed port adapter enables a technician to conveniently apply test pressures to the DP transmitter without having to loosen any of the instrument manifold bolts, tapered thread pipe connections, or impulse tube compression fittings.

When performing field checks of pressure transmitters, bleed port adapters substantially reduce the amount of time necessary to field-test pressure instruments. The following sequence of illustrations show how a bleed port adapter may be used in conjunction with a three-valve instrument manifold to isolate a DP transmitter from a process and then subject it to test pressures from a hand pump:



Note how both bleed vents must be opened, and the equalizing valve shut, in order to apply a test pressure to the DP transmitter. Although it is possible to safely bleed pressure from both sides of a DP instrument through just one bleed fitting (through the open equalizing valve), both bleeds must be open in order to perform a pressure test. If the "L" side bleed fitting is left in the shut position, some pressure may be trapped there as pressure is applied to the "H" side by the hand pump. If the equalizing valve is left open, no difference of pressure will be allowed to form across the DP instrument.


19.6.3 Pressure pulsation damping

A simple way to mitigate the effects of pulsation on a pressure gauge is to fill the inside of the gauge with a viscous liquid such as glycerin or oil. The inherent friction of this fill liquid has a “shock-absorber” quality which damps the gauge mechanism’s oscillatory motion and helps protect against damage from pulsations or from external vibration. This method is ineffectual for high-amplitude pulsations, though.

An oil-filled pressure gauge may be seen in the following photograph. Note the air bubble near the top of the gauge face, which is the only visual indication of an oil filling:




A more sophisticated method for damping pulsations seen by a pressure instrument is called a *snubber*, and it consists of a fluid restriction placed between with the pressure sensor and the process. The simplest example of a snubber is a simple *needle valve* (an adjustable valve designed for low flow rates) placed in a mid-open position, restricting fluid flow in and out of a pressure gauge:




At first, the placement of a throttling valve between the process and a pressure-measuring instrument seems rather strange, because there should not be any continuous flow in or out of the gauge for such a valve to throttle! However, a *pulsing* pressure causes a small amount of *alternating* flow in and out of the pressure instrument, owing to the expansion and contraction of the mechanical pressure-sensing element (bellows, diaphragm, or bourdon tube). The needle valve provides a restriction for this flow which, when combined with the fluid capacitance of the pressure

instrument, combine to form a low-pass filter of sorts. By impeding the flow of fluid in and out of the pressure instrument, that instrument is prevented from “seeing” the high and low peaks of the pulsating pressure. Instead, the instrument registers a much steadier pressure over time. An electrical analogy for a pressure snubber is an RC low-pass filter circuit “damping” voltage pulsations from reaching a DC voltmeter:



One potential problem with the needle valve solution is that the small orifice inside the needle valve may plug up over time with debris from dirty process fluid. This, of course, would be bad because plugging will cause the pressure instrument to respond too slowly, or not at all if the plugging is complete.

A solution to this problem is to fill the pressure sensor mechanism with a clean liquid (called a *fill fluid*) and use that fill fluid to transfer pressure from the process fluid to the pressure-sensing element using a slack diaphragm or some other membrane separating the process fluid from the fill fluid:




It should be noted that most pressure snubbers utilize a fixed-geometry orifice rather than an adjustable needle valve to dampen pressure pulsations seen at the pressure gauge.

In order for the fill fluid and isolating diaphragm to work effectively, there cannot be any gas bubbles in the fill fluid – it must be a “solid” hydraulic system from the diaphragm to the sensing element. Gas bubbles present in the filled system would make that volume compressible, which means the isolating diaphragm would have to move more than necessary to transfer pressure to the instrument’s sensing element. This would mean motion at the isolating diaphragm caused by process pressure changes would be “lost” and not fully transferred to the instrument’s sensing element, thereby introducing a pressure measurement error²¹. For this reason, isolating diaphragm systems for pressure instruments are usually “packed” with fill fluid at the point and time of manufacture, then sealed in such a way that they cannot be opened for any form of maintenance. Consequently, any fill fluid leak in such a system immediately ruins it.

²¹This concept will be immediately familiar to anyone who has ever had to “bleed” air bubbles out of an automobile brake system. With air bubbles in the system, the brake pedal has a “spongy” feel when depressed, and much pedal motion is required to achieve adequate braking force. After bleeding all air out of the brake fluid tubes, the pedal motion feels much more “solid” than before, with minimal motion required to achieve adequate braking force. Imagine the brake pedal being the isolating diaphragm, and the brake pads being the pressure sensing element inside the instrument. If enough gas bubbles exist in the tubes, the brake pedal might stop against the floor when fully pressed, preventing full force from ever reaching the brake pads. Likewise, if the isolating diaphragm hits a hard motion limit due to gas bubbles in the fill fluid, the sensing element will not experience full process pressure.


19.6.4 Remote and chemical seals

Isolating diaphragms have merit even in scenarios where pressure pulsations are not a problem. Consider the case of a food-processing system where we must remotely measure pressure inside a mixing vessel:



The presence of the tube connecting the vessel to the pressure gauge poses a hygiene problem. Stagnant process fluid (in this case, some liquid food product) inside the tube will encourage microbial growth, which will eventually contaminate the vessel no matter how well or how often the vessel is cleaned. Even automated *Clean-In-Place* and *Steam-In-Place* (*CIP* and *SIP*, respectively) protocols where the vessel is chemically purged between batches cannot prevent this problem because the cleaning agents never purge the entire length of the tubing (ultimately, to the bourdon tube or other sensing element inside the gauge).

A solution to this problem is to install an *isolating diaphragm* at the vessel, and a liquid-filled *capillary tube* to transfer sensed pressure to the instrument. Process pressure presses against this diaphragm, which in turn transfers²² pressure to the “fill fluid” inside the capillary tube. This sealed fill fluid then presses against the instrument’s sensing element (diaphragm, bourdon tube, bellows, etc.). Process (food) liquid cannot enter this sealed tube, and the isolating diaphragm will be cleaned with every CIP cycle. Thus, we completely eliminate the problem of microbial contamination:



Such systems are often referred to as *remote seals*, and they are available on a number of different pressure instruments including gauges, transmitters, and switches. If the purpose of an isolating diaphragm and fill fluid is to protect the sensitive instrument from corrosive or otherwise harsh chemicals, it is often referred to as a *chemical seal*.

²²So long as the isolating diaphragm is “slack” (i.e. has no appreciable tautness or resistance to movement), the pressure of the fill fluid inside the capillary tube *will* be equal to the pressure of whatever fluid is within the process vessel. If any pressure imbalance were to develop between the process and fill fluids, the isolating diaphragm would immediately shift position away from the higher-pressure fluid and toward the lower-pressure fluid until equal pressures were re-established. In real practice, isolating diaphragms do indeed have some stiffness opposing motion, and therefore do not *perfectly* transfer pressure from the process fluid to the fill fluid. However, this pressure difference is usually negligible.


The following photograph shows a pressure gauge equipped with a chemical seal diaphragm. Note that the chemical seal on this particular gauge is close-coupled to the gauge, since the only goal here is protection of the gauge from harsh process fluids, not the ability to remotely mount the gauge:




A view facing the bottom of the flange reveals the thin metal isolating diaphragm keeping process fluid from entering the gauge mechanism. Only inert fill fluid occupies the space between this diaphragm and the gauge's bourdon tube:



The following illustration shows how the fill fluid transfers process fluid pressure to the gauge's bourdon tube element while isolating that bourdon tube from the process fluid (shown here inside a pipe):



The only difference between this chemical-seal gauge and a remote-seal gauge is the small-diameter *capillary tubing* used to connect the gauge to a remote diaphragm. An illustration showing the internals of a remote seal system appears here:




Direct-reading gauges are not the only type of pressure instruments benefiting from remote seals. Pressure switches and pressure transmitters may also employ remote seals for the same reasons: protection of the transmitter sensor from harsh process fluid, elimination of impulse tube clogging problems, and/or the prevention of “dead-end” tube lengths where organic process fluid would stagnate and harbor microbial growths. In this photograph, you see three pressure sensing devices (gauge, transmitter, and switch, from top to bottom), each one with its own remote seal to sense fluid pressure in a large pipe:




The pressure transmitter (a Yokogawa unit) is the only instrument shown here using a capillary tube. The other two instruments use short lengths of rigid pipe. The capillary is visible in this photograph as a coiled tube covered in black plastic. It is actually a very small-diameter metal tube enclosed in a spiral-metal protective sheath which is in turn covered by black plastic.

This particular application happens to be in wastewater treatment, where sludge has a tendency to clog instrument impulse lines connected to the main piping. With remote seals in place, that problem is completely eliminated.

The following photograph shows a Rosemount model 1151 electronic pressure transmitter equipped with a remote sealing diaphragm. Here we may see the coiled metal (“armor”) sheath protecting the capillary tube from damage:



A close-up view of the sealing diaphragm shows its corrugated design, allowing the metal to easily flex and transfer pressure to the fill fluid within the capillary tubing²³:



As with the isolating diaphragms of the pressure-sensing capsule, these remote diaphragms need only transfer process fluid pressure to the fill fluid and (ultimately) to the taut sensing element inside the instrument. Therefore, these diaphragms perform their function best if they are designed to easily flex. This allows the sensing diaphragm to provide the vast majority of the opposing force to the fluid pressure, as though it were the only spring element in the fluid system.

²³Like all instrument diaphragms, this one is sensitive to damage from contact with sharp objects. If the diaphragm ever becomes nicked, dented, or creased, it will tend to exhibit hysteresis in its motion, causing calibration errors for the instrument. For this reason, isolating diaphragms are often protected from contact by a plastic plug when the instrument is shipped from the manufacturer. This plug must be removed from the instrument before placing it into service.

The connection point between the capillary tube and the transmitter's sensor capsule is labeled with a warning never to disassemble, since doing so would allow air to enter the filled system (or fill fluid to escape from the system) and thereby ruin its accuracy:




In order for a remote seal system to work, the hydraulic “connection” between the sealing diaphragm and the pressure-sensing element must be completely gas-free so there will be a “solid” transfer of motion from one end to the other²⁴. The presence of gas bubbles in the filled system will cause some of the isolating diaphragms’ motion to be “lost” rather than be fully transferred to the instrument’s pressure-sensing element. For this reason, the capillary system must remain perfectly sealed at all times! Breaching this seal, even for just a brief moment, will ruin the system.

A protective measure visible in this photograph is the orange compound painted on the screw head and on the capillary tube connector. This is simply a visual indication that the factory seal is still intact, since any motion of the screw or of the tube connector would crack the brittle orange compound and betray the breach.

²⁴Anyone familiar with “bleeding” air bubbles out of automotive hydraulic brake systems will understand this concept. In order for the pedal-operated hydraulic brakes in an automobile to function as designed, the hydraulic system must be gas-free. Incompressible liquid transfers pressure without loss of motion, whereas compressible gas bubbles will “give” in to pressure and result in lost brake pad motion for any given brake pedal motion. Thus, an hydraulic brake system with air bubbles in it will have a “spongy” feel at the brake pedal, and may not give full braking force when needed.

A potential problem with using remote diaphragms is the hydrostatic pressure generated by the fill fluid if the pressure instrument is located far away (vertically) from the process connection point. For example, a pressure gauge located far below the vessel it connects to will register a *greater* pressure than what is actually inside the vessel, because the vessel's pressure adds to the hydrostatic pressure caused by the liquid in the tubing:



This pressure may be calculated by the formula $P = \rho g h$ or $P = \gamma h$ where ρ is the mass density of the fill liquid or γ is the weight density of the fill liquid. For example, a 12 foot capillary tube height filled with a fill liquid having a weight density of 58.3 lb/ft³ will generate an elevation pressure of almost 700 lb/ft², or 4.86 PSI. If the pressure instrument is located below the process connection point, this 4.86 PSI offset must be incorporated into the instrument's calibration range. If we desire this pressure instrument to accurately measure a process pressure range of 0 to 50 PSI, we would have to calibrate it for an actual range of 4.86 to 54.86 PSI.

The reverse problem exists where the pressure instrument is located *higher* than the process connection: here the instrument will register a *lower* pressure than what is actually inside the vessel, offset by the amount predicted by the hydrostatic pressure formulae $\rho g h$ or γh .

In all fairness, this problem is not limited to remote seal systems – even non-isolated systems


where the tubing is filled with process liquid will exhibit this offset error. However, in filled-capillary systems a vertical offset is *guaranteed* to produce a pressure offset because fill fluids are always liquid, and liquids generate pressure in direct proportion to the vertical height of the liquid column and its density. Many common remote seal fill fluids have *specific gravity* ratings greater than 1 (i.e. they are denser than water) which means the offset error resulting from hydrostatic pressure will be even greater than that of a water-filled tube.

A similar problem unique to isolated-fill pressure instruments is measurement error caused by temperature extremes. Suppose the liquid-filled capillary tube of a remote seal pressure instrument is placed too near a hot steam pipe, furnace, or some other source of high temperature. The expansion of the fill fluid may cause the isolation diaphragm to extend to the point where it begins to tense and add a pressure to the fill fluid above and beyond that of the process fluid. Cold temperatures may wreak havoc with filled capillary tubes as well, if the fill fluid congeals or even freezes such that it no longer flows as it should. Fill fluid expansion/contraction effects may be mitigated by keeping the volume of the fill fluid to a minimum, which is why capillary (small-diameter) tubes are used to connect remote seals with instruments.

Another problem with remote seal pressure instruments is a time delay caused by the viscosity of the fill fluid as it moves through the small-diameter capillary tubes. This makes the pressure instrument slow to respond to changes in process fluid pressure, as though the pressure instrument were equipped with an (undesired) pressure snubber. While minimal capillary tube diameter reduces the effects of temperature changes, it increases the effect of time delay.

Proper mounting of the instrument and proper selection of the fill fluid²⁵ will help to avoid such problems. All in all, the potential for trouble with remote- and chemical-seal pressure instruments is greatly offset by their benefits in the right applications.


Some diaphragm-sealed pressure transmitters are equipped with close-coupled seals rather than remote seals to minimize hydrostatic, temperature, and time delay effects caused by fill fluid inside a capillary tube. A Rosemount extended-diaphragm pressure transmitter appears in the left-hand photograph, while a Yokogawa transmitter of the same basic design is shown installed in a working process in the right-hand photograph:



²⁵Most pressure instrument manufacturers offer a range of fill fluids for different applications. Not only is temperature a consideration in the selection of the right fill fluid, but also potential contamination of or reaction with the process if the isolating diaphragm ever suffers a leak!

19.6.5 Filled impulse lines

An alternate method for isolating a pressure-sensing instrument from direct contact with process fluid is to *fill* the impulse lines with a harmless fluid, which in turn directly contacts the process fluid. Filling impulse tubes with a static fluid works when gravity is able to keep the fill fluid in place, such as in this example of a pressure transmitter connected to a water pipe by a glycerin-filled impulse line:




A reason someone might do this is for freeze protection, since glycerin freezes at a lower temperature than water. If the impulse line were filled with process water, it might freeze solid in cold weather conditions (the water in the pipe cannot freeze so long as it is forced to flow). The greater density of glycerin keeps it placed in the impulse line, below the process water line. A fill valve is provided near the transmitter so a technician may re-fill the impulse line with glycerin (using a hand pump) if ever needed.

As with a remote diaphragm, a filled impulse line will generate its own pressure proportional to the height difference between the point of process connection and the pressure-sensing element. If the height difference is substantial, the pressure offset resulting from this difference in elevation will require compensation by means of an intentional “zero shift” of the pressure instrument when it is calibrated.

With no isolating diaphragm to separate process fluid from the fill fluid, it is critical that the fill fluid be compatible²⁶ with the process fluid. Not only does this imply a total lack of chemical reactivity between the two fluids, but it also means the two fluids should not be readily *miscible* (capable of mixing in any proportion).

²⁶Truth be told, this is a requirement for all pressure transmitter fill fluids even when isolating diaphragms are in place to prevent mixing of process and fill fluids, because no diaphragm is 100% guaranteed to seal forever. This means *every* pressure transmitter must be chosen for the application in mind, since modern DP transmitters all use fill fluid in their internal sensors, whether or not the impulse lines are also filled with a non-reactive fluid.

An important consideration in filled-line systems is how to refill the impulse line(s) without damaging the pressure instrument. Hand-operated pumps are commonly used to refill impulse lines, but such pumps are often capable of generating greater fluid pressures than the instrument can safely withstand. If we were to connect a glycerin pump to the filled system pictured previously, it would be advisable to shut the transmitter's block valve to ensure we did not accidentally over-pressure the transmitter. This is especially true if the impulse line happens to become plugged with debris, and substantial glycerin pressure from the hand pump is required to dislodge the blockage:




In fact, the issue of impulse tube plugging is another reason to consider filled-line connections between pressure instruments and process lines or vessels. If ever a plug develops in the line, re-pumping the lines with fresh fill fluid is a practical way to clear the plug without disassembling any part of the system.

For processes where impulse line plugging is a chronic problem, another solution exists called *purging* impulse lines, discussed in the next section.

19.6.6 Purged impulse lines

Another method for isolating a pressure instrument from direct contact with process fluid, particularly when the impulse line is prone to plugging, is *purging* the line with a continuous flow of clean fluid. Consider this example, where pressure is measured at the bottom of a sedimentation vessel:



In this system, a continuous flow of clean water enters through a pressure-dropping “purge valve” and flows through the impulse line, keeping it clear of sediment while still allowing the pressure instrument to sense pressure at the bottom of the vessel. A *check valve* guards against reverse flow through the purge line, in case process fluid pressure ever exceeds purge supply pressure. The continuous water purge maintains clean impulse tubing, and ensures the pressure transmitter never contacts process fluid directly.

A very important element of any purge system is a *restriction* between the purge supply and the connection with the process and pressure-sensing device. It is important that the pressure-sensing instrument senses the pressure of the process fluid and not the (higher) pressure of the purge fluid supply. In the example shown, the purge valve fulfills the role of this restriction, which is why it must be left in a partially-open condition, rather than fully-open.

If this purge restriction is not restrictive enough, the purge fluid flow rate will be too great, resulting in a dynamic pressure drop developed across the length of the impulse line. This pressure drop will add to the pressure of the process fluid inside the vessel, creating a positive pressure measurement error at the instrument (i.e. the instrument will register *more* pressure than there actually is in the vessel). The purge restriction should be set to allow just enough purge fluid flow to guard against plugging, and no more.

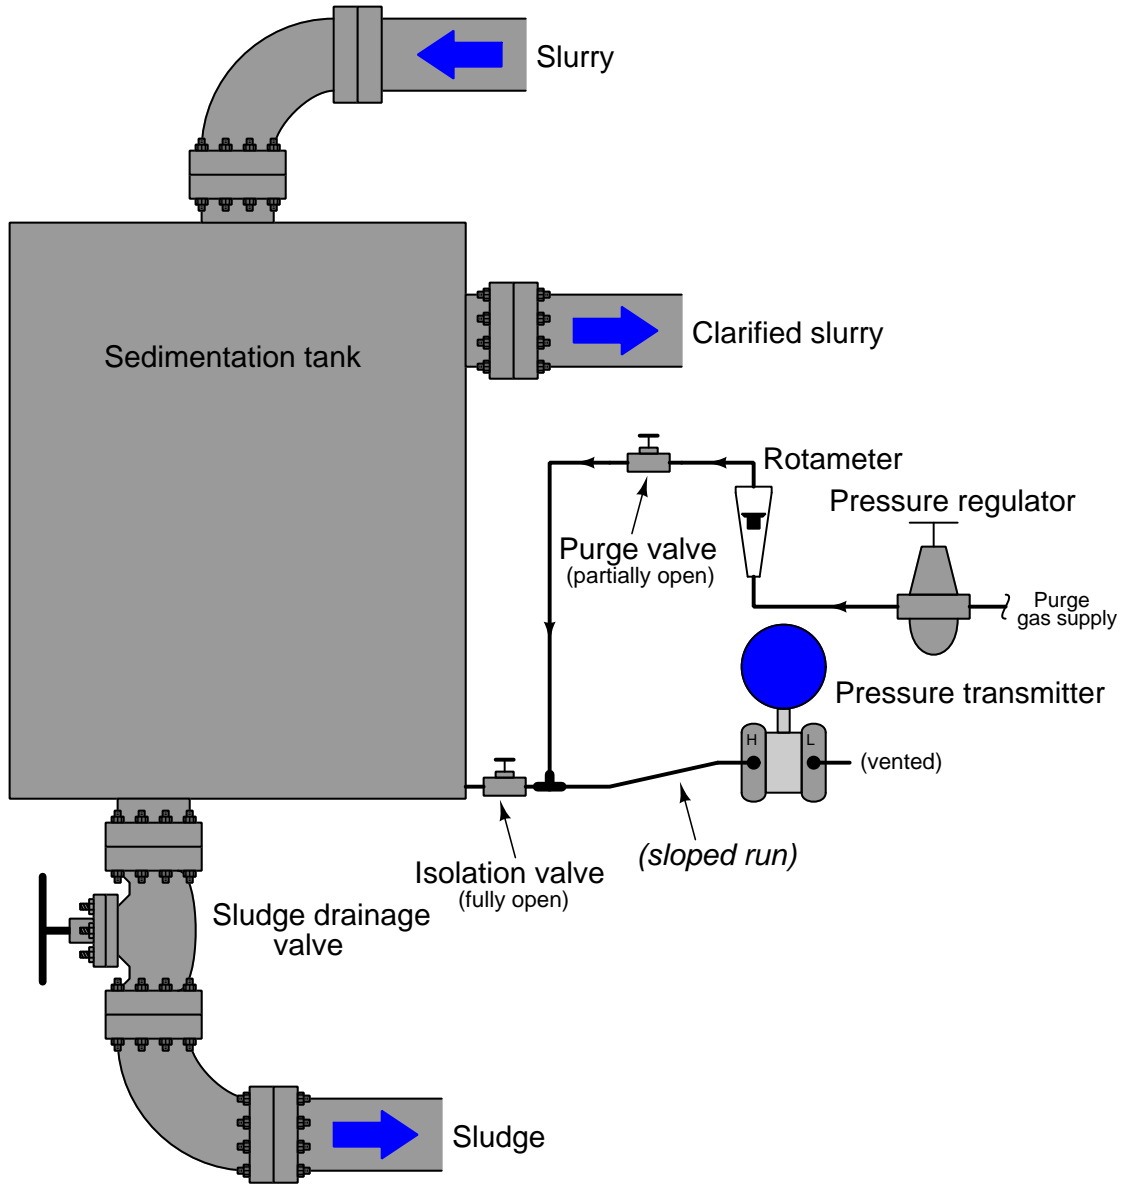
Purged systems are very useful, but a few requirements are necessary in order to ensure accurate and reliable operation:

- The purge fluid supply must be reliable: if the flow stops for any reason, the impulse line may plug!
- The purge fluid supply pressure must exceed the process pressure at all times, or else process fluid will flow backward *into* the impulse line!
- The purge fluid flow must be maintained at a low rate, to avoid pressure measurement errors.
- The purge fluid should be introduced into the impulse line as close to the process connection as possible, to minimize errors due to the purge flow rate through long lengths of tubing.
- The purge fluid must not adversely react with the process.
- The purge fluid must not contaminate the process.
- The purge fluid must be reasonable in cost, since it will be continuously consumed over time.


A useful accessory to include in any purge system is a visual flow indicator such as a *rotameter*. Such an indicator is useful for manual adjustment of purge flow rate, and also as a troubleshooting aid, to indicate if anything happens to halt the purge flow.

In the previous example, the purge fluid was clean water. Many options exist for purge fluids other than water, though. Gases such as air, nitrogen, or carbon dioxide are often used in purged systems, for both gas and liquid process applications.

An illustration of a gas-purged pressure measurement system is shown here:



An alternative to the pressure regulator, rotameter, and purge valve is a self-contained unit called a *purge flow regulator* which automatically adjusts to maintain a constant flow rate of purge gas into the purged impulse line:




It should be noted that liquid-purged impulse lines – just like filled lines and diaphragm-isolated lines – will generate hydrostatic pressure with vertical height. This is not a problem with gas-purged lines.

19.6.7 Heat-traced impulse lines

If impulse lines are filled with liquid, there may exist a possibility for that liquid to freeze in cold-weather conditions. This possibility depends, of course, on the type of liquid filling the impulse lines and how cold the weather gets in that geographic location.

One safeguard against impulse line freezing is to *trace* the impulse lines with some form of active heating medium, steam and electrical being the most common. “Steam tracing” consists of a copper tube carrying low-pressure steam, bundled alongside one or more impulse tubes, enclosed in a thermally insulating jacket.



Steam flows through the shut-off valve, through the tube in the insulated bundle, transferring heat to the impulse tube as it flows past. Cooled steam condenses into water and collects in the *steam trap* device located at the lowest elevation on the steam trace line. When the water level builds up to a certain level inside the trap, a float-operated valve opens to vent the water. This allows more steam to flow into the tracing tube, keeping the impulse line continually heated.

The steam trap naturally acts as a sort of thermostat as well, even though it only senses condensed water level and not temperature. The rate at which steam condenses into water depends on how cold the impulse tube is. The colder the impulse tube (caused by colder ambient conditions), the more heat energy drawn from the steam, and consequently the faster condensation rate of steam into water. This means water will accumulate faster in the steam trap, which means it will “blow down” more often. More frequent blow-down events means a greater flow rate of steam into the tracing tube, which adds more heat to the tubing bundle and raises its temperature. Thus, the system is naturally regulating, with its own negative feedback loop to maintain bundle temperature at a relatively stable point²⁷.

²⁷In fact, after you become accustomed to the regular “popping” and “hissing” sounds of steam traps blowing down, you can interpret the blow-down frequency as a crude ambient temperature thermometer! Steam traps seldom

The following photograph shows a picture of a steam trap:



Steam traps are not infallible, being susceptible to freezing (in *very* cold weather) and sticking open (wasting steam by venting it directly to atmosphere). However, they are generally reliable devices, capable of adding tremendous amounts of heat to impulse tubing for protection against freezing.

blow down during warm weather, but their “popping” is much more regular (one every minute or less) when ambient temperatures drop well below the freezing point of water.

Electrically traced impulse lines are an alternative solution for cold-weather problems. The “tracing” used is a twin-wire cable (sometimes called *heat tape*) that acts as a resistive heater. When power is applied, the cable heats up, thus imparting thermal energy to the impulse tubing it is bundled with. This next photograph shows the end of a section of electrical heat tape, rated at 33 watts per meter (10 watts per foot) at 10 degrees Celsius (50 degrees Fahrenheit):




This particular heat tape also has a maximum current rating of 20 amps (at 120 volts). Since heat tape is really just a continuous parallel circuit, longer lengths of it draw greater current. This maximum total current rating therefore places a limit on the usable length of the tape.

Heat tape may be self-regulating, or controlled with an external thermostat. Self-regulating heat tape exhibits an electrical resistance that varies with temperature, automatically self-regulating its own temperature without the need for external controls.

Both steam and electrical heat tracing are used to protect instruments themselves from cold weather freezing, not just the impulse lines. In these applications it is important to remember that only the liquid-filled portions of the instrument require freeze protection, not the electronic portions!

19.6.8 Self-purged impulse lines

One of the advantages of purging a pressure instrument impulse line with gas rather than with liquid is the elimination of measurement error due to the pressure generated by a vertical column of liquid. We investigated an example of this phenomenon in the “Remote and Chemical Seals” subsection where a liquid-filled vertical run of capillary tube 12 feet high created an additional hydrostatic pressure of 4.86 PSI sensed by the pressure gauge. This hydrostatic pressure caused the gauge to read falsely high, so that instead of registering 0 to 50 PSI as the process vessel pressure ranges from 0 to 50 PSI, the gauge instead reads 4.86 PSI *extra* at all points: reading 4.86 to 54.86 PSI as the process vessel pressure goes from 0 to 50 PSI:




This measurement error may be compensated by shifting the “zero” calibration of the pressure gauge by 4.86 PSI, forcing it to register 4.86 PSI *less* than the pressure it senses at its input port. Only by custom-calibrating the pressure gauge in this manner can we solve the problem created by that 4.86 PSI hydrostatic pressure ($P_{elevation}$).

We also learned in the “Filled Impulse Lines” subsection that this hydrostatic effect is not limited to remote-seal capillary systems but is endemic to *any* significant vertical length of tube filled with

any liquid. A gas-purged impulse line, by contrast, generates negligible pressure due to elevation differences simply because the density of most gases is negligibly small.


An interesting variation on the theme of gas-purging for instrument impulse lines is the use of an external heat source on those impulse lines to cause the process liquid to *boil* and vaporize within the lines. This technique, of course, only works for process liquids that are easily vaporized with modest applications of heat, but in many processes this is practical. Examples of process liquids amenable to this treatment are propane, butane, and any cryogenic²⁸ liquid. If we use heat-traced impulse lines, the thermal energy added to the lines maintains their interiors in a gaseous state rather than a liquid state, eliminating any vertical liquid columns inside the lines and therefore eliminating any pressure measurement error resulting from elevation between the pressure instrument and the process connection point.

Self-purging works best in installations where the pressure-sensing instrument is mounted *above* the process liquid level, and where the presence of liquid inside the vertical run of impulse line extending down from the pressure instrument to the vessel would otherwise create a *negative* pressure offset. Shown here is an example of an installation where a vertical impulse line creates a negative pressure measurement error:



²⁸“Cryogenic” simply refers to a condition of extremely low temperature required to condense a gas into liquid. Such liquids will flash into vapor if raised to room temperature, and so it is quite easy to make impulse lines self-purging in such cases.

Heating this impulse line causes any liquid inside of it to vaporize, forcing remaining liquid to flow out the bottom of the line and back into the process vessel, “self-purging” the line with vapor and thereby ensuring the pressure instrument senses actual process vessel pressure:



The ideal impulse line heat-trace temperature is greater than the *critical temperature* of the process fluid, so that no amount of process pressure can make it liquefy. This will ensure a liquid-free state inside the heated impulse lines even if process pressure happens to increase.

If any *other* liquids exist inside the vessel that will not vaporize at the same temperature, it becomes necessary to install a liquid *trap* at the bottom of the impulse line where this other liquid can accumulate without filling up the impulse line. Periodically draining this trap of accumulated liquid then becomes a regular maintenance task.


Self-purging does not work as well in installations where the process vessel is located above of the pressure-sensing instrument because liquid will continuously find its way into impulse line by gravity, where the sudden expansion from liquid into vapor will create pressure surges inside of the line. This will cause the pressure instrument to register intermittent “surges” of pressure, which is a worse problem²⁹ than having an hydrostatic offset.

²⁹At least in the case of a liquid-filled impulse line generating its own hydrostatic pressure, that pressure is constant

19.6.9 Water traps and pigtail siphons

Many industrial processes utilize high-pressure steam for direct heating, performing mechanical work, combustion control, and as a chemical reactant. Measuring the pressure of steam is important both for its end-point use and its generation (in a boiler). One problem with doing this is the relatively high temperature of steam at the pressures common in industry, which can cause damage to the sensing element of a pressure instrument if directly connected.

A simple yet effective solution to this problem is to intentionally create a “low” spot in the impulse line where condensed steam (water) will accumulate and act as a liquid barrier to prevent hot steam from reaching the pressure instrument. The principle is much the same as a plumber’s trap used underneath sinks, creating a liquid seal to prevent noxious gases from entering a home from the sewer system. A loop of tube or pipe called a *pigtail siphon* achieves the same purpose:



and may be compensated by “zero-shifting” the range of the pressure instrument. An impulse line that generates random *surges* of pressure cannot be compensated at all!


The following photograph shows a pigtail siphon connected to a pressure gauge sensing pressure on a steam line:



19.6.10 Mounting brackets

An accessory specifically designed for a variety of field-mounted instruments including DP transmitters is the *2 inch pipe mounting bracket*. Such a bracket is manufactured from heavy-gauge sheet metal and equipped with a U-bolt designed to clamp around any 2 inch black iron pipe. Holes stamped in the bracket match mounting bolts on the capsule flanges of most common DP transmitters, providing a mechanically stable means of attaching a DP transmitter to a framework in a process area.

The following photographs show several different instruments mounted to pipe sections using these brackets:



19.6.11 Heated enclosures

In installations where the ambient temperature may become very cold, a protective measure against fluid freezing inside a pressure transmitter is to house the transmitter in an insulated, heated enclosure. The next photograph shows just such an enclosure with the cover removed:



Not surprisingly, this installation works well to protect all kinds of temperature-sensitive instruments from extreme cold. Here, we see an explosive gas sensor mounted inside a slightly different style of insulated enclosure, with the lid opened up for inspection:



19.7 Process/instrument suitability

On a fundamental level, pressure is universal. Regardless of the fluid in question; liquid or gas, hot or cold, corrosive or inert, pressure is nothing more than the amount of force exerted by that fluid over a unit area:

$$P = \frac{F}{A}$$

It should come as no surprise, then, that the common mechanical sensing elements for measuring pressure (bellows, diaphragm, bourdon tube, etc.) are equally applicable to all fluid pressure measurement applications, at least in principle. It is normally a matter of proper material selection and element strength (material thickness) to make a pressure instrument suitable for any range of process fluids.


Fill fluids used in pressure instruments – whether it be the dielectric liquid inside a differential capacitance sensor, the fill liquid of a remote or chemical seal system, or liquid used to fill a vertical section of impulse tubing – must be chosen so as to not adversely react with or contaminate the process.

Pure oxygen processes require that no system component have traces of hydrocarbon fluids present. While oxygen itself is not explosive, it greatly accelerates the combustion and explosive potential of any flammable substance. Therefore, a pressure gauge calibrated using oil as the working fluid in a deadweight tester would definitely *not* be suitable for pure oxygen service! The same may be said for a DP transmitter with a hydrocarbon-based fill inside its pressure-sensing capsule³⁰.

Pharmaceutical, medical, and food manufacturing processes require strict purity and the ability to disinfect all elements in the process system at will. Stagnant lines are not allowed in such processes, as microbe cultures may flourish in such “dead end” piping. Remote seals are very helpful in overcoming this problem, but the fill fluids used in remote systems must be chosen such that a leak in the isolating diaphragm will not contaminate the process.

³⁰Although this fluid would not *normally* contact pure oxygen in the process, it could if the isolating diaphragm inside the transmitter were to ever leak.

Manometers, of course, are rather limited in their application, as their operation depends on direct contact between process fluid and manometer liquid. In the early days of industrial instrumentation, liquid mercury was a very common medium for process manometers, and it was not unusual to see a mercury manometer used in direct contact with a process fluid such as oil or water to provide pressure indication:



Thankfully, those days are gone. Mercury (chemical symbol "Hg") is a toxic metal and therefore hazardous to work with. Calibration of these manometers was also challenging due to the column height of the process liquid in the impulse line and the range tube. When the process fluid is a gas, the difference in mercury column height directly translates to sensed pressure by the hydrostatic pressure formula $P = \rho gh$ or $P = \gamma h$. When the process fluid is a liquid, though, the shifting of mercury columns also creates a change in height of the process liquid column, which means the indicated pressure is a function of the height difference (h) and the difference in density between the process liquid and mercury. Consequently, the indications provided by mercury manometers in liquid pressure applications were subject to correction according to process liquid density.

19.8 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Definition of pressure:** $P = \frac{F}{A}$ (pressure is the amount of force applied over a specified area by a fluid).
- **Pascal's principle:** changes in fluid pressure are transmitted evenly throughout an enclosed fluid volume. Relevant to pressure measurement, as fluid pressure in all parts of an enclosed system will experience the same changes in pressure.
- **Hydrostatic pressure:** fluids having substantial weight generate pressure proportional to their density and to their vertical height ($P = \gamma h$ and $P = \rho gh$). Relevant to pressure offsets generated in vertical spans of impulse or capillary tubing, causing a pressure instrument to register more or less pressure than that at the process vessel connection.
- **Self-balancing pneumatic mechanisms:** all self-balancing pneumatic instruments work on the principle of negative feedback maintaining a nearly constant baffle-nozzle gap. Force-balance mechanisms maintain this constant gap by balancing force against force with negligible motion, like a tug-of-war. Motion-balance mechanisms maintain this constant gap by balancing one motion with another motion, like two dancers moving in unison.

References

Beckerath, Alexander von; Eberlein, Anselm; Julien, Hermann; Kersten, Peter; and Kretzner, Jochem, *WIKA-Handbook, Pressure and Temperature Measurement*, WIKA Alexander Wiegand GmbH & Co., Klingenberg, Germany, 1995.

Calkins, A.B., “Steam and Pressure Gauges”, *Cassier’s Magazine*, Volume 1, London, England, 1891.

“Digital Sensor Technology” (PowerPoint slideshow presentation), Yokogawa Corporation of America.

Fribance, Austin E., *Industrial Instrumentation Fundamentals*, McGraw-Hill Book Company, New York, NY, 1962.

Kallen, Howard P., *Handbook of Instrumentation and Controls*, McGraw-Hill Book Company, Inc., New York, NY, 1961.

Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.

Patrick, Dale R. and Patrick, Steven R., *Pneumatic Instrumentation*, Delmar Publishers, Inc., Albany, NY, 1993.

“Rosemount 1199 Diaphragm Seal Systems”, product data sheet 00813-0100-4016 revision HA, Rosemount, Emerson Process Management, January 2008.

Technical Note: “Rosemount 1199 Fill Fluid Specifications”, Rosemount, Emerson Process Management, 2005.

Chapter 20

Continuous level measurement

Many industrial processes require the accurate measurement of fluid or solid (powder, granule, etc.) height within a vessel. Some process vessels hold a stratified combination of fluids, naturally separated into different layers by virtue of differing densities, where the height of the *interface* point between liquid layers is of interest.

A wide variety of technologies exist to measure the level of substances in a vessel, each exploiting a different principle of physics. This chapter explores the major level-measurement technologies in current use.

20.1 Level gauges (sightglasses)


Level gauges are perhaps the simplest indicating instrument for liquid level in a vessel. They are often found in industrial level-measurement applications, even when another level-measuring instrument is present, to serve as a direct indicator for an operator to monitor in case there is doubt about the accuracy of the other instrument.

20.1.1 Basic concepts of sightglasses

The *level gauge*, or *sightglass* is to liquid level measurement as manometers are to pressure measurement: a very simple and effective technology for direct visual indication of process level. In its simplest form, a level gauge is nothing more than a clear tube through which process liquid may be seen. The following photograph shows a simple example of a sightglass:



A functional diagram of a sightglass shows how it visually represents the level of liquid inside a vessel such as a storage tank:



A level gauge is not unlike a U-tube manometer, with equal pressures applied to both liquid columns (one column being the liquid in the gauge sightglass, the other column being the liquid in the vessel).

Level gauge valves exist to allow replacement of the glass tube without emptying or depressurizing the process vessel. These valves are usually equipped with flow-limiting devices in the event of a tube rupture, so too much process fluid does not escape even when the valves are fully open.

Some level gauges called *reflex gauges* are equipped with special optics to facilitate the viewing of clear liquids, which is problematic for simple glass-tube sightglasses.

A weakness of glass-tube level gauges is the glass tube itself. The tube must be kept in a clean condition in order for the liquid level to be clearly visible, which may be a problem in a dirty-liquid service. Also, glass tubes may rupture if subjected to thermal or mechanical shock. One solution to this problem is to eliminate the glass tube entirely, replacing it with a non-magnetic metal tube (e.g. stainless steel) containing a magnetized float, with magnet-sensing indicator flags outside of this tube to visually indicate level. Here is one example of such a level gauge, manufactured by MagTech:




In this instrument, you can see red-colored flags toward the bottom of the scale which have been “flipped” by the motion of the magnetic float inside the stainless-steel tube. The height of the red zone (i.e. how many flags have been flipped to show their red sides) indicates the height of the liquid inside the tube.

Some magnetic level gauges even have high- and low-level magnetic switches located at strategic points along the tube’s height, providing discrete sensing capability for alarms and/or shutdown controls, if the liquid level ever goes outside of safe operating limits. These switches will open and close as the magnetic float passes by, remotely signaling liquid level at that height.

20.1.2 Interface problems

As simple and apparently trouble-free as level gauges may seem, there are special circumstances where they will register incorrectly. One such circumstance is in the presence of a lighter liquid layer existing between the connection ports of the gauge. If a lighter (less dense) liquid exists above a heavier (denser) liquid in the process vessel, the level gauge may not show the proper interface, if at all:




A practical application of level measurement where two liquids form an interface is where water exists in the presence of petroleum substances, such as diesel fuel. Water is a denser liquid than most oils, and these two liquids are immiscible¹, which means the denser water forms a separate layer beneath the lighter oil. Another application of interface measurement is in the oil and gas extraction industry, where water must be separated from petroleum fluids coming out of wells drilled deep into the earth. Fluid from an oil well enters a special vessel called a “separator” where gravity causes the water to separate from the petroleum liquids, and petroleum vapors to separate from all liquids. These water/oil/gas separator vessels are critical components of any petroleum well system, with the liquid-liquid interface between water and oil being an important process variable to measure and control.


In the above illustration we see how a column of water in the sightglass shows less (total) level than the combination of water and oil inside a process vessel. Since the oil lies between the two level gauge ports into the vessel (sometimes called *nozzles*), it cannot enter the sightglass tube, and therefore the level gauge will continue to show just water.

¹Liquids are considered “miscible” if they may be mixed in any proportion to each other to form a solution. Immiscible liquids refuse to mix thoroughly, and therefore tend to separate.


If by chance some oil does find its way into the sightglass tube – either by the interface level dropping below the lower nozzle or the total level rising above the upper nozzle – the oil/water interface shown inside the level gauge may not continue to reflect the true interface inside the vessel once the interface and total levels return to their previous positions:



Recall that the level gauge and vessel together form a U-tube manometer. So long as the pressures from each liquid column are the same, the columns balance each other. The problem is, many different liquid-liquid interface columns can have the same hydrostatic pressure without being identical to one another:




The only way to ensure proper two-part liquid interface level indication in a sightglass is to keep both ports (nozzles) submerged:



20.1.3 Temperature problems

Another troublesome scenario for level gauges is when the liquid inside the vessel is substantially hotter than the liquid in the gauge, causing the densities to be different. This is commonly seen on boiler level gauges, where the water inside the sightglass cools off substantially from its former temperature inside the boiler drum:




Looking at the sightglass as a U-tube manometer again, we see that unequal-height liquid columns may indeed balance each other's hydrostatic pressures if the two columns are comprised of liquids with different densities. The weight density of water is 62.4 lb/ft^3 at standard temperature, but may be as low as only 36 lb/ft^3 at temperatures common for power generation boilers.

20.2 Float

Perhaps the simplest form of solid or liquid level measurement is with a *float*: a device that rides on the surface of the fluid or solid within the storage vessel. The float itself must be of substantially lesser density than the substance of interest, and it must not corrode or otherwise react with the substance.


Floats may be used for manual “gauging” of level, as illustrated here:




A person lowers a float down into a storage vessel using a flexible measuring tape, until the tape goes slack due to the float coming to rest on the material surface. At that point, the person notes the length indicated on the tape (reading off the lip of the vessel access hole). This distance is called the *ullage*, being the distance from the top of the vessel to the surface of the process material. *Fillage* of the vessel may be determined by subtracting this “ullage” measurement from the known height of the vessel.

Obviously, this method of level measurement is tedious and may pose risk to the person conducting the measurement. If the vessel is pressurized, this method is simply not applicable.

If we automate the person's function using a small winch controlled by a computer – having the computer automatically lower the float down to the material surface and measure the amount of cable played out at each measurement cycle – we may achieve better results without human intervention. Such a level gauge may be enclosed in such a way to allow pressurization of the vessel, too:



A simpler version of this technique uses a spring-reel to constantly tension the cable holding the float, such that the float continuously rides on the surface of the liquid in the vessel²:




²A spring-loaded cable float only works with liquid level measurement, while a retracting float will measure liquids and solids with equal ease. The reason for this limitation is simple: a float that always contacts the material surface is likely to become buried if the material in question is a solid (powder or granules), which must be fed into the vessel from above.

The following photograph shows the “measurement head” of a spring-reel tape-and-float liquid level transmitter, with the vertical pipe housing the tape on its way to the top of the storage tank where it will turn 180 degrees via two pulleys and attach to the float inside the tank:



The spring reel’s angular position may be measured by a multi-turn potentiometer or a rotary encoder (located inside the “head” unit), then converted to an electronic signal for transmission to a remote display, control, and/or recording system. Such systems are used extensively for measurement of water and fuel in storage tanks.


If the liquid inside the vessel is subject to turbulence, *guide wires* may be necessary to keep the float cable in a vertical orientation:



The guide wires are anchored to the floor and roof of the vessel, passing through ring lugs on the float to keep it from straying laterally.

One of the potential disadvantages of tape-and-float level measurement systems is fouling of the tape (and guide wires) if the substance is sticky or unclean.

A variation on the theme of float level measurement is to place a small float inside the tube of a sightglass-style level gauge:



The float's position inside the tube may be readily detected by ultrasonic waves, magnetic sensors or any other applicable means. Locating the float inside a tube eliminates the need for guide wires or a sophisticated tape retraction or tensioning system. If no visual indication is necessary, the

level gauge tube may be constructed out of metal instead of glass, greatly reducing the risk of tube breakage. All the problems inherent to sightglasses, however, still apply to this form of float instrument.

The following photograph shows just such a float-style level indicator, manufactured by K-Tek. This particular level indicator is removed from the process, sitting on a concrete floor for the photograph. A magnetic float rides on the surface of a liquid column inside a non-magnetic metal tube, while a brightly-colored indicating flag tracks the magnetic float's position inside a transparent plastic tube for convenient viewing:



Two major advantages of a magnetically-coupled float is increased pressure rating and safety (since the float tube need not be constructed of clear material such as plastic or glass), and increased readability (since the viewing tube will never get dirty with process fluid residue, and the float may be brightly colored).

This particular level indicator also has a special pneumatic valve mounted to the side of the non-magnetic metal tube. This valve is actuated by the magnetic field of the float, turning a pneumatic “circuit” on and off based on the float’s position:



This is just one example of auxiliary functions possible with magnetic float level indicators. Such a pneumatic valve may be used to control a larger process valve to redirect liquid to or from the vessel based on level, to trip an operator alarm, or any number of other automatic functions.

Another variation on the theme of auxiliary float functions is a principle called *magnetostriction* to detect the position of the float along a metal guide rod called a *waveguide*. This instrument design is discussed in significant detail later in this chapter (see [20.5.4](#) beginning on page [1486](#)).


20.3 Hydrostatic pressure

A vertical column of fluid generates a pressure at the bottom of the column owing to the action of gravity on that fluid. The greater the vertical height of the fluid, the greater the pressure, all other factors being equal. This principle allows us to infer the level (height) of liquid in a vessel by pressure measurement.

20.3.1 Pressure of a fluid column

A vertical column of fluid exerts a pressure due to the column's weight. The relationship between column height and fluid pressure at the bottom of the column is constant for any particular fluid (density) regardless of vessel width or shape.

This principle makes it possible to infer the height of liquid in a vessel by measuring the pressure generated at the bottom:



The mathematical relationship between liquid column height and pressure is as follows:

$$P = \rho gh$$

$$P = \gamma h$$

Where,

P = Hydrostatic pressure

ρ = Mass density of fluid in kilograms per cubic meter (metric) or slugs per cubic foot (British)

g = Acceleration of gravity

γ = Weight density of fluid in newtons per cubic meter (metric) or pounds per cubic foot (British)

h = Height of vertical fluid column above point of pressure measurement

For example, the pressure generated by a column of oil 12 feet high (h) having a weight density of 40 pounds per cubic foot (γ) is:

$$P = \gamma h$$

$$P_{oil} = \left(\frac{40 \text{ lb}}{\text{ft}^3} \right) \left(\frac{12 \text{ ft}}{1} \right) = \frac{480 \text{ lb}}{\text{ft}^2}$$

Note the cancellation of units, resulting in a pressure value of 480 pounds per square foot (PSF). To convert into the more common pressure unit of pounds per square inch, we may multiply by the proportion of square feet to square inches, eliminating the unit of square feet by cancellation and leaving square inches in the denominator:

$$P_{oil} = \left(\frac{480 \text{ lb}}{\text{ft}^2} \right) \left(\frac{1^2 \text{ ft}^2}{12^2 \text{ in}^2} \right)$$

$$P_{oil} = \left(\frac{480 \text{ lb}}{\text{ft}^2} \right) \left(\frac{1 \text{ ft}^2}{144 \text{ in}^2} \right)$$

$$P_{oil} = \frac{3.33 \text{ lb}}{\text{in}^2} = 3.33 \text{ PSI}$$

Thus, a pressure gauge attached to the bottom of the vessel holding a 12 foot column of this oil would register 3.33 PSI. It is possible to customize the scale on the gauge to read directly in feet of oil (height) instead of PSI, for convenience of the operator who must periodically read the gauge. Since the mathematical relationship between oil height and pressure is both linear and direct, the gauge's indication will always be proportional to height.

An alternative method for calculating pressure generated by a liquid column is to relate it to the pressure generated by an equivalent column of water, resulting in a pressure expressed in units of water column (e.g. inches W.C.) which may then be converted into PSI or any other unit desired.

For our hypothetical 12-foot column of oil, we would begin this way by calculating the *specific gravity* (i.e. how dense the oil is compared to water). With a stated weight density of 40 pounds per cubic foot, the specific gravity calculation looks like this:

$$\text{Specific Gravity of oil} = \frac{\gamma_{oil}}{\gamma_{water}}$$

$$\text{Specific Gravity of oil} = \frac{40 \text{ lb/ft}^3}{62.4 \text{ lb/ft}^3}$$

$$\text{Specific Gravity of oil} = 0.641$$

The hydrostatic pressure generated by a column of water 12 feet high, of course, would be 144 inches of water column (144 "W.C.). Since we are dealing with an oil having a specific gravity of 0.641 instead of water, the pressure generated by the 12 foot column of oil will be only 0.641 times (64.1%) that of a 12 foot column of water, or:

$$P_{oil} = (P_{water})(\text{Specific Gravity})$$

$$P_{oil} = (144 \text{ "W.C.})(0.641)$$

$$P_{oil} = 92.3 \text{ "W.C.}$$

We may convert this pressure value into units of PSI simply by dividing by 27.68, since we know 27.68 inches of water column is equivalent to 1 PSI:

$$P_{oil} = \left(\frac{92.3 \text{ "W.C.}}{1} \right) \left(\frac{1 \text{ PSI}}{27.68 \text{ "W.C.}} \right)$$

$$P_{oil} = 3.33 \text{ PSI}$$

As you can see, we arrive at the same result as when we applied the $P = \gamma h$ formula. Any difference in value between the two methods is due to imprecision of the conversion factors used (e.g. 27.68 "W.C., 62.4 lb/ft³ density for water).

Any type of pressure-sensing instrument may be used as a liquid level transmitter by means of this principle. In the following photograph, you see a Rosemount model 1151 pressure transmitter being used to measure the height of colored water inside a clear plastic tube:




In most level-measurement applications, we are concerned with knowing the *volume* of the liquid contained within a vessel, and we infer this volume by using instruments to sense the *height* of the fluid column. So long as the vessel's cross-sectional area is constant throughout its height, liquid height will be directly proportional to stored liquid volume. Pressure measured at the bottom of a vessel can give us a proportional indication of liquid height if and only if the density of that liquid is known and constant. This means liquid density is a critically important factor for volumetric measurement when using hydrostatic pressure-sensing instruments. If liquid density is subject to random change, the accuracy of any hydrostatic pressure-based level or volume instrument will be correspondingly unreliable.

It should be noted, though, that changes in liquid density will have absolutely no effect on hydrostatic measurement of liquid *mass*, so long as the vessel has a constant cross-sectional area throughout its entire height. A simple thought experiment proves this: imagine a vessel partially full of liquid, with a pressure transmitter attached to the bottom to measure hydrostatic pressure. Now imagine the temperature of that liquid increasing, such that its volume expands and has a lower density than before. Assuming no addition or loss of liquid to or from the vessel, any increase in liquid level will be strictly due to volume expansion (density decrease). Liquid level inside this vessel will rise, but the transmitter will sense the exact same hydrostatic pressure as before, since the rise in level is precisely countered by the decrease in density (if h increases by the same factor

that γ decreases, then $P = \gamma h$ must remain the same!). In other words, hydrostatic pressure is seen to be directly proportional to the amount of liquid *mass* contained within the vessel, regardless of changes in liquid density. This is useful to know in applications where true mass measurement of a liquid (rather than volume measurement) is either preferable or necessary³.


Differential pressure transmitters are the most common pressure-sensing device used in this capacity to infer liquid level within a vessel. In the hypothetical case of the oil vessel just considered, the transmitter would connect to the vessel in this manner (with the high side toward the process and the low side vented to atmosphere):



Connected as such, the differential pressure transmitter functions as a gauge pressure transmitter, responding to hydrostatic pressure exceeding ambient (atmospheric) pressure. As liquid level increases, the hydrostatic pressure applied to the “high” side of the differential pressure transmitter also increases, driving the transmitter’s output signal higher.

³We may prove this mathematically by algebraic substitution. Given that the total mass (m) of any liquid sample is equal to the product of that liquid’s mass density and its sample volume ($m = \rho V$), that volume (V) for any vessel of constant cross-sectional area (A) is given by the expression $V = Ah$, and that hydrostatic pressure is equal to $P = \rho gh$, we may combine these three equations to arrive at $m = \frac{AP}{g}$. This final equation demonstrates how the total mass of liquid stored in a vessel (m) of constant cross-sectional area (A) is directly proportional to pressure (P), and independent of density (ρ).

Some pressure-sensing instruments are built specifically for hydrostatic measurement of liquid level in vessels, eliminating with impulse tubing altogether in favor of a special kind of sealing diaphragm extending slightly into the vessel through a flanged pipe entry (commonly called a *nozzle*). A Rosemount hydrostatic level transmitter with an extended diaphragm is shown here:




The calibration table for a transmitter close-coupled to the bottom of an oil storage tank would be as follows, assuming a zero to twelve foot measurement range for oil height, an oil density of 40 pounds per cubic foot, and a 4-20 mA transmitter output signal range:

Oil level	Percent of range	Hydrostatic pressure	Transmitter output
0 ft	0 %	0 PSI	4 mA
3 ft	25 %	0.833 PSI	8 mA
6 ft	50 %	1.67 PSI	12 mA
9 ft	75 %	2.50 PSI	16 mA
12 ft	100 %	3.33 PSI	20 mA

20.3.2 Bubbler systems


An interesting variation on this theme of direct hydrostatic pressure measurement is the use of a purge gas to measure hydrostatic pressure in a liquid-containing vessel. This eliminates the need for direct contact of the process liquid against the pressure-sensing element, which can be advantageous if the process liquid is corrosive.

Such systems are often called *bubble tube* or *dip tube* systems, the former name being appropriately descriptive for the way purge gas bubbles out the end of the tube as it is submerged in process liquid. A very simple bubbler system may be simulated by gently blowing air through a straw into a glass of water, maintaining a steady rate of bubbles exiting the straw while changing the depth of the straw's end in the water:



The deeper you submerge the straw, the harder it becomes to blow bubbles out the end with your breath. The hydrostatic pressure of the water at the straw's tip becomes translated into air pressure in your mouth as you blow, since the air pressure must just exceed the water's pressure in order to escape out the end of the straw. So long as the flow rate of air is modest (no more than a few bubbles per second), the air pressure will be very nearly equal to the water pressure, allowing measurement of water pressure (and therefore water depth) at any point along the length of the air tube.


If we lengthen the straw and measure pressure at all points throughout its length, it will be the same as the pressure at the submerged tip of the straw (assuming negligible friction between the moving air molecules and the straw's interior walls):



This is how industrial “bubbler” level measurement systems work: a *purge gas* is slowly introduced into a “dip tube” submerged in the process liquid, so that no more than a few bubbles per second of gas emerge from the tube’s end. Gas pressure inside all points of the tubing system will (very nearly) equal the hydrostatic pressure of the liquid at the tube’s submerged end. Any pressure-measuring device tapped anywhere along the length of this tubing system will sense this pressure and be able to infer the depth of the liquid in the process vessel without having to directly contact the process liquid.

Bubbler-style liquid level measurement systems are especially useful when the process liquid in question is highly corrosive, prone to plugging sample ports, or in any other way objectionable to have in direct contact with a pressure sensor. Unlike pressure sensors which must use diaphragms or other flexible (usually metallic) sensing elements and therefore may only be constructed from a limited range of materials, a dip tube need not be flexible and therefore may be constructed of *any* material capable of withstanding the process liquid. A process liquid so corrosive that non-metallic vessels are required to hold it would preclude direct contact with a metal pressure gauge or pressure transmitter, but would be easily measured with a bubbler system provided the dip tube were made out of plastic, ceramic, or some other material immune to corrosion. A process liquid so laden with solids that it plugs up any non-flowing port would preclude pressure measurement via a sample port and impulse line, but would be easily measured by a bubbler system where the dip tube is continuously purged with clean gas. Level measurement applications where direct contact with the pressure sensor would render access to that sensor inconvenient or even impossible are made much more practical by using a bubbler system, where the pressure sensor may be located *anywhere* along the dip tube’s length and therefore easily located where maintenance personnel can access it.

Excessive purge gas flow through the tube will result in additional pressure caused by frictional pressure drop along the tube's length, causing the pressure-sensing instrument to falsely register high. A key detail of any practical bubble tube system, therefore, is some means to monitor and control gas flow through the tube. A common construction uses either a *rotameter* or a *sightfeed bubbler* to monitor purge gas flow rate, with a needle valve to restrict that flow:



A more sophisticated solution to the problem of purge gas flow rate is to install a *flow-regulator* in lieu of a pressure regulator and needle valve, a mechanism designed to automatically monitor and throttle gas flow to maintain a constant purge rate. Flow regulators compensate for changes in dip tube pressure and in gas supply pressure, eliminating the need for a human operator to periodically adjust a needle valve.

Limiting the flow of purge gas is also important if that purge gas is expensive to obtain. For bottled gases such as nitrogen (necessary in processes requiring a non-reactive purge), the cost of purchasing tanks of compressed gas is obvious. For air-purged systems the cost is still present, but not so obvious: the cost of running an air compressor to maintain continuous purge air pressure. Either way, limiting the flow rate of purge gas in a bubbler system yields economic benefits aside from increased measurement accuracy.


As with all purged systems, certain criteria must be met for successful operation. Listed here are some of them:

- The purge gas supply must be reliable: if the flow stops for any reason, the level measurement will cease to be accurate, and the dip tube may even plug with debris!
- The purge gas supply pressure must exceed the hydrostatic pressure at all times, or else the level measurement range will fall below the actual liquid level.
- The purge gas flow must be maintained at a low rate, to avoid pressure drop errors (i.e. excess pressure measured due to friction of the purge gas through the tube).
- The purge gas must not adversely react with the process.
- The purge gas must not contaminate the process.
- The purge gas must be reasonable in cost, since it will be continuously consumed over time.

One measurement artifact of a bubble tube system is a slight variation in pressure each time a new bubble breaks away from the end of the tube. The amount of pressure variation is approximately equal to the hydrostatic pressure of process fluid at a height equal to the diameter of the bubble, which in turn will be approximately equal to the diameter of the bubble tube. For example, a $\frac{1}{4}$ inch diameter dip tube will experience pressure oscillations with a peak-to-peak amplitude of approximately $\frac{1}{4}$ inch elevation of process liquid. The frequency of this pressure oscillation, of course, will be equal to the rate at which individual bubbles escape out the end of the dip tube.

Usually, this is a small variation when considered in the context of the measured liquid height in the vessel. A pressure oscillation of approximately $\frac{1}{4}$ inch compared to a measurement range of 0 to 10 feet, for example, is only about 0.2% of span. Modern pressure transmitters have the ability to “filter” or “damp” pressure variations over time, which is a useful feature for minimizing the effect such a pressure variation will have on system performance.

A way to help minimize this effect is to place small V-shaped notches at the end of the dip tube, to help bubbles escape at sizes smaller than the tube’s diameter:




20.3.3 Transmitter suppression and elevation

A very common scenario for liquid level measurement is where the pressure-sensing instrument is not located at the same level as the 0% measurement point. The following photograph shows an example of this, where a Rosemount model 3051 differential pressure transmitter is being used to sense hydrostatic pressure of colored water inside a (clear) vertical plastic tube:




Consider the example of a pressure sensor measuring the level of liquid ethanol in a storage tank. The measurement range for liquid height in this ethanol storage tank is 0 to 40 feet, but the transmitter is located 30 feet below the tank:



This means the transmitter's impulse line contains a 30-foot elevation head of ethanol, so the transmitter "sees" 30 feet of ethanol when the tank is empty and 70 feet of ethanol when the tank is full. A 3-point calibration table for this instrument would look like this, assuming a 4 to 20 mA DC output signal range:

Ethanol level in tank	Percent of range	Pressure (inches of water)	Pressure (PSI)	Output (mA)
0 ft	0 %	284 "W.C.	10.3 PSI	4 mA
20 ft	50 %	474 "W.C.	17.1 PSI	12 mA
40 ft	100 %	663 "W.C.	24.0 PSI	20 mA

Another common scenario is where the transmitter is mounted at or near the vessel's bottom, but the desired level measurement range does not extend to the vessel bottom:




In this example, the transmitter is mounted exactly at the same level as the vessel bottom, but the level measurement range begins at 4 feet up from the vessel bottom. At the level of castor oil deemed 0%, the transmitter “sees” a hydrostatic pressure of 1.68 PSI (46.5 inches of water column) and at the 100% castor oil level the transmitter “sees” a pressure of 3.78 PSI (105 inches water column). Thus, these two pressure values would define the transmitter’s lower and upper range values (LRV and URV), respectively.

The term for describing either of the previous scenarios, where the lower range value (LRV) of the transmitter’s calibration is a positive number, is called *zero suppression*⁴. If the zero offset is reversed (e.g. the transmitter mounted at a location *higher* than the 0% process level), it is referred to as *zero elevation*⁵.

⁴Or alternatively, *zero depression*.

⁵There is some disagreement among instrumentation professionals as to the definitions of these two terms. According to Béla G. Lipták’s *Instrument Engineers’ Handbook, Process Measurement and Analysis* (Fourth Edition, page 67), “suppressed zero range” refers to the transmitter being located below the 0% level (the LRV being a positive pressure value), while “suppression,” “suppressed range,” and “suppressed span” mean exactly the opposite (LRV is a negative value). The Yokogawa Corporation defines “suppression” as a condition where the LRV is a positive pressure (“Autolevel” Application Note), as does the Michael MacBeth in his CANDU Instrumentation & Control course (lesson 1, module 4, page 12), Foxboro’s technical notes on bubble tube installations (pages 4 through 7), and Rosemount’s product manual for their 1151 Alphaline pressure transmitter (page 3-7). Interestingly, the Rosemount document defines “zero range suppression” as synonymous with “suppression,” which disagrees with Lipták’s distinction. My advice: draw a picture if you want the other person to clearly understand what you mean!

If the transmitter is elevated above the process connection point, it will most likely “see” a negative pressure (vacuum) with an empty vessel owing to the pull of liquid in the line leading down from the instrument to the vessel. It is vitally important in elevated transmitter installations to use a *remote seal* rather than an open impulse line, so liquid cannot dribble out of this line and into the vessel⁶:




In this example, we see a remote seal system with a fill fluid having a density of 58.3 lb/ft^3 , and a process level measurement range of 0 to 11 feet of sea water (density = 64 lb/ft^3). The transmitter elevation is 6 feet, which means it will “see” a vacuum of -2.43 PSI (-67.2 inches of water column) when the vessel is completely empty. This, of course, will be the transmitter’s calibrated lower range value (LRV). The upper range value (URV) will be the pressure “seen” with 11 feet of sea water in the vessel. This much sea water will contribute an additional 4.89 PSI of hydrostatic pressure at the level of the remote seal diaphragm, causing the transmitter to experience a pressure of $+2.46 \text{ PSI}$ ⁷.

⁶As you are about to see, the calibration of an elevated transmitter depends on us knowing how much hydrostatic pressure (or vacuum, in this case) is generated within the tube connecting the transmitter to the process vessel. If liquid were to ever escape from this tube, the hydrostatic pressure would be unpredictable, and so would be the accuracy of our transmitter as a level-measuring instrument. A remote seal diaphragm guarantees no fill fluid will be lost if and when the process vessel goes empty.

⁷The sea water’s positive pressure at the remote seal diaphragm adds to the negative pressure already generated by the downward length of the capillary tube’s fill fluid (-2.43 PSI), which explains why the transmitter only “sees” 2.46 PSI of pressure at the 100% full mark.

20.3.4 Compensated leg systems


The simple and direct relationship between liquid height in a vessel and pressure at the bottom of that vessel is ruined if another source of pressure exists inside the vessel other than hydrostatic (elevation head). This is virtually guaranteed to be the case if the vessel in question is unvented. Any gas or vapor pressure accumulation in an enclosed vessel will add to the hydrostatic pressure at the bottom, causing any pressure-sensing instrument to falsely register a high level:



A pressure transmitter has no way of “knowing” how much of the sensed pressure is due to liquid elevation and how much of it is due to pressure existing in the vapor space above the liquid. Unless a way can be found to compensate for any non-hydrostatic pressure in the vessel, this extra pressure will be interpreted by the transmitter as additional liquid level.

Moreover, this error will change as gas pressure inside the vessel changes, so it cannot simply be “calibrated away” by a static zero shift within the instrument. The only way to hydrostatically measure liquid level inside an enclosed (non-vented) vessel is to continuously compensate for gas pressure.


Fortunately, the capabilities of a *differential* pressure transmitter make this a simple task. All we need to do is connect a second impulse line (called a *compensating leg*), from the “Low” port of the transmitter to the top of the vessel, so the “Low” side of the transmitter experiences nothing but the gas pressure enclosed by the vessel, while the “High” side experiences the *sum* of gas and hydrostatic pressures. Since a differential pressure transmitter responds only to *differences* in pressure between “High” and “Low” sides, it will naturally subtract the gas pressure (P_{gas}) to yield a measurement based solely on hydrostatic pressure (γh):



$$(P_{gas} + \gamma h) - P_{gas} = \gamma h$$

The amount of gas pressure inside the vessel now becomes completely irrelevant to the transmitter’s indication, because its effect is canceled at the differential pressure instrument’s sensing element. If gas pressure inside the vessel were to increase while liquid level remained constant, the pressure sensed at *both* ports of the differential pressure transmitter would increase by the exact same amount, with the pressure *difference* between the “high” and “low” ports remaining absolutely constant with the constant liquid level. This means the instrument’s output signal is a representation of hydrostatic pressure only, which represents liquid height (assuming a known liquid density γ).

Unfortunately, it is common for enclosed vessels to hold condensable vapors, which may over time fill a compensating leg full of liquid. If the tube connecting the “Low” side of a differential pressure transmitter fills completely with a liquid, this will add a hydrostatic pressure to that side of the transmitter, causing another calibration shift. This *wet leg* condition makes level measurement more complicated than a *dry leg* condition where the only pressure sensed by the transmitter’s “Low” side is gas pressure (P_{gas}):




$$(P_{gas} + \gamma_1 h_1) - (P_{gas} + \gamma_2 h_2) = \gamma_1 h_1 - \gamma_2 h_2$$

Gas pressure still cancels due to the differential nature of the pressure transmitter, but now the transmitter’s output indicates a difference of hydrostatic pressures between the vessel and the wet leg, rather than just the hydrostatic pressure of the vessel’s liquid level. Fortunately, the hydrostatic pressure generated by the wet leg will be constant, so long as the density of the condensed vapors filling that leg (γ_2) is constant. If the wet leg’s hydrostatic pressure is constant, we can compensate for it by calibrating the transmitter with an intentional zero shift, so it indicates as though it were measuring hydrostatic pressure on a vented vessel.

$$\text{Differential pressure} = \gamma_1 h_1 - \text{Constant}$$

We may ensure a constant density of wet leg liquid by intentionally filling that leg with a liquid known to be denser than the densest condensed vapor inside the vessel and non-miscible with the process fluid. We could also use a differential pressure transmitter with remote seals and capillary tubes filled with liquid of known density:




Remote seals are very useful in applications such as this, as the wet leg never requires re-filling.

An actual level transmitter installation using two remote seals (in this case, a Foxboro IDP10 differential pressure transmitter) appears in this photograph:



The vessel itself is insulated, and covered in sheet aluminum to protect the thermal insulation from impact and weather-related damage. White-painted flanged “nozzles” protrude from the vessel through the insulation to provide places for the level-sensing instrument to connect. You can see the two flanged remote seals (painted blue) where the armored (stainless steel) capillary tubes terminate. Note how the long capillary tube on the “wet” leg is neatly coiled to reduce the possibility of damage by snagging on any moving object.


An accessory commonly used with non-sealed (non-capillary) “wet leg” systems is a *seal pot*. This is a chamber at the top of the wet leg joining the wet leg line with the impulse line to the upper connection point on the process vessel. This “seal pot” maintains a small volume of liquid in it to allow for occasional liquid loss during transmitter maintenance procedures without greatly affecting the height of the liquid column in the wet leg:



Regular operation of the transmitter’s three-valve manifold (and drain valve) during routine instrument maintenance inevitably releases some liquid volume from the wet leg. Without a seal pot, even a small loss of liquid in the wet leg may create a substantial loss in liquid column height within that tube, given the tube’s small diameter. With a seal pot, the (comparatively) large liquid volume held by the pot allows for some liquid loss through the transmitter’s manifold without substantially affecting the height of the liquid column within the wet leg.

Seal pots are standard on level measurement systems for boiler steam drums, where steam readily condenses in the upper impulse tube to naturally form a wet leg. Although steam will condense over time to refill the wet leg following a loss of water in that leg, the level measurements taken during that re-fill time will be in error. The presence of a seal pot practically eliminates this error as the steam condenses to replenish the water lost from the pot, since the amount of height change inside the pot due to a small volume loss is trivial compared to the height change in a wet leg lacking a seal pot.

The following example shows the calibration table for a compensated-leg (wet) hydrostatic level measurement system, for a gasoline storage vessel and water as the wet leg fill fluid. Here, I am assuming a density of 41.0 lb/ft^3 for gasoline and 62.4 lb/ft^3 for water, with a 0 to 10 foot measurement range and an 11 foot wet leg height:




Gasoline level	Percent of range	Differential pressure at transmitter	Transmitter output
0 ft	0 %	-4.77 PSI	4 mA
2.5 ft	25 %	-4.05 PSI	8 mA
5 ft	50 %	-3.34 PSI	12 mA
7.5 ft	75 %	-2.63 PSI	16 mA
10 ft	100 %	-1.92 PSI	20 mA

Note that due to the superior density and height of the wet (water) leg, the transmitter *always* sees a negative pressure (pressure on the “Low” side exceeds pressure on the “High” side).

With some older differential pressure transmitter designs, this negative pressure was a problem. Consequently, it is common to see “wet leg” hydrostatic transmitters installed with the “Low” port connected to the bottom of the vessel and the “High” port connected to the compensating leg. In fact, it is *still* common to see modern differential pressure transmitters installed in this manner⁸, although modern transmitters may be ranged for negative pressures just as easily as for positive pressures. It is vitally important to recognize that any differential pressure transmitter connected as such (for any reason) will respond in reverse fashion to increases in liquid level. That is to say, as the liquid level in the vessel rises, the transmitter’s output signal will *decrease* instead of increase:

High side of DP transmitter connected to the compensating impulse leg




Either way of connecting the transmitter to the vessel will suffice for measuring liquid level, so long as the instrumentation receiving the transmitter’s signal is properly configured to interpret the signal. The choice of which way to connect the transmitter to the vessel should be driven by fail-safe system design, which means to design the measurement system such that the most probable system failures – including broken signal wires – result in the control system “seeing” the most dangerous process condition and therefore taking the safest action.

⁸Sometimes this is done out of habit, other times because instrument technicians do not know the capabilities of new technology.

20.3.5 Tank expert systems


An alternative to using a compensating leg to subtract gas pressure inside an enclosed vessel is to simply use a second pressure transmitter and electronically subtract the two pressures in a computing device:



This approach enjoys the distinct advantage of avoiding a potentially wet compensating leg, but suffers the disadvantages of extra cost and greater error due to the potential calibration drift of *two* transmitters rather than just one. Such a system is also impractical in applications where the gas pressure is substantial compared to the hydrostatic (elevation head) pressure⁹.

⁹This is due to limited transmitter resolution. Imagine an application where the elevation head was 10 PSI (maximum) yet the vapor space pressure was 200 PSI. The majority of each transmitter's working range would be "consumed" measuring gas pressure, with hydrostatic head being a mere 5% of the measurement range. This would make precise measurement of liquid level very difficult, akin to trying to measure the sound intensity of a whisper in a noisy room.

If we add a third pressure transmitter to this system, located a known distance (x) above the bottom transmitter, we have all the pieces necessary for what is called a *tank expert system*. These systems are used on large storage tanks operating at or near atmospheric pressure, and have the ability to measure infer liquid height, liquid density, total liquid volume, and total liquid mass stored in the tank:



The pressure difference between the bottom and middle transmitters will change only if the liquid density changes¹⁰, since the two transmitters are separated by a known and fixed height difference.

¹⁰ Assuming the liquid level is equal to or greater than x . Otherwise, the pressure difference between P_{bottom} and P_{middle} will depend on liquid density *and* liquid height. However, this condition is easy to check: the level computer simply checks to see if P_{middle} and P_{top} are unequal. If so, then the computer knows the liquid level exceeds x and it is safe to calculate density. If not, and P_{middle} registers the same as P_{top} , the computer knows those two transmitters are both registering gas pressure only, and it knows to stop calculating density.

Algebraic manipulation shows us how the measured pressures may be used by the level computer (LY) to continuously calculate liquid density (γ):

$$P_{bottom} - P_{middle} = (P_{gas} + \gamma h) - [P_{gas} + \gamma(h - x)]$$

$$P_{bottom} - P_{middle} = P_{gas} + \gamma h - P_{gas} - \gamma(h - x)$$

$$P_{bottom} - P_{middle} = P_{gas} + \gamma h - P_{gas} - \gamma h + \gamma x$$

$$P_{bottom} - P_{middle} = \gamma x$$

$$\frac{P_{bottom} - P_{middle}}{x} = \gamma$$

Once the computer knows the value of γ , it may calculate the height of liquid in the tank with great accuracy based on the pressure measurements taken by the bottom and top transmitters:

$$P_{bottom} - P_{top} = (P_{gas} + \gamma h) - P_{gas}$$

$$P_{bottom} - P_{top} = \gamma h$$

$$\frac{P_{bottom} - P_{top}}{\gamma} = h$$

With all the computing power available in the LY, it is possible to characterize the tank such that this height measurement converts to a precise volume measurement¹¹ (V), which may then be converted into a total mass (m) measurement based on the mass density of the liquid (ρ) and the acceleration of gravity (g). First, the computer calculates mass density based on the proportionality between mass and weight (shown here starting with the equivalence between the two forms of the hydrostatic pressure formula):

$$\rho g h = \gamma h$$

$$\rho g = \gamma$$

$$\rho = \frac{\gamma}{g}$$

¹¹The details of this math depend entirely on the shape of the tank. For vertical cylinders – the most common shape for vented storage tanks – volume and height are related by the simple formula $V = \pi r^2 h$ where r is the radius of the tank's circular base. Other tank shapes and orientations may require much more sophisticated formulae to calculate stored volume from height. See section 26.3 beginning on page 2065, for more details on this subject.

Armed with the mass density of the liquid inside the tank, the computer may now calculate total liquid mass stored inside the tank:

$$m = \rho V$$

Dimensional analysis shows how units of mass density and volume cancel to yield only units of mass in this last equation:

$$[\text{kg}] = \left[\frac{\text{kg}}{\text{m}^3} \right] [\text{m}^3]$$

Here we see a vivid example of how several measurements may be inferred from just a few actual process (in this case, pressure) measurements. Three pressure measurements on this tank allow us to compute four inferred variables: liquid density, liquid height, liquid volume, and liquid mass.

The accurate measurement of liquids in storage tanks is not just useful for process operations, but also for conducting business. Whether the liquid represents raw material purchased from a supplier, or a processed product ready to be pumped out to a customer, both parties have a vested interest in knowing the exact quantity of liquid bought or sold. Measurement applications such as this are known as *custody transfer*, because they represent the transfer of custody (ownership) of a substance exchanged in a business agreement. In some instances, both buyer and seller operate and maintain their own custody transfer instrumentation, while in other instances there is only one instrument, the calibration of which validated by a neutral party.

A photograph showing the two lower pressure transmitters of a tank expert system on a refrigerated butane storage vessel appears here:




The upper and lower instruments are pressure transmitters, while the middle instrument is a temperature sensor used to report the temperature of the refrigerated butane to the control system.

20.3.6 Hydrostatic interface level measurement


Hydrostatic pressure sensors may be used to detect the level of a liquid-liquid interface, if and only if the total height of liquid never drops down into the sensing range of the interface level instrument. This is critically important because any single hydrostatic-based level instrument cannot discriminate between a changing interface level and a changing total level within the same range, and therefore the latter must be fixed in order to measure the former.

One way to ensure this condition is to fix the total liquid height to some constant value by using an overflow pipe, and ensuring drain flow is always less than incoming flow (forcing some flow to always go through the overflow pipe). This strategy naturally lends itself to separation processes, where a mixture of light and heavy liquids are separated by their differing densities:



Here we see a practical application for liquid-liquid interface level measurement. If the goal is to separate two liquids of differing densities from one another, we need only the light liquid to exit out the overflow pipe and only the heavy liquid to exit out the drain pipe. This means we must control the interface level to stay between those two piping points on the vessel. If the interface drifts too far up, heavy liquid will be carried out the overflow pipe; and if we let the interface drift too far down, light liquid will flow out of the drain pipe. The first step in controlling any process variable is to measure that variable, and so here we are faced with the necessity of measuring the interface point between the light and heavy liquids.

Another way to address the problem of varying total liquid height is to use a compensating leg located at a point on the vessel below the total liquid height. In this example, a transmitter with remote seals is used:



Since both sides of the differential pressure transmitter “see” the hydrostatic pressure generated by the liquid column above the top connection point ($\gamma_2 h_3$), this term naturally cancels. The result of this is that total liquid height becomes irrelevant, so long as it remains above the upper connection point:

$$(\gamma_1 h_1 + \gamma_2 h_2 + \gamma_2 h_3) - (\gamma_4 h_4 + \gamma_2 h_3)$$


$$\gamma_1 h_1 + \gamma_2 h_2 + \gamma_2 h_3 - \gamma_4 h_4 - \gamma_2 h_3$$

$$\gamma_1 h_1 + \gamma_2 h_2 - \gamma_4 h_4$$


The hydrostatic pressure in the compensating leg is constant ($\gamma_4 h_4 = \text{Constant}$), since the fill fluid never changes density and the compensating leg’s height never changes. This means the transmitter’s sensed pressure will differ from that of an uncompensated transmitter merely by a constant offset, which may be “calibrated out” so as to have no impact on the measurement:

$$\gamma_1 h_1 + \gamma_2 h_2 - \text{Constant}$$

At first, it may seem as though determining the calibration points (lower- and upper-range values: LRV and URV) for a hydrostatic interface level transmitter is impossibly daunting given all the different pressures involved. A recommended problem-solving technique to apply here is that of a *thought experiment*, where we imagine what the process will “look like” at lower-range value condition and at the upper-range value condition, drawing two separate illustrations:




For example, suppose we must calibrate a differential pressure transmitter to measure the interface level between two liquids having specific gravities of 1.1 and 0.78, respectively, over a span of 3 feet. The transmitter is equipped with remote seals, each containing a halocarbon fill fluid with a specific gravity of 1.09. The physical layout of the system is as follows:



As the first step in our “thought experiment,” we imagine what the process would look like with the interface at the LRV level, calculating hydrostatic pressures seen at each side of the transmitter. For this I recommend actually *sketching* a simple diagram showing what the fluid levels would look like with the interface at the LRV point, so the height dimensions of each liquid layer become obvious:

LRV interface level condition



We know from our previous exploration of this setup that any hydrostatic pressure resulting from liquid level *above* the top remote seal location is irrelevant to the transmitter, since it is “seen” on both sides of the transmitter and thus cancels out. All we must do, then, is calculate hydrostatic pressures as though the total liquid level stopped at that upper diaphragm connection point.

First, calculating the hydrostatic pressure “seen” at the high port of the transmitter¹²:

$$P_{high} = 4.5 \text{ feet of heavy liquid} + 4.5 \text{ feet of light liquid}$$

$$P_{high} = 54 \text{ inches of heavy liquid} + 54 \text{ inches of light liquid}$$

$$P_{high} \text{ "W.C.} = (54 \text{ inches of heavy liquid})(1.1) + (54 \text{ inches of light liquid})(0.78)$$

$$P_{high} \text{ "W.C.} = 59.4 \text{ "W.C.} + 42.12 \text{ "W.C.}$$

$$P_{high} = 101.52 \text{ "W.C.}$$

Next, calculating the hydrostatic pressure “seen” at the low port of the transmitter:

$$P_{low} = 9 \text{ feet of fill fluid}$$

$$P_{low} = 108 \text{ inches of fill fluid}$$

$$P_{low} \text{ "W.C.} = (108 \text{ inches of fill fluid})(1.09)$$

$$P_{low} = 117.72 \text{ "W.C.}$$


The differential pressure applied to the transmitter in this condition is the difference between the high and low port pressures, which becomes the lower range value (LRV) for calibration:

$$P_{LRV} = 101.52 \text{ "W.C.} - 117.72 \text{ "W.C.} = -16.2 \text{ "W.C.}$$

¹²Here I will calculate all hydrostatic pressures in units of inches water column. This is relatively easy because we have been given the specific gravities of each liquid, which make it easy to translate actual liquid column height into column heights of pure water.

As the second step in our “thought experiment,” we imagine what the process would look like with the interface at the URV level, calculating hydrostatic pressures seen at each side of the transmitter. Again, sketching a simple diagram of what the liquid layers would look like in this condition helps us visualize the scenario:

URV interface level condition



$$P_{high} = 7.5 \text{ feet of heavy liquid} + 1.5 \text{ feet of light liquid}$$

$$P_{high} = 90 \text{ inches of heavy liquid} + 18 \text{ inches of light liquid}$$

$$P_{high} \text{ "W.C.} = (90 \text{ inches of heavy liquid})(1.1) + (18 \text{ inches of light liquid})(0.78)$$

$$P_{high} \text{ "W.C.} = 99 \text{ "W.C.} + 14.04 \text{ "W.C.}$$

$$P_{high} = 113.04 \text{ "W.C.}$$

The hydrostatic pressure of the compensating leg is exactly the same as it was before: 9 feet of fill fluid having a specific gravity of 1.09, which means there is no need to calculate it again. It will still be 117.72 inches of water column. Thus, the differential pressure at the URV point is:

$$P_{URV} = 113.04 \text{ "W.C.} - 117.72 \text{ "W.C.} = -4.68 \text{ "W.C.}$$

Using these two pressure values and some interpolation, we may generate a 5-point calibration table (assuming a 4-20 mA transmitter output signal range) for this interface level measurement system:

Interface level	Percent of range	Differential pressure at transmitter	Transmitter output
4.5 ft	0 %	-16.2 "W.C.	4 mA
5.25 ft	25 %	-13.32 "W.C.	8 mA
6 ft	50 %	-10.44 "W.C.	12 mA
6.75 ft	75 %	-7.56 "W.C.	16 mA
7.5 ft	100 %	-4.68 "W.C.	20 mA

When the time comes to bench-calibrate this instrument in the shop, the easiest way to do so will be to set the two remote diaphragms on the workbench (at the same level), then apply 16.2 to 4.68 inches of water column pressure to the *low* remote seal diaphragm with the other diaphragm at atmospheric pressure to simulate the desired range of negative differential pressures¹³.

The more mathematically inclined reader will notice that the span of this instrument (Span = URV - LRV) is equal to the span of the interface level (3 feet, or 36 inches) multiplied by the difference in specific gravities (1.1 - 0.78):

$$\text{Span in "W.C.} = (36 \text{ inches})(1.1 - 0.78)$$

$$\text{Span} = 11.52 \text{ "W.C.}$$

Looking at our two “thought experiment” illustrations, we see that the only difference between the two scenarios is the type of liquid filling that 3-foot region between the LRV and URV marks. Therefore, the only difference between the transmitter’s pressures in those two conditions will be the difference in height multiplied by the difference in density. Not only is this an easy way for us to quickly calculate the necessary transmitter span, but it also is a way for us to check our previous work: we see that the difference between the LRV and URV pressures is indeed a difference of 11.52 inches water column just as this method predicts.


20.4 Displacement

Displacer level instruments exploit *Archimedes’ Principle* to detect liquid level by continuously measuring the weight of an object (called the *displacer*) immersed in the process liquid. As liquid level increases, the displacer experiences a greater buoyant force, making it appear lighter to the sensing instrument, which interprets the loss of weight as an increase in level and transmits a proportional output signal.

¹³Remember that a differential pressure instrument cannot “tell the difference” between a positive pressure applied to the low side, an equal vacuum applied to the high side, or an equivalent difference of two positive pressures with the low side’s pressure exceeding the high side’s pressure. Simulating the exact process pressures experienced in the field to a transmitter on a workbench would be exceedingly complicated, so we “cheat” by simplifying the calibration setup and applying the equivalent difference of pressure only to the “low” side.

20.4.1 Buoyant-force instruments

In practice a displacer level instrument usually takes the following form. Process piping in and out of the vessel has been omitted for simplicity – only the vessel and its displacer level instrument are shown:



The displacer itself is usually a sealed metal tube, weighted sufficiently so it cannot float in the process liquid. It hangs within a pipe called a “cage” connected to the process vessel through two block valves and nozzles. These two pipe connections ensure the liquid level inside the cage matches the liquid level inside the process vessel, much like a sightglass.

If liquid level inside the process vessel rises, the liquid level inside the cage rises to match. This will submerge more of the displacer’s volume, causing a buoyant force to be exerted upward on the displacer. Remember that that displacer is too heavy to float, so it does not “bob” on the surface of the liquid nor does it rise the same amount as the liquid’s level – rather, it hangs in place inside the cage, becoming “lighter¹⁴” as the buoyant force increases. The weight-sensing mechanism detects this buoyant force when it perceives the displacer becoming lighter, interpreting the decreased (apparent) weight as an increase in liquid level. The displacer’s apparent weight reaches a minimum when it is fully submerged, when the process liquid has reached the 100% point inside the cage.

It should be noted that static pressure inside the vessel will have negligible effect on a displacer instrument’s accuracy. The only factor that matters is the density of the process fluid, since buoyant force is directly proportional to fluid density ($F = \gamma V$).

¹⁴This is not unlike the experience of feeling lighter when you are standing in a pool of water just deep enough to submerge most of your body with your feet touching the bottom. This reduction of apparent weight is due to the buoyant force of the water upward on your body, equal to the weight of water that your body displaces.

The following photograph shows a Fisher “Level-Trol” model pneumatic transmitter measuring condensate level in a *knockout drum*¹⁵ for natural gas service. The instrument itself appears on the right-hand side of the photo, topped by a grey-colored “head” with two pneumatic pressure gauges visible. The displacer “cage” is the vertical pipe immediately behind and below the head unit. Note that a sightglass level gauge appears on the left-hand side of the knockout chamber (or *condensate boot*) for visual indication of condensate level inside the process vessel:



The purpose of this particular displacer instrument is to measure the amount of condensate liquid collected inside the “boot.” This model of Fisher Level-Trol comes complete with a pneumatic controller mechanism sending an air pressure signal to a drain valve to automatically drain the condensate out of the boot.

¹⁵So-called for its ability to “knock out” (separate and collect) condensable vapors from the gas stream. This particular photograph was taken at a natural gas compression facility, where it is very important the gas to be compressed is dry (since liquids are essentially incompressible). Sending even relatively small amounts of liquid into a compressor may cause the compressor to catastrophically fail!


Two photos of a disassembled Level-Trol displacer instrument appear here, showing how the displacer fits inside the cage pipe:



The cage pipe is coupled to the process vessel through two block valves, allowing isolation from the process. A drain valve allows the cage to be emptied of process liquid for instrument service and zero calibration.

Some displacer-type level sensors do not use a cage, but rather hang the displacer element directly in the process vessel. These are called “cageless” sensors. Cageless instruments are of course simpler than cage-style instruments, but they cannot be serviced without de-pressurizing (and perhaps even emptying) the process vessel in which they reside. They are also susceptible to measurement errors and “noise” if the liquid inside the vessel is agitated, either by high flow velocities in and out of the vessel, or by the action of motor-turned impellers installed in the vessel to provide thorough mixing of the process liquid(s).

Full-range calibration may be performed by flooding the cage with process liquid (a *wet* calibration), or by suspending the displacer with a string and precise scale (a *dry* calibration), pulling upward on the displacer at just the right amount to simulate buoyancy at 100% liquid level:



Calculation of this buoyant force is a simple matter. According to Archimedes' Principle, buoyant force is always equal to the weight of the fluid volume displaced. In the case of a displacer-based level instrument at full range, this usually means the entire volume of the displacer element is submerged in the liquid. Simply calculate the volume of the displacer (if it is a cylinder, $V = \pi r^2 l$, where r is the cylinder radius and l is the cylinder length) and multiply that volume by the weight density (γ):

$$F_{buoyant} = \gamma V$$

$$F_{buoyant} = \gamma \pi r^2 l$$

For example, if the weight density of the process fluid is 57.3 pounds per cubic foot and the displacer is a cylinder measuring 3 inches in diameter and 24 inches in length, the necessary force to simulate a condition of buoyancy at full level may be calculated as follows:

$$\gamma = \left(\frac{57.3 \text{ lb}}{\text{ft}^3} \right) \left(\frac{1 \text{ ft}^3}{12^3 \text{ in}^3} \right) = 0.0332 \frac{\text{lb}}{\text{in}^3}$$

$$V = \pi r^2 l = \pi (1.5 \text{ in})^2 (24 \text{ in}) = 169.6 \text{ in}^3$$

$$F_{\text{buoyant}} = \gamma V = \left(0.0332 \frac{\text{lb}}{\text{in}^3} \right) (169.6 \text{ in}^3) = 5.63 \text{ lb}$$


Note how important it is to maintain consistency of units! The liquid density was given in units of pounds per cubic *foot* and the displacer dimensions in *inches*, which would have caused serious problems without a conversion between feet and inches. In my example work, I opted to convert density into units of pounds per cubic inch, but I could have just as easily converted the displacer dimensions into feet to arrive at a displacer volume in units of cubic feet.

In a “wet” calibration, the 5.63 pound buoyant force will be created by the liquid itself, the technician ensuring there is enough liquid inside the cage to simulate a 100% level condition. In a “dry” calibration, the buoyant force will be simulated by tension applied upward on the displacer with a hand scale and string, the technician pulling with an upward force of 5.63 pounds to make the instrument “think” it is sensing 100% liquid level when in fact the displacer is completely dry, hanging in air.

20.4.2 Torque tubes

An interesting design problem for displacement-type level transmitters is how to transfer the sensed weight of the displacer to the transmitter mechanism while positively sealing process vapor pressure from that same mechanism. The most common solution to this problem is an ingenious mechanism called a *torque tube*. Unfortunately, torque tubes can be rather difficult to understand unless you have direct hands-on access to one, and so this section will explore the concept in more detail than is customarily available in reference manuals.

Imagine a solid, horizontal, metal rod with a flange at one end and a perpendicular lever at the other end. The flange is mounted to a stationary surface, and a weight suspended from the end of the lever. A dashed-line circle shows where the rod is welded to the center of the flange:




The downward force of the weight acting on the lever imparts a twisting force (torque) to the rod, causing it to slightly twist along its length. The more weight hung at the end of the lever, the more the rod will twist¹⁶. So long as the torque applied by the weight and lever never exceeds the elastic limit of the rod, the rod will continue to act as a spring. If we know the “spring constant” of the rod, and measured its torsional deflection, we can in fact use this slight motion to measure the magnitude of the weight hung at the end of the lever.

Applied to a displacer-type level instrument, a displacer takes the place of the weight at the lever’s end, the torsional deflection of this rod serving to indicate buoyant force. As liquid rises, buoyant force on the displacer increases, making the displacer seem lighter from the rod’s perspective. The rod’s slight motion resulting from this apparent weight change, then, indicates liquid level.


¹⁶To anyone familiar with the front suspension of a 1960’s vintage Chevrolet truck, or the suspension of the original Volkswagen “Beetle” car, the concept of a *torsion bar* should be familiar. These vehicles used straight, spring-steel rods to provide suspension force instead of the more customary coil springs used in modern vehicles. However, even the familiar coil spring is an example of torsional forces at work: a coil spring is nothing more than a torsion bar bent in a coil shape. As a coil spring is stretched or compressed, torsional forces develop along the circumferential length of the spring coil, which is what makes the spring “try” to maintain a fixed height.

Now imagine drilling a long hole through the rod, lengthwise, that almost reaches the end where the lever attaches. In other words, imagine a *blind hole* through the center of the rod, starting at the flange and ending just shy of the lever:




The presence of this long hole does not change much about the behavior of the assembly, except perhaps to alter the rod's spring constant. With less solid metal, the rod will be a weaker spring, and will twist to a greater degree with applied weight at the end of the lever. More importantly for the purpose of this discussion, though, the long hole transforms the rod into a *tube* with a sealed end. Instead of being a "torsion bar," the rod is now more properly called a *torque tube*, twisting ever so slightly with applied weight at the end of the lever.

In order to give the torque tube vertical support so it does not sag downward with the applied weight, a supporting *knife-edge bearing* is often placed underneath the end of the lever where it attaches to the torque tube. The purpose of this fulcrum is to provide vertical support for the weight while forming a virtually frictionless pivot point, ensuring the only stress applied to the torque tube is *torque* from the lever:



Finally, imagine another solid metal rod (slightly smaller diameter than the hole) spot-welded to the far end of the blind hole, extending out beyond the end of the flange:



The purpose of this smaller-diameter rod is to transfer the twisting motion of the torque tube's far end to a point past the flange where it may be sensed. Imagine the flange anchored to a vertical wall, while a variable weight tugs downward at the end of the lever. The torque tube will flex in a twisting motion with the variable force, but now we are able to see just how much it twists by watching the rotation of the smaller rod on the near side of the wall. The weight and lever may be completely hidden from our view by this wall, but the small rod's twisting motion nevertheless reveals how much the torque tube yields to the force of the weight.



We may apply this torque tube mechanism to the task of measuring liquid level in a pressurized vessel by replacing the weight with a displacer, attaching the flange to a nozzle welded to the vessel, and aligning a motion-sensing device with the small rod end to measure its rotation. As liquid level rises and falls, the apparent weight of the displacer varies, causing the torque tube to slightly twist. This slight twisting motion is then sensed at the end of the small rod, in an environment isolated from the process fluid pressure.

A photograph taken of a real torque tube from a Fisher “Level-Trol” level transmitter shows its external appearance:




The dark-colored metal is the elastic steel used to suspend the weight by acting as a torsional spring, while the shiny portion is the inner rod used to transfer motion. As you can see, the torque tube itself is not very wide in diameter. If it were, it would be far too stiff of a spring to be of practical use in a displacer-type level instrument, since the displacer is not typically very heavy, and the lever is not long.


Looking closer at each end of the torque tube reveals the open end where the small-diameter rod protrudes (left) and the “blind” end of the tube where it attaches to the lever (right):



If we were to slice the torque tube assembly in half, lengthwise, its cross-section would look something like this:



This next illustration shows the torque tube as part of a whole displacement-style level transmitter¹⁷:



As you can see from this illustration, the torque tube serves three distinct purposes when applied to a displacer-type level measurement application: (1) to serve as a torsional spring suspending the weight of the displacer, (2) to seal off process fluid pressure from the position-sensing mechanism, and (3) to transfer motion from the far end of the torque tube into the sensing mechanism.

¹⁷This illustration is simplified, omitting such details as access holes into the cage, block valves between the cage and process vessel, and any other pipes or instruments attached to the process vessel. Also, the position-sensing mechanism normally located at the far left of the assembly is absent from this drawing.

In pneumatic level transmitters, the sensing mechanism used to convert the torque tube's twisting motion into a pneumatic (air pressure) signal is typically of the *motion-balance* design. The Fisher Level-Trol mechanism, for example, uses a C-shaped bourdon tube with a nozzle at the end to follow a baffle attached to the small rod. The center of the bourdon tube is aligned with the center of the torque tube. As the rod rotates, the baffle advances toward the nozzle at the bourdon tube tip, causing backpressure to rise, which in turn causes the bourdon tube to flex. This flexing draws the nozzle away from the advancing baffle until a balanced condition exists. Rod motion is therefore balanced by bourdon tube motion, making this a motion-balance pneumatic system:



20.4.3 Displacement interface level measurement

Displacer level instruments may be used to measure liquid-liquid interfaces just the same as hydrostatic pressure instruments. One important requirement is that the displacer always be fully submerged (“flooded”). If this rule is violated, the instrument will not be able to discriminate between a low (total) liquid level and a low interface level. This criterion is analogous to the use of compensated-leg differential pressure instruments to measure liquid-liquid interface levels: in order for the instrument to solely respond to changes in interface level and not be “fooled” by changes in total liquid level, both process connection points must be submerged.

If the displacer instrument has its own “cage,” it is important that both pipes connecting the cage to the process vessel (sometimes called “nozzles”) be submerged. This ensures the liquid interface inside the cage matches the interface inside the vessel. If the upper nozzle ever goes dry, the same problem can happen with a caged displacer instrument as with a “sightglass” level gauge (see section 20.1.2 beginning on page 1408 for a detailed explanation of this problem.).


Calculating buoyant force on a displacer element due to a combination of two liquids is not as difficult as it may sound. Archimedes’ Principle still holds: that buoyant force is equal to the weight of the fluid(s) displaced. All we need to do is calculate the combined weights and volumes of the displaced liquids to calculate buoyant force. For a single liquid, buoyant force is equal to the weight density of that liquid (γ) multiplied by the volume displaced (V):

$$F_{buoyant} = \gamma V$$

For a two-liquid interface, the buoyant force is equal to the sum of the two liquid weights displaced, each liquid weight term being equal to the weight density of that liquid multiplied by the displaced volume of that liquid:

$$F_{buoyant} = \gamma_1 V_1 + \gamma_2 V_2$$

Assuming a displacer of constant cross-sectional area throughout its length, the volume for each liquid's displacement is simply equal to the same area (πr^2) multiplied by the length of the displacer submerged in that liquid:




$$F_{buoyant} = \gamma_1 \pi r^2 l_1 + \gamma_2 \pi r^2 l_2$$

Since the area (πr^2) is common to both buoyancy terms in this equation, we may factor it out for simplicity's sake:

$$F_{buoyant} = \pi r^2 (\gamma_1 l_1 + \gamma_2 l_2)$$

Determining the calibration points of a displacer-type level instrument for interface applications is relatively easy if the LRV and URV conditions are examined as a pair of “thought experiments” just as we did with hydrostatic interface level measurement. First, we imagine what the displacer’s condition would “look like” with the interface at the lower range value, then we imagine a different scenario with the interface at the upper range value. Sketching illustrations of each scenario is recommended for clarity.

Suppose we have a displacer instrument measuring the interface level between two liquids having specific gravities of 0.850 and 1.10, with a displacer length of 30 inches and a displacer diameter of 2.75 inches (radius = 1.375 inches). Let us further suppose the LRV in this case is where the interface is at the displacer’s bottom and the URV is where the interface is at the displacer’s top. The placement of the LRV and URV interface levels at the extreme ends of the displacer’s length simplifies our LRV and URV calculations, as the LRV “thought experiment” will simply be the displacer completely submerged in light liquid and the URV “thought experiment” will simply be the displacer completely submerged in heavy liquid.



Calculating the LRV buoyant force:

$$F_{buoyant} \text{ (LRV)} = \gamma_2 V = \gamma_2 \pi r^2 l$$

Calculating the URV buoyant force:

$$F_{buoyant} \text{ (URV)} = \gamma_1 V = \gamma_1 \pi r^2 l$$

Showing the actual calculations for this hypothetical example:

$$\gamma_1 = \left(62.4 \frac{\text{lb}}{\text{ft}^3} \right) (1.10) = 68.6 \frac{\text{lb}}{\text{ft}^3} = 0.0397 \frac{\text{lb}}{\text{in}^3}$$

$$\gamma_2 = \left(62.4 \frac{\text{lb}}{\text{ft}^3} \right) (0.85) = 53.0 \frac{\text{lb}}{\text{ft}^3} = 0.0307 \frac{\text{lb}}{\text{in}^3}$$

$$F_{buoyant} \text{ (LRV)} = \left(0.0307 \frac{\text{lb}}{\text{in}^3} \right) \pi (1.375 \text{ in})^2 (30 \text{ in}) = 5.47 \text{ lb}$$

$$F_{buoyant} \text{ (URV)} = \left(0.0397 \frac{\text{lb}}{\text{in}^3} \right) \pi (1.375 \text{ in})^2 (30 \text{ in}) = 7.08 \text{ lb}$$

The buoyancy for any measurement percentage between the LRV (0%) and URV (100%) may be calculated by interpolation:

Interface level (inches)	Buoyant force (pounds)
0	5.47
7.5	5.87
15	6.27
22.5	6.68
30	7.08

20.5 Echo

A completely different way of measuring liquid level in vessels is to bounce a traveling wave off the surface of the liquid – typically from a location at the top of the vessel – using the time-of-flight for the waves as an indicator of distance¹⁸, and therefore an indicator of liquid height inside the vessel. Echo-based level instruments enjoy the distinct advantage of immunity to changes in liquid density, a factor crucial to the accurate calibration of hydrostatic and displacement level instruments. In this regard, they are quite comparable with float-based level measurement systems.

From a historical perspective, hydrostatic and displacement level instruments have a richer pedigree. These instruments are simpler in nature than echo-based instruments, and were practical long before the advent of modern electronic technology. Echo-based instruments require precision timing and wave-shaping circuitry, plus sensitive (and rugged!) transceiver elements, demanding a much higher level of technology. However, modern electronic design and instrument manufacturing practices are making echo-based level instruments more and more practical for industrial applications. At the time of this writing (2008), it is common practice in some industries to replace old displacer level instruments with guided-wave radar instruments, even in demanding applications operating at high pressures¹⁹.

Liquid-liquid interfaces may also be measured with some types of echo-based level instruments, most commonly guided-wave radar.

The single most important factor to the accuracy of any echo-based level instrument is the speed at which the wave travels en route to the liquid surface and back. This wave propagation speed is as fundamental to the accuracy of an echo instrument as liquid density is to the accuracy of a hydrostatic or displacer instrument. So long as this velocity is known and stable, good level measurement accuracy is possible. Although it is true that the calibration of an echo-based level instrument does not depend on process fluid density for the reason it does in hydrostatic- or displacement-based level instruments, this does not necessarily mean the calibration of an echo-based level instrument remains fixed as process fluid density changes. The propagation velocity of the wave used in an echo-based level instrument may indeed be subject to change as the process fluids change temperature or composition. For ultrasonic (sound) echo instruments, the speed of sound is a function of medium density. Thus, an ultrasonic level transmitter measuring time-of-flight through a vapor above the liquid may drift out of calibration if the speed of sound through that vapor changes substantially, which may happen if the vapor's temperature or pressure happens to change. If the sound wave time-of-flight is measured while the waves pass through liquid, the calibration may drift if the speed of sound in that liquid changes substantially, which may happen if the liquid's temperature changes. For radar (radio wave) echo instruments, the speed of radio wave propagation varies according to the dielectric permittivity of the medium. Permittivity is also affected by changes in density for the fluid medium, and so even radar level instruments may suffer calibration drift with process fluid density changes.

To summarize these effects, the speed of sound through any medium is a function of density and bulk modulus (the “compressibility” of the medium), with density generally being the more variable of the two. For gases and vapors, this means the speed of sound is strongly affected by changes in

¹⁸The general term for this form of measurement is *time domain reflectometry*.

¹⁹My own experience with this trend is within the oil refining industry, where legacy displacer instruments (typically Fisher brand “Level-Trol” units) are being replaced with new guided-wave radar transmitters, both for vapor-liquid and liquid-liquid interface applications.

gas pressure and/or gas temperature. For liquids, this means the speed of sound is strongly affected by temperature. For solids, this means the speed of sound is weakly affected by temperature. The degree to which the speed of sound will be affected by temperature changes is directly related to the degree the medium's density changes with temperature: solid materials generally expand and contract less than liquids over the same temperature range, thus the strong temperature effect for liquids and the weak temperature effect for solids.


Radio wave velocity is a function of dielectric permittivity, which is also a function of density. However, the degree of change in dielectric permittivity resulting from changes in pressure and/or temperature are generally much less than the degree of change in speed of sound for the same media and the same changes in pressure and/or temperature. This means that – all other factors being equal – an echo-based level instrument using radio waves will suffer far less calibration error than an echo-based level instrument using sound waves when process fluid pressure and/or temperature change. However, it should be noted that process fluid *composition* (i.e. its chemical make-up) may have a strong effect on radio wave propagation, not just on its time-of-flight but also on its ability to produce an adequate echo at the interface between two fluids.

Echo-based level instruments may also be “fooled” by layers of foam resting on top of the liquid, and the liquid-to-liquid interface detection models may have difficulty detecting non-distinct interfaces (such as emulsions). Irregular structures residing within the vapor space of a vessel (such as access portals, mixer paddles and shafts, ladders, etc.) may wreak havoc with echo-based level instruments by casting false echoes back to the instrument, although this problem may be mitigated by installing guide tubes for the waves to travel in, or using wave probes as in the cases of guided-wave radar instruments. Liquid streams pouring in to the vessel through the vapor space may similarly cause problems for an echo instrument. Additionally, all echo-based instruments have *dead zones* where liquid level is too close to the transceiver to be accurately measured or even detected (the echo time-of-flight being too short for the receiving electronics to distinguish from the incident pulse).

As you can see, echo-based level instruments have strengths and weaknesses just like any other type of level instrument. There is no “perfect” level instrument, but rather a wide array of choices from which the end-user must judiciously select for the particular application in mind. Beware of sales pitches urging you to buy the “perfect” level meter! The wise approach is to first research the underlying physics of the instrument, then determine how strongly its accuracy will be affected by realistic changes in process conditions (e.g. pressure, temperature, composition).

20.5.1 Ultrasonic level measurement

Ultrasonic level instruments measure the distance from the transmitter (located at some high point) to the surface of a process material located farther below using reflected sound waves. The frequency of these waves extend beyond the range of human hearing, which is why they are called *ultrasonic*. The time-of-flight for a sound pulse indicates this distance, and is interpreted by the transmitter electronics as process level. These transmitters may output a signal corresponding either to the fullness of the vessel (*fillage*) or the amount of empty space remaining at the top of a vessel (*ullage*).



Ullage is the “natural” mode of measurement for this sort of level instrument, because the sound wave’s time-of-flight is a direct function of how much empty space exists between the liquid surface and the top of the vessel. Total tank height will always be the sum of fillage and ullage, though. If the ultrasonic level transmitter is programmed with the vessel’s total height, it may calculate fillage via simple subtraction:


$$\text{Fillage} = \text{Total height} - \text{Ullage}$$

If a sound wave encounters a sudden change in the material’s speed of sound, some of that wave’s energy will be reflected in the form of another wave in the opposite direction. In other words, the sound wave will “echo” when it encounters a material having a different sonic velocity²⁰. This is the basis of all ultrasonic ranging devices. Thus, in order for an ultrasonic level transmitter to function reliably, the difference in sonic velocities at the interface between liquid and gas must be large. Distinct interfaces of liquid and gas almost always exhibit huge differences in their speeds of sound, and so are relatively easy to detect using ultrasonic waves. Liquids with a heavy layer of foam floating on top are more difficult, since the foam is less dense than the liquid, but considerably

²⁰The speed of sound through any substance is a function of both the substance’s density and its bulk modulus (i.e. the compressibility of a substance). Mathematically, $c = \sqrt{\frac{B}{\rho}}$ where c is the sonic velocity, B is the bulk modulus, and ρ is the mass density. Water and air provide an excellent illustration of this principle: the speed of sound through water happens to be much faster than the speed of sound through air despite the vastly greater mass density of water, only because of the even greater disparity in bulk modulus between water and air.

denser than the gas above. A weak echo will be generated at the interface of foam and gas, and another generated at the interface of liquid and foam, with the foam acting to scatter and dissipate much of the second echo's energy.


The instrument itself consists of an electronics module containing all the power, computation, and signal processing circuits; plus an ultrasonic transducer²¹ to send and receive the sound waves. This transducer is typically piezoelectric in nature, being the equivalent of a very high-frequency audio speaker. The following photographs show a typical electronics module (left) and sonic transducer (right):



The ISA-standard designations for each component would be “LT” (level transmitter) for the electronics module and “LE” (level element) for the transducer, respectively. Even though we call the device responsible for transmitting and receiving the sound waves a *transducer* (in the scientific sense of the word), its function as a process instrument is to be the *primary sensing element* for the level measurement system, and therefore it is more properly designated a “level element” (LE).


²¹In the industrial instrumentation world, the word “transducer” usually has a very specific meaning: a device used to process or convert standardized instrumentation signals, such as 4-20 mA converted into 3-15 PSI, etc. In the general scientific world, however, the word “transducer” describes any device converting one form of energy into another. It is this latter definition of the word that I am using when I describe an ultrasonic “transducer” – a device used to convert electrical energy into ultrasonic sound waves, and vice-versa.

This photograph shows a typical installation for an ultrasonic level-sensing element (LE), here sensing the level of wastewater in an open channel:



Electrical conduit serves to protect the signal cable from exposure to the elements as it routes back to an indoor location where the level transmitter (LT) is located.

If the ultrasonic transducer is rugged enough, and the process vessel sufficiently free of sludge and other sound-damping materials accumulating at the vessel bottom, the transducer may be mounted at the bottom of the vessel, bouncing sound waves off the liquid surface through the liquid itself rather than through the vapor space. As stated previously, any significant difference in sonic velocity between the two materials is sufficient to reflect a sound wave. This being the case, it shouldn't matter which material the incident sound wave propagates through *first*:




This arrangement makes fillage the natural measurement, and ullage a derived measurement (calculated by subtraction from total vessel height).

$$\text{Ullage} = \text{Total height} - \text{Fillage}$$

As mentioned previously, the calibration of an ultrasonic level transmitter depends on the speed of sound through the medium between the transducer and the interface. For top-mounted transducers, this is the speed of sound through the air (or vapor) over the liquid, since this is the medium through which the incident and reflected wave travel time is measured. For bottom-mounted transducers, this is the speed of sound through the liquid. In either case, to ensure good accuracy, one must make sure the speed of sound through the “timed” travel path remains reasonably constant (or else compensate for changes in the speed of sound through that medium by use of temperature or pressure measurements and a compensating algorithm).

Ultrasonic level instruments enjoy the advantage of being able to measure the height of solid materials such as powders and grains stored in vessels, not just liquids. Again, the fundamental criterion for detecting a level of material is that the speeds of sound through the upper and lower materials must differ (the greater the difference, the stronger the echo). A unique challenge to solids measurement is the distinct possibility of uneven material profiles. A classic problem encountered when measuring the level of a powdered or granular material in a vessel is the *angle of repose* formed by the material as a result of being fed into the vessel at one point:




This angled surface is difficult for an ultrasonic device to detect because it tends to scatter the sound waves laterally instead of reflecting them strongly back toward the instrument. However, even if the scattering problem is not significant, there still remains the problem of interpretation: what is the instrument actually measuring? The detected level near the vessel wall will certainly register less than at the center, but the level detected mid-way between the vessel wall and vessel center may not be an accurate average of those two heights. Moreover, this angle may decrease over time if mechanical vibrations cause the material to “flow” and tumble from center to edge.

For this reason, solids storage measurement applications demanding high accuracy generally use other techniques, such as weight-based measurement (see section 20.6 for more information) or three-dimensional scanning (see section 26.3 for more information).

20.5.2 Radar level measurement

*Radar*²² level instruments measure the distance from the transmitter (located at some high point) to the surface of a process material located farther below in much the same way as ultrasonic transmitters – by measuring the time-of-flight of a traveling wave. The fundamental difference between a radar instrument and an ultrasonic instrument is the type of wave used: radio waves instead of sound waves. Radio waves are electromagnetic in nature (comprised of alternating electric and magnetic fields), and very high frequency (in the microwave frequency range – GHz). Sound waves are *mechanical* vibrations (transmitted from molecule to molecule in a fluid or solid substance) and of much lower frequency (tens or hundreds of kilohertz – still too high for a human being to detect as a tone) than radio waves.

Some radar level instruments use waveguide “probes” to guide the electromagnetic waves to and from the process liquid while others send electromagnetic waves out through open space to reflect off the process material. The instruments using waveguides are called *guided-wave radar* instruments, whereas the radar instruments relying on open space for signal propagation are called *non-contact radar*. The differences between these two varieties of radar instruments is shown in the following illustration:



²²“Radar” is an acronym: RAdio Detection And Ranging. First used as a method for detecting enemy ships and aircraft at long distances over the ocean in World War II, this technology is used for detecting the presence, distance, and/or speed of objects in a wide variety of applications.


Photographs of non-contact (left) and guided-wave (right) radar level transmitters are shown below. The non-contact transmitter is placed on a table for inspection while the guided-wave transmitter is installed in a “cage”²³ similar to that of a displacement-style level transmitter attached to the vessel by two pipes:



Non-contact radar devices suffer much more signal loss than guided-wave radar devices, due to the natural tendency of electromagnetic radiation to disperse over space. Waveguides combat this signal loss by channeling the radio energy along a straight-line path. Probes used in guided-wave radar instruments may be single metal rods, parallel pairs of metal rods, or a coaxial metal rod-and-tube structure. Single-rod probes suffer the greatest energy losses, while coaxial probes excel at guiding the microwave energy to the liquid interface and back. However, single-rod probes are much more tolerant of process fouling than two-rod or (especially) coaxial probes, where sticky masses of viscous liquid and/or solid matter cling to the probe. Such fouling deposits, if severe enough, will cause electromagnetic wave reflections that “look” to the transmitter like the reflection from an actual liquid level or interface.

²³In fact, it is a common retrofit practice to install a guided-wave radar level transmitter in the exact same cage that once housed a displacement-style level transmitter.

Non-contact radar instruments rely on antennas to direct microwave energy into the vessel, and to receive the echo (return) energy. These antennas must be kept clean and dry, which may be a problem if the liquid being measured emits condensable vapors. For this reason, non-contact radar instruments are often separated from the vessel interior by means of a *dielectric window* (made of some substance such as plastic that is relatively “transparent” to electromagnetic waves yet acts as an effective vapor barrier):



Electromagnetic waves travel at the speed of light²⁴, 2.9979×10^8 meters per second in a perfect vacuum. The velocity of an electromagnetic wave through space depends on the dielectric permittivity (symbolized by the Greek letter “epsilon,” ϵ) of that space. A formula relating wave velocity (v) to relative permittivity (the ratio of a substance’s permittivity to that of a perfect vacuum, symbolized as ϵ_r and sometimes called the *dielectric constant* of the substance) and the speed of light in a perfect vacuum (c) is shown here²⁵:

$$v = \frac{c}{\sqrt{\epsilon_r}}$$

²⁴In actuality, both radio waves and light waves are electromagnetic in nature. The only difference between the two is frequency: while the radio waves used in radar systems are classified as “microwaves” with frequencies in the gigahertz (GHz) region, visible light waves range in the hundred of terahertz (THz)!

²⁵This formula assumes lossless conditions: that none of the wave’s energy is converted to heat while traveling through the dielectric. For many situations, this is true enough to assume.

As mentioned previously, the calibration of any echo-based level transmitter depends on knowing the speed of wave propagation through the medium separating the instrument from the process fluid interface. For radar transmitters sensing a single liquid below a gas or vapor, this speed is the speed of light through that gas or vapor space, which we know to be a function of electrical permittivity.

The relative permittivity of air at standard pressure and temperature is very nearly unity (1). This means the speed of light in air under atmospheric pressure and ambient temperature will very nearly be the same as it is for a perfect vacuum (2.9979×10^8 meters per second). If, however, the vapor space above the liquid is not ambient air, and is subject to large changes in temperature and/or pressure²⁶ which cause the vapor's density to change, the permittivity of that vapor may substantially change and consequently skew the speed of light, and therefore the calibration of the level instrument. This calibration shift is sometimes referred to as the *gas phase effect*.

A formula useful for calculating the permittivity of any gas or vapor based on both pressure and temperature is shown here²⁷:

$$\epsilon_r = 1 + (\epsilon_{ref} - 1) \frac{PT_{ref}}{P_{ref}T}$$

Where,

ϵ_r = Relative permittivity of a gas at a given pressure (P) and temperature (T)

ϵ_{ref} = Relative permittivity of the same gas at standard pressure (P_{ref}) and temperature (T_{ref})

P = Absolute pressure of gas (bars)

P_{ref} = Absolute pressure of gas under standard conditions (≈ 1 bar)

T = Absolute temperature of gas (Kelvin)

T_{ref} = Absolute temperature of gas under standard conditions (≈ 273 K)


This formula is based on the principle that bulk permittivity is a function of density. We may see why this is by running a "thought experiment" in which a sample of gas becomes denser. As gas density increases, more gas molecules will become packed into the same volume of space. If each gas molecule's permittivity is greater than the permittivity of empty space, then having more of those gas molecules present will mean the permittivity of that volume increases. Greater permittivity, of course, decreases the velocity of light through the gas, and thereby affects the calibration of the radar instrument.

Relating this concept to pressure and temperature variations in the gas, we can see that the permittivity of a gas increases with increasing pressure (by increasing gas density), and decreases with increasing temperature (by decreasing gas density). This means the speed of light through a gas decreases with increasing pressure, and increases with increasing temperature. For radar level instruments operating in gas environments subject to significant pressure and temperature (i.e. density) variations, the consequent variations in the speed of light through that gas will compromise the instrument's accuracy.

²⁶Or if the chemical composition of the gas or vapor changes dramatically.

²⁷The pressure and temperature factors in this formula come from the Ideal Gas Law ($PV = nRT$), manipulating that equation to express molecular gas density in terms of pressure and temperature ($\rho = \frac{n}{V} = \frac{P}{RT}$). The fraction $\frac{PT_{ref}}{P_{ref}T}$ expresses a ratio of molecular densities: $\frac{\rho}{\rho_{ref}}$.

With ultrasonic level instruments, the necessary condition for an echo to occur is that the sound wave encounters a sudden change in material density. With radar level instruments, the necessary condition for wave reflection is a sudden change in dielectric permittivity²⁸ (ϵ). When an electromagnetic wave encounters a sudden change in dielectric permittivity, some of that wave's energy will be reflected in the form of another wave traveling the opposite direction, while the balance of the wave's energy continues forward to propagate into the new material. The strength of the reflected signal depends on how greatly the two materials' permittivities differ:



This same principle explains reflected signals in copper transmission lines as well. Any discontinuities (sudden changes in characteristic impedance) along the length of a transmission line will reflect a portion of the electrical signal's power back to the source. In a transmission line, continuities may be formed by pinches, breaks, or short-circuits. In a radar level measurement system, any sudden change in electrical permittivity is a discontinuity that reflects some of the incident wave energy back to the source. Thus, radar level instruments function best when there is a large difference in permittivity between the two substances at the interface. As shown in the previous illustration, air and water meet this criterion, having an 80:1 permittivity ratio.

²⁸Once again, it is closer to the truth to say that radio waves reflect when they encounter any sudden change in the *velocity of light*. Light and radio waves are both oscillating electromagnetic fields, and so they propagate in the same ways. Dielectric permittivity is one of the factors determining the speed of light through a substance, but not the only one. The material's *magnetic permeability* is another factor, but it is far more common to encounter interfaces of gas-liquid or liquid-liquid where differences in permittivity rather than differences in permeability constitute the major reason for differences in radio wave velocity.

The ratio of reflected power to incident (transmitted) power at any interface of materials is called the *power reflection factor* (R). This may be expressed as a unitless ratio, or more often as a decibel figure. The relationship between dielectric permittivity and reflection factor is as follows:

$$R = \frac{(\sqrt{\epsilon_{r2}} - \sqrt{\epsilon_{r1}})^2}{(\sqrt{\epsilon_{r2}} + \sqrt{\epsilon_{r1}})^2}$$

Where,

R = Power reflection factor at interface, as a unitless ratio

ϵ_{r1} = Relative permittivity (dielectric constant) of the first medium

ϵ_{r2} = Relative permittivity (dielectric constant) of the second medium

The fraction of incident power transmitted through the interface ($\frac{P_{forward}}{P_{incident}}$) is, of course, the mathematical complement of the power reflection factor: $1 - R$.

For situations where the first medium is air or some other low-permittivity gas, the formula simplifies to the following form (with ϵ_r being the relative permittivity of the reflecting substance):

$$R = \frac{(\sqrt{\epsilon_r} - 1)^2}{(\sqrt{\epsilon_r} + 1)^2}$$

In the previous illustration, the two media were air ($\epsilon_r \approx 1$) and water ($\epsilon_r \approx 80$) – a nearly ideal scenario for strong signal reflection. Given these relative permittivity values, the power reflection factor has a value of 0.638 (63.8%), or -1.95 dB. This means well over half the incident power reflects off the air/water interface to form a strong echo signal, with the remaining 0.362 (36.2%) of the wave’s power traveling through the air-water interface and propagating into water. If the liquid in question is gasoline rather than water (having a rather low relative permittivity value of approximately 2), the power reflection ratio will only be 0.0294 (2.94%) or -15.3 dB, with the vast majority of the wave’s power successfully penetrating the air-gasoline interface.


The longer version of the power reflection factor formula suggests liquid-liquid interfaces should be detectable using radar, and indeed they are. All that is needed is a sufficiently large difference in permittivity between the two liquids to create a strong enough echo to reliably detect. Liquid-liquid interface level measurement with radar works best when the upper liquid has a substantially lesser permittivity value than the lower liquid²⁹. A layer of hydrocarbon oil on top of water (or any aqueous solution such as an acid or a caustic) is a good candidate for guided-wave radar level measurement. An example of a liquid-liquid interface that would be very difficult for a radar instrument to detect is water ($\epsilon_r \approx 80$) above glycerin ($\epsilon_r \approx 42$).

If the radar instrument uses a digital network protocol to communicate information with a host system (such as HART or any number of “fieldbus” standards), it may perform as a multi-variable transmitter, transmitting *both* the interface level measurement and the total liquid level measurement simultaneously. This capability is rather unique to guided-wave radar transmitters, and is very useful in some processes because it eliminates the need for multiple instruments measuring multiple levels.

²⁹Rosemount’s “Replacing Displacers with Guided Wave Radar” technical note states that the difference in dielectric constant between the upper and lower liquids *must* be at least 10.

One reason why a lesser- ϵ fluid above a greater- ϵ fluid is easier to detect than the inverse is due to the necessity of the signal having to travel through a gas-liquid interface above the liquid-liquid interface. With gases and vapors having such small ϵ values, the signal would have to pass through the gas-liquid interface first in order to reach the liquid-liquid interface. This gas-liquid interface, having the greatest difference in ϵ values of any interface within the vessel, will be *most* reflective to electromagnetic energy *in both directions*. Thus, only a small portion of the incident wave will ever reach the liquid-liquid interface, and a similarly small portion of the wave reflected off the liquid-liquid interface (which itself is a fraction of the forward wave power that made it through the gas-liquid interface on its way down) will ever make it through the gas-liquid interface on its way *back up* to the instrument. The situation is much improved if the ϵ values of the two liquid layers are inverted, as shown in this hypothetical comparison (all calculations³⁰ assume no power dissipation along the way, only reflection at the interfaces):

Signal power strengths en route and reflected off of the liquid-liquid interface



As you can see in the illustration, the difference in power received back at the instrument is nearly two to one, just from the upper liquid having the lesser of two identical ϵ values. Of course, in real life you do not have the luxury of *choosing* which liquid will go on top of the other (this being determined by fluid density), but you do have the luxury of choosing the appropriate liquid-liquid interface level measurement technology, and as you can see here certain orientations of ϵ values are less detectable with radar than others.

Another factor working against radar as a liquid-liquid interface measurement technology for interfaces where the upper liquid has a greater dielectric constant is that fact that many high- ϵ liquids are aqueous in nature, and water readily dissipates microwave energy. This fact is exploited in microwave ovens, where microwave radiation excites water molecules in the food, dissipating energy in the form of heat. For a radar-based level measurement system consisting of gas/vapor over water over some other (heavier) liquid, the echo signal will be extremely weak because the signal must pass through the “lossy” water layer *twice* before it returns to the radar instrument.

³⁰ $R = 0.5285$ for the 1/40 interface; $R = 0.02944$ for the 40/80 interface; and $R = 0.6382$ for the 1/80 interface, all based on the formula $R = \frac{(\sqrt{\epsilon_r} - 1)^2}{(\sqrt{\epsilon_r} + 1)^2}$ using the pair of permittivity values at each interface.

Electromagnetic energy losses are important to consider in radar level instrumentation, even when the detected interface is simply gas (or vapor) over liquid. The power reflection factor formula only predicts the ratio of reflected power to incident power *at an interface of substances*. Just because an air-water interface reflects 63.8% of the incident power does not mean 63.8% of the incident power will actually return to the transceiver antenna! Any dissipative losses between the transceiver and the interface(s) of concern will weaken the signal, to the point where it may become difficult to distinguish from noise.

Another important factor in maximizing reflected power is the degree to which the microwaves disperse on their way to the liquid interface(s) and back to the transceiver. Guided-wave radar instruments receive a far greater percentage of their transmitted power than non-contact radar instruments because the metal probe used to guide the microwave signal pulses help prevent the waves from spreading (and therefore weakening) throughout the liquids as they propagate. In other words, the probe functions as a transmission line to direct and focus the microwave energy, ensuring a straight path from the instrument into the liquid, and a straight echo return path from the liquid back to the instrument. This is why guided-wave radar is the only practical radar technology for measuring liquid-liquid interfaces.

A critically important factor in accurate level measurement using radar instruments is that the dielectric permittivity of every substance lying between the radar instrument and the interface of interest be accurately known. The reason for this is rooted in the dependence of electromagnetic wave propagation velocity to relative permittivity. Recalling the wave velocity formula shown earlier:

$$v = \frac{c}{\sqrt{\epsilon_r}}$$

Where,

- v = Velocity of electromagnetic wave through a particular substance
- c = Speed of light in a perfect vacuum ($\approx 3 \times 10^8$ meters per second)
- ϵ_r = Relative permittivity (dielectric constant) of the substance

In the case of a single-liquid application where nothing but gas or vapor exists above the liquid, the permittivity of that gas or vapor must be precisely known. In the case of a two-liquid interface with gas or vapor above, the relative permittivities of *both* gas and upper liquids must be accurately known in order to accurately measure the liquid-liquid interface. Changes in dielectric constant value of the medium or media through which the microwaves must travel and echo will cause the microwave radiation to propagate at different velocities. Since all radar measurement is based on time-of-flight through the media separating the radar transceiver from the echo interface, changes in wave velocity through this media will affect the amount of time required for the wave to travel from the transceiver to the echo interface, and reflect back to the transceiver. Therefore, changes in dielectric constant are relevant to the accuracy of any radar level measurement³¹.

Factors influencing the dielectric constant of gases include pressure and temperature, which means the accuracy of a radar level instrument will vary as gas pressure and/or gas temperature vary! This is often referred to as the *gas phase effect*. Whether or not this variation is substantial

³¹It should be noted that the dielectric constant of the lowest medium (the liquid in a simple, non-interface, level measurement application) is irrelevant for calibration purposes. All we are concerned with is the propagation time of the signal to and from the level of interest, nothing below it.


enough to consider for any application depends on the desired measurement accuracy and the degree of permittivity change from one pressure/temperature extreme to another. In no case should a radar instrument be considered for any level measurement application unless the dielectric constant value(s) of the upper media are precisely known. This is analogous to the dependence on liquid density that hydrostatic level instruments face. It is futile to attempt level measurement based on hydrostatic pressure if liquid density is unknown or widely varying, and it is just as futile to attempt level measurement based on radar if the dielectric constants are unknown³² or varies widely.

One way to compensate for the gas phase effect in radar level instruments is to equip the instrument with a *reference probe* of fixed length oriented in such a way that its entire length is always above the liquid level (i.e. it only senses gas). If the permittivity of the gas is constant, the echo time along this reference probe will remain the same. If, however, the gas permittivity changes, the reference probe's echo time will correspondingly change, allowing the instrument's microprocessor to measure gas permittivity and consequently adjust calculations for liquid level based on this known change. This concept is analogous to the *compensating probe* sometimes used in capacitive level sensors, designed to measure fluid permittivity so as to compensate for any changes in this critical parameter.

As with ultrasonic level instruments, radar level instruments can sense the level of solid substances in vessels (e.g. powders and granules) and not just liquids. The same caveat of repose angle applicable to ultrasonic level measurement (see section 20.5.1 beginning on page 1472), however, is a factor for radar measurement as well. Also, low particulate solid density (i.e. significant amounts of air between the solid particles) tends to reduce the material's dielectric constant and thereby weaken the radar echo.

³²For vented-tank level measurement applications where air is the only substance above the point of interest, the relative permittivity is so close to a value of 1 that there is little need for further consideration on this point. Where the permittivity of fluids becomes a problem for radar is in high-pressure (non-air) gas applications and liquid-liquid interface applications, especially where the upper substance composition is subject to change.


Modern radar level instruments provide a wealth of diagnostic information to aid in troubleshooting. One of the most informative is the *echo curve*, showing each reflected signal received by the instrument along the incident signal's path of travel. The following image is a screen capture of a computer display, from software used to configure a Rosemount model 3301 guided-wave radar level transmitter with a coaxial probe:



To view a flip-book animation showing how a guided-wave radar (GWR) instrument detects both liquid surface level and liquid-liquid interface level, turn to Appendix A.5 beginning on page 3043.

Pulse P1 is the *reference* or *fiducial* pulse, resulting from the change in dielectric permittivity between the extended “neck” of the probe (connecting the transmitter to the probe tube) and the coaxial probe itself. This pulse marks the top of the probe, thereby establishing a point of reference for ullage measurement.

This next screen capture shows the same level transmitter measuring a water level that is 8 inches higher than before. Note how pulse P2 is further to the left (indicating an echo received sooner in time), indicating a lesser ullage (greater level) measurement:



Several *threshold* settings determine how the transmitter categorizes each received pulse. Threshold T1 for this particular radar instrument defines which pulse is the reference (fiducial). Thus, the first echo in time to exceed the value of threshold T1 is interpreted by the instrument to be the reference point. Threshold T2 defines the upper product level, so the first echo in time to exceed this threshold value is interpreted as the vapor/liquid interface point. Threshold T3 for this particular transmitter is used to define the echo generated by a liquid-liquid interface. However, threshold T3 does not appear in this echo plot because the interface measurement option was disabled during this experiment. The last threshold, T4, defines the end-of-probe detection. Set at a negative value (just like the reference threshold T1), threshold T4 looks for the first pulse in time to exceed that value and interprets that pulse as the one resulting from the signal reaching the probe's end.


All along the echo curve you can see weak echo signals showing up as bumps. These echoes may be caused by discontinuities along the probe (solid deposits, vent holes, centering spacers, etc.), discontinuities in the process liquid (suspended solids, emulsions, etc.), or even discontinuities in the surrounding process vessel (for non-coaxial probes which exhibit varying degrees of sensitivity to surrounding objects). A challenge in configuring a radar level transmitter is to set the threshold values such that “false” echoes are not interpreted as real liquid or interface levels.

A simple way to eliminate false echoes near the reference point is to set a *null zone* where any echoes are ignored. The upper null zone (UNZ) setting on the Rosemount 3301 radar level transmitter whose screen capture image was shown previously was set to zero, meaning it would be sensitive to any and all echoes near the reference point. If a false echo from a tank nozzle or some other discontinuity near the probe's entry point into the process vessel created a measurement problem, the upper null zone (UNZ) value could be set just beyond that point so the false echo would not be interpreted as a liquid level echo, regardless of the threshold value for T2. A “null zone” is sometimes referred to as a *hold-off distance*.

Some radar level instruments allow thresholds to be set as curves themselves rather than straight lines. Thus, thresholds may be set high during certain periods along the horizontal (time/distance)

axis to ignore false echoes, and set low during other periods to capture legitimate echo signals.

Regardless of how null zones and thresholds are set for any guided-wave radar level transmitter, the technician must be aware of *transition zones* near each end of the probe. Measurements of liquid level or interface level within these zones may not be accurate or even linearly responsive. Thus, it is strongly advised to range the instrument in such a way that the lower- and upper-range values (LRV and URV) lie between the transition zones:



The size of these transition zones depends on both the process substances and the probe type³³. The instrument manufacturer will provide you with appropriate data for determining transition zone dimensions.

³³Probe mounting style will also influence the lower transition zone, in the case of flexible probes anchored to the bottom of the process vessel.

20.5.3 Laser level measurement

The least-common form of echo-based level measurement is *laser*, which uses pulses of laser light reflected off the surface of a liquid to detect the liquid level. Perhaps the most limiting factor with laser measurement is the necessity of having a sufficiently reflective surface for the laser light to “echo” off of. Many liquids are not reflective enough for this to be a practical measurement technique, and the presence of dust or thick vapors in the space between the laser and the liquid will disperse the light, weakening the light signal and making the level more difficult to detect.

However, lasers have been applied with great success in measuring distances between objects. Applications of this technology include motion control on large machines, where a laser points at a moving reflector, the laser’s electronics calculating distance to the reflector based on the amount of time it takes for the laser “echo” to return. The advent of mass-produced, precision electronics has made this technology practical and affordable for many applications. At the time of this writing (2008), it is even possible for the average American consumer to purchase laser “tape measures” for use in building construction.

20.5.4 Magnetostrictive level measurement

A variation on the theme of echo-based level instruments, where the level of some process material in a vessel is measured by timing the travel of a wave between the instrument and the material interface, is one applied to float-type instruments: *magnetostriction*.


In a magnetostrictive level instrument, liquid level is sensed by a lightweight, donut-shaped float containing a magnet. This float is centered around a long metal rod called a *waveguide*, hung vertically in the process vessel (or hung vertically in a protective cage like the type used for displacement-style level instruments) so that the float may rise and fall with process liquid level. The magnetic field from the float's magnet at that point, combined with the magnetic field produced by an electric current pulse periodically sent through the rod, generates a torsional stress pulse³⁴ at the precise location of the float. This torsional (twisting) stress travels at the speed of sound through the rod toward either end. At the bottom end is a dampener device designed to absorb the mechanical wave³⁵.

One might argue that a magnetostrictive instrument is not an “echo” technology in the strictest sense of the word. Unlike ultrasonic, radar, and laser instruments, we are not reflecting a wave off a discontinuous interface between materials. Instead, a mechanical wave (pulse) is *generated* at the location of a magnetic float in response to an electrical pulse. However, the principle of measuring distance by the wave's travel time is the same. At the top end of the rod (above the process liquid level) is a sensor and electronics package designed to detect the arrival of the mechanical wave. A precision electronic timing circuit measures the time elapsed between the electric current pulse (called the *interrogation pulse*) and the received mechanical pulse. So long as the speed of sound through the metal waveguide rod remains fixed, the time delay is strictly a function of distance between the float and the sensor, which we already know is called *ullage*.

³⁴An approximate analogy for understanding the nature of this pulse may be performed using a length of rope. Laying a long piece of rope in a straight line on the ground, pick up one end and quickly move it in a tight circle using a “flip” motion of your wrist. You should be able to see the torsional pulse travel down the length of the rope until it either dies out from dissipation or it reaches the rope's end. As with the torsional pulse in a magnetostrictive waveguide, this pulse in the rope is mechanical in nature: a movement of the rod's (rope's) molecules. As a mechanical wave, it may be properly understood as a form of *sound*.

³⁵This “dampener” is the mechanical equivalent of a *termination resistor* in an electrical transmission line: it makes the traveling wave “think” the waveguide is infinitely long, preventing any reflected pulses. For more information on electrical transmission lines and termination resistors, see section 5.10 beginning on page 475.

The following photograph (left) and illustration (right) show a magnetostrictive level transmitter³⁶ propped up against a classroom wall and the same style of transmitter installed in a liquid-holding vessel, respectively:




The design of this instrument is reminiscent of a guided-wave radar transmitter, where a metal *waveguide* hangs vertically into the process liquid, guiding a pulse to the sensor head where the sensitive electronic components are located. The major difference here is that the pulse is a sonic vibration traveling through the metal of the waveguide rod, not an electromagnetic pulse as is the case with radar. Like all sound waves, the torsional pulse in a magnetostriction-based level transmitter is much slower-traveling³⁷ than electromagnetic waves.

³⁶This particular transmitter happens to be one of the “M-Series” models manufactured by MTS.

³⁷One reference gives the speed of sound in a magnetostrictive level instrument as 2850 meters per second. Rounding this up to 3×10^3 m/s, we find that the speed of sound in the magnetostrictive waveguide is at least *five orders of magnitude* slower than the speed of light in a vacuum (approximately 3×10^8 m/s). This relative slowness of

It is even possible to measure liquid-liquid interfaces with magnetostrictive instruments. If the waveguide is equipped with a float of such density that it floats on the interface between the two liquids (i.e. the float is denser than the light liquid and less dense than the heavy liquid), the sonic pulse generated in the waveguide by that float's position will represent interface level. Magnetostrictive instruments may even be equipped with two floats: one to sense a liquid-liquid interface, and the other to sense the liquid-vapor interface, so that it may measure both the interface and total levels simultaneously just like a guided-wave radar transmitter:



With such an instrument, each electrical “interrogation” pulse returns *two* sonic pulses to the sensor head: the first pulse representing the total liquid level (upper, light float) and the second pulse representing the interface level (lower, heavy float). If the instrument has digital communication capability (e.g. HART, FOUNDATION Fieldbus, Profibus, etc.), both levels may be reported to the control system over the same wire pair, making it a “multivariable” instrument.

wave propagation is a good thing for our purposes here, as it gives more time for the electronic timing circuit to count, yielding a more precise measurement of distance traveled by the wave. This fact grants superior resolution of measurement to magnetostrictive level sensors over radar-based and laser-based level sensors. Open-air ultrasonic level instruments deal with propagation speeds even slower than this (principally because the bulk moduli of gases and vapors is far less than that of a solid metal rod) which at first might seem to give these level sensors the advantage in precision. However, open-air level sensors experience far greater propagation velocity variations caused by changes in pressure and temperature than magnetostrictive sensors. Unlike the speed of sound in gases or liquids, the speed of sound in a solid metal rod is very stable over a large range of process temperatures, and practically constant for a large range of process pressures. Another factor adding to the calibration stability of magnetostrictive instruments is that the composition of the medium never changes. With instruments measuring time-of-flight through process fluids, the chemical composition of those fluids often affects the wave velocity. In a magnetostrictive instrument, the waves are always traveling through the same material – the metal of the waveguide bar – and thus are not subject to variation with process changes.

Perhaps the greatest limitation of magnetostrictive level instruments is mechanical interference between the float and the rod. In order for the magnetostrictive effect to be strong, the magnet inside the float must be in close proximity to the rod. This means the inside diameter of the donut-shaped float must fit closely to the outside diameter of the waveguide. Any fouling of the waveguide's or float's surfaces by suspended solids, sludge, or other semi-solid materials may cause the float to bind and therefore not respond to changes in liquid level.

20.6 Weight

Weight-based level instruments sense process level in a vessel by directly measuring the weight of the vessel. If the vessel's empty weight (*tare weight*) is known, process weight becomes a simple calculation of total weight minus tare weight. Obviously, weight-based level sensors can measure both liquid and solid materials, and they have the benefit of providing inherently linear mass storage measurement³⁸. *Load cells* (strain gauges bonded to a steel element of precisely known modulus) are typically the primary sensing element of choice for detecting vessel weight. As the vessel's weight changes, the load cells compress or relax on a microscopic scale, causing the strain gauges inside to change resistance. These small changes in electrical resistance become a direct indication of vessel weight.

The following photograph shows three bins used to store powdered milk, each one supported by pillars equipped with load cells near their bases:



³⁸Regardless of the vessel's shape or internal structure, the measurement provided by a weight-sensing system is based on the true mass of the stored material. Unlike height-based level measurement technologies (float, ultrasonic, radar, etc.), no characterization will ever be necessary to convert a measurement of height into a measurement of mass.

A close-up photograph shows one of the load cell units in detail, near the base of a pillar:



When multiple load cells are used to measure the weight of a storage vessel, the signals from all load cell units must be added together (“summed”) to produce a signal representative of the vessel’s *total* weight. Simply measuring the weight at one suspension point is insufficient³⁹, because one can never be sure the vessel’s weight is distributed equally amongst all the supports.

³⁹If we happened to know, somehow, that the vessel’s weight *was* in fact equally shared by all supports, it would be sufficient to simply measure stress at one support to infer total vessel weight. In such an installation, assuming three supports, the total vessel weight would be the stress at any one support multiplied by three.

This next photograph shows a smaller-scale load cell installation used to measure the quantity of material fed into a beer-brewing process⁴⁰:




Weight-based measurements are often employed where the true mass of a quantity must be ascertained, rather than the level. So long as the material's density is a known constant, the relationship between weight and level for a vessel of constant cross-sectional area will be linear and predictable. Constant density is not always the case, especially for solid materials, and so weight-based inference of vessel level may be problematic.

In applications where batch mass is more important than height (level), weight-based measurement is often the preferred method for portioning batches. You will find weight-based portion measurements used frequently in the food processing industries (e.g. consistently filling bags and boxes with product), and also for custody transfer of certain materials (e.g. coal and metal ore).

⁴⁰The particular "micro-brewery" process shown here is at the Pike's Place Market in downtown Seattle, Washington. Three load cells measure the weight of a hopper filled with ingredients prior to brewing in the "mash tun" vessel.

One very important caveat for weight-based level instruments is to isolate the vessel from any external mechanical stresses generated by pipes or machinery. The following illustration shows a typical installation for a weight-based measurement system, where all pipes attaching to the vessel do so through flexible couplings, and the weight of the pipes themselves is borne by outside structures through *pipe hangers*:



Stress relief is very important because any forces acting upon the storage vessel will be interpreted by the load cells as more or less material stored in the vessel. The only way to ensure that the load cell's measurement is a direct indication of material held inside the vessel is to ensure that no other forces act upon the vessel except the gravitational weight of the material.

A similar concern for weight-based batch measurement is *vibration* produced by machinery surrounding (or on) the vessel. Vibration is nothing more than oscillatory *acceleration*, and the acceleration of any mass produces a reaction force ($F = ma$). Any vessel suspended by weight-sensing elements such as load cells will induce oscillating forces on those load cells if shaken by vibration. This concern in particular makes it quite difficult to install and operate *agitators* or other rotating machinery on a weighed vessel⁴¹.

An interesting problem associated with load cell measurement of vessel weight arises if there are ever electric currents traveling through the load cell(s). This is not a normal state of affairs,

⁴¹One practical solution to this problem is to shut down the source of vibration (e.g. agitator motor, pump, etc.) for a long enough time to take a sample weight measurement, then run the machine again between measurements. So long as intermittent weight measurement is adequate for the needs of the process, the interference of machine vibration may be dealt with in this manner.

but it can happen if maintenance workers incorrectly attach arc welding equipment to the support structure of the vessel, or if certain electrical equipment mounted on the vessel such as lights or motors develop ground faults. The electronic amplifier circuits interpreting a load cell's resistance will detect voltage drops created by such currents, interpreting them as changes in load cell resistance and therefore as changes in material level. Sufficiently large currents may even cause permanent damage to load cells, as is often the case when the currents in question are generated by arc welding equipment.

A variation on this theme is the so-called *hydraulic load cell* which is a piston-and-cylinder mechanism designed to translate vessel weight directly into hydraulic (liquid) pressure. A normal pressure transmitter then measures the pressure developed by the load cell and reports it as material weight stored in the vessel. Hydraulic load cells completely bypass the electrical problems associated with resistive load cells, but are more difficult to network for the calculation of total weight (using multiple cells to measure the weight of a large vessel).

20.7 Capacitive

Capacitive level instruments measure electrical capacitance of a conductive rod inserted vertically into a process vessel. As process level increases, capacitance increases between the rod and the vessel walls, causing the instrument to output a greater signal.

The basic principle behind capacitive level instruments is the capacitance equation:

$$C = \frac{\epsilon A}{d}$$

Where,

C = Capacitance


ϵ = Permittivity of dielectric (insulating) material between plates

A = Overlapping area of plates

d = Distance separating plates

The amount of capacitance exhibited between a metal rod inserted into the vessel and the metal walls of that vessel will vary only with changes in permittivity (ϵ), area (A), or distance (d). Since A is constant (the interior surface area of the vessel is fixed, as is the area of the rod once installed), only changes in ϵ or d can affect the probe's capacitance.


Capacitive level probes come in two basic varieties: one for conductive liquids and one for non-conductive liquids. If the liquid in the vessel is conductive, it cannot be used as the dielectric (insulating) medium of a capacitor. Consequently, capacitive level probes designed for conductive liquids are coated with plastic or some other dielectric substance, so the metal probe forms one plate of the capacitor and the conductive liquid forms the other:



In this style of capacitive level probe, the variables are permittivity (ϵ) and distance (d), since a rising liquid level displaces low-permittivity gas and essentially acts to bring the vessel wall

electrically closer to the probe. This means total capacitance will be greatest when the vessel is full (ϵ is greatest and effective distance d is at a minimum), and least when the vessel is empty (ϵ of the gas is in effect, and over a much greater distance).

If the liquid is non-conductive, it may be used as the dielectric itself, with the metal wall of the storage vessel forming the second capacitor plate. The probe is just a bare metal cable or rod:



In this style of capacitive level probe, the only variable affecting probe capacitance is permittivity (ϵ), provided the liquid has a substantially greater permittivity than the vapor space above the liquid. This means total capacitance will be greatest when the vessel is full (average permittivity ϵ is at a maximum), and least when the vessel is empty. Distance (d) is constant with a non-conducting process liquid, being the radius of the vessel (assuming the probe is mounted in the center).

Permittivity of the process substance is a critical variable in the non-conductive style of capacitance level probe, and so good accuracy may be obtained with this kind of instrument only if the process material permittivity is accurately known. A clever way to ensure good level measurement accuracy when the material's permittivity is not stable over time is to equip the instrument with a special *compensating probe* (sometimes called a *composition probe*) below the LRV point in the vessel that will always be submerged. Since this compensating probe is always immersed, and always experiences the same A and d dimensions, its capacitance is purely a function of the substance's permittivity (ϵ). This gives the instrument a way to continuously measure material permittivity, which it then uses to calculate the level of that material in the vessel based on the capacitance of the main probe. The inclusion of a compensating probe to measure and compensate for changes in permittivity is analogous to the inclusion of a third pressure transmitter in a hydrostatic *tank expert* system to continuously measure and compensate for density. It is a way to correct for changes in the one remaining system variable that is not related to changes in level.

Capacitive level instruments may be used to measure the level of solids (powders and granules) in

addition to liquids. In these applications, the material in question is almost always non-conductive, and therefore the permittivity of the substance becomes a factor in measurement accuracy. This can be problematic, as moisture content variations in the solid may greatly affect permittivity, as can variations in granule size. Compensating probes may not be very useful, either, because their location (at the bottom of the vessel) may not expose them to the same degree of material granularity and moisture content experienced by the main probe.

Capacitive level instruments are generally found in applications where precision is not important. These instruments tend to suffer from errors arising from changes in process substance permittivity, changes in process vapor-space permittivity, and errors caused by stray capacitance in probe cables.

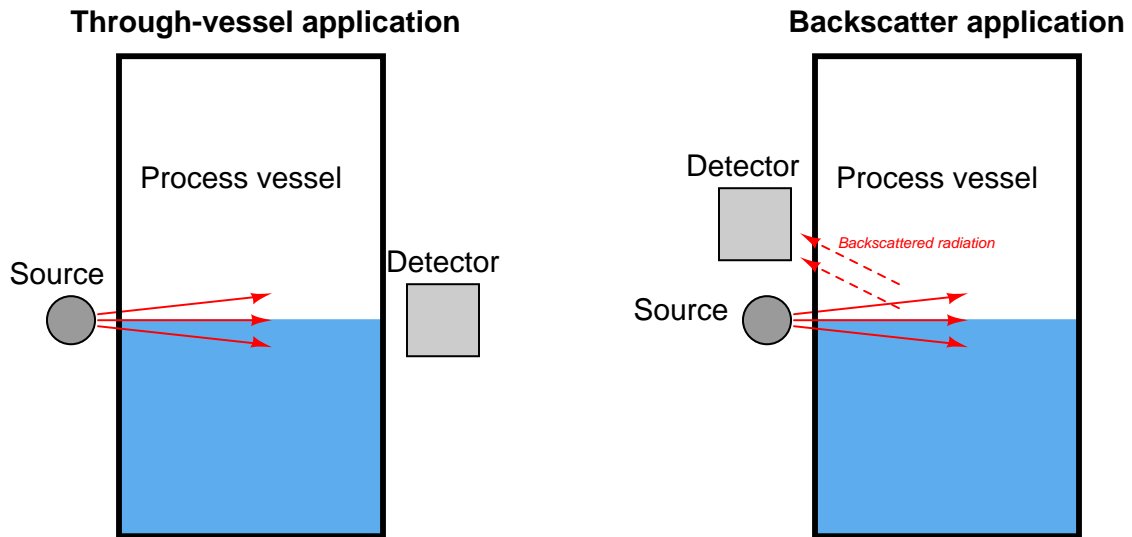
20.8 Radiation

Certain types of nuclear radiation easily penetrates the walls of industrial vessels, but is attenuated by traveling through the bulk of material stored within those vessels. By placing a radioactive source on one side of the vessel and measuring the radiation reaching the other side of the vessel, an approximate indication of level within that vessel may be obtained. Other types of radiation are *scattered* by process material in vessels, which means the level of process material may be sensed by sending radiation into the vessel through one wall and measuring *back-scattered* radiation returning through the same wall.

The four most common forms of nuclear radiation are *alpha particles* (α), *beta particles* (β), *gamma rays* (γ), and *neutrons* (n). Alpha particles are helium nuclei (2 protons bound together with 2 neutrons) ejected at high velocity from the nuclei of certain decaying atoms. They are easy to detect, but have very little penetrating power and so are not used for industrial level measurement. Beta particles are electrons⁴² ejected at high velocity from the nuclei of certain decaying atoms. Like alpha particles, though, they have little penetrating power and so are not used for industrial level measurement. Gamma rays are electromagnetic in nature (like X-rays and light waves) and have great penetrating power. Neutron radiation also penetrates metal very effectively, but is attenuated and scattered by any substance containing hydrogen (e.g. water, hydrocarbons, and many other industrial fluids), which makes it almost ideal for detecting the presence of a great many process fluids. These latter two forms of radiation (gamma rays and neutrons) are the most common in industrial measurement, with gamma rays used in through-vessel applications and neutrons typically used in backscatter applications.

⁴²Beta particles are *not* orbital electrons, but rather than product of elementary particle decay in an atom's nucleus. These electrons are spontaneously generated and subsequently ejected from the nucleus of the atom.

Through-vessel and *backscatter* nuclear level instrument applications appear contrasted in these two illustrations:




Nuclear radiation sources consist of radioactive samples contained in a shielded box. The sample itself is a small piece of radioactive substance encased in a double-wall stainless steel cladding, typically resembling a medicinal pill in size and shape. The specific type and quantity of radioactive source material depends on the nature and intensity of radiation required for the application. The basic rule here is that less is better: the smallest source capable of performing the measurement task is the best one for the application.

Common source types for gamma-ray applications are Cesium-137 and Cobalt-60. The numbers represent the *atomic mass* of each isotope: the sum total of protons and neutrons in the nucleus of each atom. These isotopes' nuclei are unstable, decaying over time to become different elements (Barium-137 and Nickel-60, respectively). Cobalt-60 has a relatively short half-life⁴³ of 5.3 years, whereas Cesium-137 has a much longer half-life of 30 years. This means radiation-based sensors using Cesium will be more stable over time (i.e. less calibration drift) than sensors using Cobalt. The trade-off is that Cobalt emits more powerful gamma rays than Cesium, which makes it better suited to applications where the radiation must penetrate thick process vessels or travel long distances (across wide process vessels).


⁴³The *half-life* of a radioactive substance is the amount of time it takes for one-half of the original quantity to experience radioactive decay. To illustrate, a 10-gram quantity consisting of 100% Cobalt-60 atoms will only contain 5 grams of Cobalt-60 after 5.3 years, and then only 2.5 grams of Cobalt-60 after another 5.3 years (10.6 years from the start), and so on. The actual mass of the sample does not change significantly over this time period because the Cobalt atoms have decayed into atoms of Nickel, which still have the same atomic mass value. However, the intensity of the gamma radiation emitted by the sample decreases over time, proportional to the percentage of Cobalt remaining therein.

One of the most effective methods of shielding against gamma ray radiation is with very dense substances such as lead or concrete. This is why the source boxes holding gamma-emitting radioactive pellets are lined with lead, so the radiation escapes only in the direction intended:



Radioactive sources naturally emit radiation, requiring no source of energy such as electricity to do their job. As such, they are “always-on” devices and may be locked out for testing and maintenance only by dropping a lead *shutter* over the “window” of the box. The lever actuating the shutter typically has provisions for lock-out/tag-out so a maintenance person may place a padlock on the lever and prevent anyone else from “turning on” the source during maintenance. For point-level (level switch) applications, the source shutter acts as a simple simulator for either a full vessel (in the case of a through-vessel installation) or an empty vessel (in the case of a backscatter installation). A full vessel may be simulated for neutron backscatter instruments by placing a sheet of plastic (or other hydrogen-rich substance) between the source box and the process vessel wall.

The detector for a radiation-based instrument is by far the most complex and expensive component of the system. Many different detector designs exist, the most common at the time of this writing being *ionization chambers* such as the Geiger-Muller (G-M) tube. In such devices, a thin metal wire centered in a metal cylinder sealed and filled with inert gas is energized with high voltage DC. Any ionizing radiation such as alpha, beta, or gamma radiation entering the tube causes gas molecules to ionize, allowing a pulse of electric current to travel between the wire and tube wall. A sensitive electronic circuit detects and counts these pulses, with a greater pulse rate corresponding to a greater intensity of detected radiation.



The following photograph shows an aluminum Geiger-Muller tube connected to a portable, battery-powered counter. This *Geiger counter* may be used as a piece of test equipment to measure radiation intensity while diagnosing problems in nuclear level measurement systems:



Geiger-style radiation detectors used as part of permanently-installed level measurement systems are housed in rugged housings, internally similar to the portable G-M tube shown in the photograph but designed for the rigors of continuous use in harsh industrial environments.

Neutron radiation is notoriously difficult to electronically detect, since neutrons are non-ionizing. Ionization tubes specifically made for neutron radiation detection are typically based on the Geiger-Muller design, but using tubes filled with special substances known to react with neutron radiation to produce (secondary) ionizing radiation. One example of such a detector is the so-called *fission chamber*, which is an ionization chamber lined with a fissile material such as uranium-235 (^{235}U). When a neutron enters the chamber and is captured by a fissile nucleus, that nucleus undergoes fission (splits into separate pieces) with a subsequent emission of gamma rays and charged particles, which are then detected by ionization in the chamber. Another variation on this theme is to fill an ionization tube with boron trifluoride gas. When a boron-10 (^{10}B) nucleus captures a neutron, it transmutes into lithium-7 (^7Li) and ejects an alpha particle and several beta particles, both of which cause detectable ionization in the chamber.


The accuracy of radiation-based level instruments varies with the stability of process fluid density, vessel wall coating, source decay rates, and detector drift. The multitude of error variables in radiation-based level measurement is one reason why they are more typically found as point-level (i.e. level switch) devices rather than continuous level (i.e. transmitter) measurement applications.

With their generally poor accuracy and the additional need for NRC (Nuclear Regulatory Commission) licensing to operate such instruments at an industrial facility, radiation instruments are typically used where no other instrument is practical. Examples include the level measurement of highly corrosive or toxic process fluids where penetrations into the vessel must be minimized and where piping requirements make weight-based measurement impractical (e.g. hydrocarbon/acid separators in alkylation processes in the oil refining industry), as well as processes where the internal conditions of the vessel are too physically violent for any instrument to survive (e.g. delayed coking vessels in the oil refining industry, where the coke is “drilled” out of the vessel by a high-pressure water jet).


20.9 Level sensor accessories

Disturbances in the liquid tend to complicate liquid level measurement. These disturbances may result from liquid introduced into a vessel above the liquid level (splashing into the liquid's surface), the rotation of agitator paddles, and/or turbulent flows from mixing pumps. Any source of turbulence for the liquid surface (or liquid-liquid interface) is especially problematic for echo-type level sensors, which *only* sense interfaces between vapors and liquids, or liquids and liquids.


If it is not possible to eliminate disturbances inside the process vessel, a relatively simple accessory one may add to the process vessel is a vertical length of pipe called a *stilling well*. To understand the principle of a stilling well, first consider the application of a hydraulic oil reservoir where we wish to continuously measure oil level. The oil flow in and out of this reservoir will cause problems for the displacer element:



A section of vertical pipe installed in the reservoir around the displacer will serve as a shield to all the turbulence in the rest of the reservoir. The displacer element will no longer be subject to a horizontal blast of oil entering the reservoir, nor any wave action to make it bob up and down. This section of pipe *quiets*, or *stills*, the oil surrounding the displacer, making it easier to measure oil level:




Stilling wells may be used in conjunction with many types of level instruments: floats, displacers, ultrasonic, radar, and laser to name a few. If the process application necessitates liquid-liquid interface measurement, however, the stilling well must be properly installed to ensure the interface level inside the well match the interface levels in the rest of the vessel. Consider this example of using a stilling well in conjunction with a tape-and-float system for interface measurement:




In the left-hand installation where the stilling well is completely submerged, the interface levels will always match. In the right-hand installation where the top of the stilling well extends above the total liquid level, however, the two levels may not always match.

This potential problem for the non-submerged stilling well is graphically illustrated here:



The problem here is analogous to what we see with sightglass-style level gauges: interfaces may be reliably indicated if and only if both ends of the sightglass are submerged (see section 20.1.2 beginning on page 1408 for an illustrated explanation of the problem).

If it is not possible or practical to ensure complete submersion of the stilling well, an alternative technique is to drill holes or cut slots in the well to allow interface levels to equalize inside and outside of the well tube:



Such equalization ports are commonly found as a standard design feature on coaxial probes for guided-wave radar level transmitters, where the outer tube of the coaxial transmission line acts as a sort of stilling well for the fluid. Coaxial probes are typically chosen for liquid-liquid interface radar measurement applications because they do the best job of preventing dispersion of the electromagnetic wave energy⁴⁴, but the “stilling well” property of a coaxial probe practically necessitates these equalization ports to ensure the interface level within the probe always matches the interface level in the rest of the vessel.

⁴⁴So much of the incident power is lost as the radar signal partially reflects off the gas-liquid interface, then the liquid-liquid interface, then *again* through the gas-liquid interface on its return trip to the instrument that every care must be taken to ensure optimum received signal strength. While twin-lead probes have been applied in liquid-liquid interface measurement service, the coaxial probe design is still the best for maintaining radar signal integrity.

20.10 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Definition of pressure:** $P = \frac{F}{A}$ (pressure is the amount of force applied over a specified area by a fluid).
- **Pascal's principle:** changes in fluid pressure are transmitted evenly throughout an enclosed fluid volume. Relevant to pressure measurement, as fluid pressure in all parts of an enclosed system will experience the same changes in pressure.
- **Hydrostatic pressure:** fluids having substantial weight generate pressure proportional to their density and to their vertical height ($P = \gamma h$ and $P = \rho gh$). Relevant to pressure offsets generated in vertical spans of impulse or capillary tubing, causing a pressure instrument to register more or less pressure than that at the process vessel connection.
- **Archimedes' principle:** the buoyant force experienced by an object submerged in liquid is equal to the weight of the fluid that object displaces, which is equal to the volume displaced multiplied by the weight density of the fluid ($F_{buoyant} = \gamma V$). Relevant to displacer-type instruments, which work by sensing the buoyant force exerted on an object as liquid rises around it.
- **Time, velocity, and distance:** $x = vt$, describing the relationship between velocity (v), time of travel (t), and distance traveled (x). Relevant to all types of "echo" level instruments, where travel time of a wave is used to measure distance.
- **Transmission lines:** short-duration (pulsed) electrical signals travel along a cable at nearly the speed of light, reflecting off of any discontinuity along the cable. Relevant to guided-wave radar level-sensing instruments, where the waveguide serves the same purpose as a transmission line, and the fluid-fluid boundary constitutes a discontinuity (sudden change in electrical permittivity) causing a reflected signal to arise.
- **Ideal Gas Law:** $PV = nRT$, used to calculate corrections to gas permittivity. Relevant to the "gas phase effect" of radar level instruments.
- **Capacitance:** $C = \frac{\epsilon A}{d}$, capacitance being proportional to the area of two overlapping conductors (A), the permittivity of the insulating (dielectric) substance between them (ϵ), and the distance (d) separating the conductors. Relevant to capacitive level sensing, where changes in liquid level alter the effective area, permittivity, and/or distance.

References

- “Autolevel” Application Note AN 01C22A01-01E, Yokogawa Electric Corporation, 2006.
- “Boiler Drum Level Transmitter Calibration”, application data sheet 00800-0100-3055, Rosemount, Inc., Chanhassen, MN, 2001.
- Brumbi, Detlef, *Fundamentals of Radar Technology for Level Gauging*, 4th Edition, Krohne Messtechnik GmbH & Co. KG, Duisburg, Germany, 2003.
- “Bubble Tube Installations For Liquid Level, Density, and Interface Measurements”, document MI 020-328, The Foxboro Company, Foxboro, MA, 1988.
- “DOE Fundamentals Handbook, Instrumentation and Control, Volume 2 of 2”, document DOE-HDBK-1013/2-92, U.S. Department of Energy, Washington, D.C., 1992.
- Fribance, Austin E., *Industrial Instrumentation Fundamentals*, McGraw-Hill Book Company, New York, NY, 1962.
- Kallen, Howard P., *Handbook of Instrumentation and Controls*, McGraw-Hill Book Company, Inc., New York, NY, 1961.
- “Level Measurement Technology: Radar”, document 00816-0100-3209, revision AA, Rosemount, Inc., Chanhassen, MN, 1999.
- Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.
- MacBeth, Michael, *IAEA CANDU Instrumentation & Control Course*, SNERDI, Shanghai, 1998.
- “Model 1151 Alphaline Pressure Transmitters”, product manual 00809-0100-4360, revision AA, Rosemount, Inc., Chanhassen, MN, 1997.
- “Replacing Displacers with Guided Wave Radar”, technical note 3300.2.02_CA, Rosemount, Inc., Chanhassen, MN, 2003.
- “The Art of Tank Gauging For Safety And Precision”, IN 4416.650, revision 6, Enraf B.V., The Netherlands.

Chapter 21

Continuous temperature measurement

Temperature is the measure of average molecular kinetic energy within a substance. The concept is easiest to understand for gases under low pressure, where gas molecules randomly shuffle about. The average kinetic (motional) energy of these gas molecules defines temperature for that quantity of gas. There is even a formula expressing the relationship between average kinetic energy ($\overline{E_k}$) and temperature (T) for a monatomic (single-atom molecule) gas:

$$\overline{E_k} = \frac{3kT}{2}$$

Where,

$\overline{E_k}$ = Average kinetic energy of the gas molecules (joules)

k = Boltzmann's constant (1.38×10^{-23} joules/Kelvin)

T = Absolute temperature of gas (Kelvin)

Thermal energy is a different concept: the quantity of *total kinetic energy* for this random molecular motion. If the average kinetic energy is defined as $\frac{3kT}{2}$, then the total kinetic energy for all the molecules in a monatomic gas must be this quantity times the total number of molecules (N) in the gas sample:

$$E_{\text{thermal}} = \frac{3NkT}{2}$$

This may be equivalently expressed in terms of the number of *moles* of gas rather than the number of molecules (a staggeringly large number for any realistic sample):

$$E_{\text{thermal}} = \frac{3nRT}{2}$$

Where,

E_{thermal} = Total thermal energy for a gas sample (joules)

n = Quantity of gas in the sample (moles)

R = Ideal gas constant (8.315 joules per mole-Kelvin)

T = Absolute temperature of gas (Kelvin)

Heat is defined as the exchange of thermal energy from one sample to another, by way of conduction (direct contact), convection (transfer via a moving fluid), or radiation (emitted energy); although you will often find the terms *thermal energy* and *heat* used interchangeably.

Illustrating by way of example, a single molecule of gas moving at a constant velocity will have a definite temperature. Two or three molecules moving at the same speed will have the same *temperature*, but together represent a greater *thermal energy* than any one of them considered alone. Heat is either the reduction or increase of thermal energy by transfer of energy. If these gas molecules happen to collide with slower-moving gas molecules, the faster molecules will lose velocity while the slower molecules will gain velocity. Thus, the higher-temperature molecules cool down while the lower-temperature molecules warm up: a transfer of *heat*.

Temperature is a more easily detected quantity than either thermal energy or heat. In fact, when we need to measure either thermal energy or heat, we do so by measuring temperature and then inferring the desired variable based on the laws of thermodynamics.

There are many different ways to measure temperature, from a simple glass-bulb mercury thermometer to sophisticated infrared optical sensor systems. Like all other areas of measurement, there is no single technology that is best for all applications. Each temperature-measurement technique has its own strengths and weaknesses. One responsibility of the instrument technician is to know these pros and cons so as to choose the best technology for the application, and this knowledge is best obtained through understanding the operational principles of each technology.

21.1 Bi-metal temperature sensors

Solids tend to expand when heated. The amount that a solid sample will expand with increased temperature depends on the size of the sample, the material it is made of, and the amount of temperature rise. The following formula relates linear expansion to temperature change:

$$l = l_0(1 + \alpha\Delta T)$$

Where,

l = Length of material after heating

l_0 = Original length of material

α = Coefficient of linear expansion


ΔT = Change in temperature

Here are some typical values of α for common metals:


- Aluminum = 25×10^{-6} per degree C
- Copper = 16.6×10^{-6} per degree C
- Iron = 12×10^{-6} per degree C
- Tin = 20×10^{-6} per degree C
- Titanium = 8.5×10^{-6} per degree C

As you can see, the values for α are quite small. This means the amount of expansion (or contraction) for modest temperature changes are almost too small to see unless the sample size (l_0) is huge. We can readily see the effects of thermal expansion in structures such as bridges, where expansion joints must be incorporated into the design to prevent serious problems due to changes in ambient temperature. However, for a sample the size of your hand the change in length from a cold day to a warm day will be microscopic.

One way to amplify the motion resulting from thermal expansion is to bond two strips of dissimilar metals together, such as copper and iron. If we were to take two equally-sized strips of copper and iron, lay them side-by-side, and then heat both of them to a higher temperature, we would see the copper strip lengthen slightly more than the iron strip:



If we bond these two strips of metal together, this differential growth will result in a bending motion greatly exceeding the linear expansion. This device is called a *bi-metal strip*:



This bending motion is significant enough to drive a pointer mechanism, activate an electromechanical switch, or perform any number of other mechanical tasks, making this a very simple and useful *primary sensing element* for temperature. Older home thermostats often used this principle to both indicate room temperature and to actuate electrical switches for controlling room temperature. Electric hot water heater units still use this type of device (usually in the form of a convex bi-metal disk) to sense over-temperature conditions and automatically shut off power to the heater if the water temperature exceeds a pre-set limit.


If a bi-metallic strip is twisted over a long length, it will tend to un-twist as it heats up. This twisting motion may be used to directly drive the needle of a temperature gauge. This is the operating principle of the temperature gauge shown in the following photograph:




21.2 Filled-bulb temperature sensors

Filled-bulb systems exploit the principle of fluid expansion to measure temperature. If a fluid is enclosed in a sealed system and then heated, the molecules in that fluid will exert a greater pressure on the walls of the enclosing vessel. By measuring this pressure, and/or by allowing the fluid to expand under constant pressure, we may infer the temperature of the fluid.

Class I and Class V systems use a liquid fill fluid (class V is mercury). Here, the volumetric expansion of the liquid drives an indicating mechanism to show temperature:




Class III systems use a gas fill fluid instead of liquid. Here, the change in pressure with temperature (as described by the Ideal Gas Law) allows us to sense the bulb's temperature:



In these systems, it is quite critical that the tube connecting the sensing bulb to the indicating element be of minimal volume, so the fluid expansion is primarily due to changes in temperature at the bulb rather than changes in temperature along the length of the tube. It is also important to realize that the fluid volume contained by the bellows (or bourdon tube or diaphragm . . .) is also subject to expansion and contraction due to temperature changes at the indicator. This means the temperature indication varies somewhat as the indicator temperature changes, which is not desirable, since we intend the device to measure temperature (exclusively) at the bulb. Various methods of compensation exist for this effect (for example, a bi-metal spring inside the indicator mechanism to automatically offset the indication as ambient temperature changes), but it may be permanently offset through a simple “zero” adjustment provided that the ambient temperature at the indicator does not change much.

A fundamentally different class of filled-bulb system is the Class II, which uses a volatile liquid/vapor combination to generate a temperature-dependent fluid expansion:




Given that the liquid and vapor are in direct contact with each other, the pressure in the system will be precisely equal to the *saturated vapor pressure* at the vapor/liquid interface. This makes the Class II system sensitive to temperature only at the bulb and nowhere else along the system's volume. Because of this phenomenon, a Class II filled-bulb system requires no compensation for temperature changes at the indicator.

Class II systems do have one notable idiosyncrasy, though: they have a tendency to switch from Class IIA to Class IIB when the temperature of the sensing bulb crosses the ambient temperature at the indicator. Simply put, the liquid tends to seek the colder portion of a Class II system while the vapor tends to seek the warmer portion. This causes problems when the indicator and sensing bulb exchange identities as warmer/colder. The rush of liquid up (or down) the capillary tubing as the system tries to reach a new equilibrium causes intermittent measurement errors. Class II filled-bulb systems designed to operate in either IIA or IIB mode are classified as *IIC*.

One calibration problem common to all systems with liquid-filled capillary tubes is an offset in temperature measurement due to hydrostatic pressure (or suction) resulting from a difference in height between the measurement bulb and the indicator. This represents a "zero" shift in calibration, which may be permanently offset by a "zero" adjustment at the time of installation. Class III (gas-filled) and Class IIB (vapor-filled) systems, of course, suffer no such problem because there is no liquid in the capillary tube to generate a pressure due to height.

A photograph of a pneumatic temperature transmitter using a filled-bulb as the sensing element appears here:



This transmitter happens to be a Moore Products “Nullmatic” model. The capillary tube connecting the fluid-filled bulb to the transmitter mechanism is protected by a spiral-metal jacket. The bulb itself is located at the very end of the stainless steel “wand” which inserts into the process fluid to be measured:




Instead of directly actuating a pointer mechanism, the fluid pressure in this instrument actuates a self-balancing pneumatic mechanism to produce a 3 to 15 PSI air pressure signal representing process temperature.

Filled-bulb temperature sensors are seldom used in industrial applications anymore, chiefly due to the superiority of electrical sensors. The only significant advantage filled-bulb sensors hold over electrical sensors is not needing electricity¹ to function, but this is usually not a serious consideration within a modern industrial facility.

¹Even this advantage is not always true. It is possible to build self-powered thermocouple temperature indicators, where an analog meter movement is driven by the electrical energy a thermocouple sensing junction outputs. Here, no external electrical power source is required! However, the accuracy of self-powered thermocouple systems is poor, as is the ability to measure small temperature ranges.

21.3 Thermistors and Resistance Temperature Detectors (RTDs)

One of the simplest classes of temperature sensor is one where temperature effects a change in electrical resistance. With this type of primary sensing element, a simple ohmmeter is able to function as a thermometer, interpreting the resistance as a temperature measurement:



Thermistors are devices made of metal oxide which either increase in resistance with increasing temperature (a *positive temperature coefficient*) or decrease in resistance with increasing temperature (a *negative temperature coefficient*). *RTDs* are devices made of pure metal wire (usually platinum or copper) which always increase in resistance with increasing temperature. The major difference between thermistors and RTDs is linearity: thermistors are highly sensitive and nonlinear, whereas RTDs are relatively insensitive but very linear. For this reason, thermistors are typically used where high accuracy is unimportant. Many consumer-grade devices use thermistors for temperature sensors.

21.3.1 Temperature coefficient of resistance (α)

A *Resistive Temperature Detector* (RTD) is a special temperature-sensing element made of fine metal wire, the electrical resistance of which changes with temperature as approximated by the following formula:

$$R_T = R_{ref}[1 + \alpha(T - T_{ref})]$$

Where,

R_T = Resistance of RTD at given temperature T (ohms)

R_{ref} = Resistance of RTD at the reference temperature T_{ref} (ohms)

α = Temperature coefficient of resistance (ohms per ohm/degree)

The following example shows how to use this formula to calculate the resistance of a “100 ohm” platinum RTD with a temperature coefficient value of 0.00392 at a temperature of 35 degrees Celsius:

$$R_T = 100 \Omega[1 + (0.00392)(35^\circ \text{C} - 0^\circ \text{C})]$$

$$R_T = 100 \Omega[1 + 0.1372]$$

$$R_T = 100 \Omega[1.1372]$$

$$R_T = 113.72 \Omega$$

Due to nonlinearities in the RTD’s behavior, this linear RTD formula is only an approximation. A better approximation is the *Callendar-van Dusen formula*, which introduces second, third, and fourth-degree terms for a better fit: $R_T = R_{ref}(1 + AT + BT^2 - 100CT^3 + CT^4)$ for temperatures ranging $-200^\circ \text{C} < T < 0^\circ \text{C}$ and $R_T = R_{ref}(1 + AT + BT^2)$ for temperatures ranging $0^\circ \text{C} < T < 661^\circ \text{C}$, both assuming $T_{ref} = 0^\circ \text{C}$. The A , B , and C coefficients vary with the specific type of RTD, equivalent in role to α in the linear RTD formula.

Water’s melting/freezing point is the standard reference temperature for most RTDs. Here are some typical values of α for common metals:

- Nickel = $0.00672 \Omega/\Omega^\circ \text{C}$
- Tungsten = $0.0045 \Omega/\Omega^\circ \text{C}$
- Silver = $0.0041 \Omega/\Omega^\circ \text{C}$
- Gold = $0.0040 \Omega/\Omega^\circ \text{C}$
- Platinum = $0.00392 \Omega/\Omega^\circ \text{C}$
- Copper = $0.0038 \Omega/\Omega^\circ \text{C}$


As mentioned previously, platinum is a common wire material for industrial RTD construction. The alpha (α) value for platinum varies according to the alloying of the metal. For “reference grade” platinum wire, the most common alpha value is 0.003902. Industrial-grade platinum alloy RTD wire is commonly available in two different coefficient values: 0.00385 (the “European” alpha value) and 0.00392 (the “American” alpha value), of which the “European” value of 0.00385 is most commonly used (even in the United States!).

An alternative to mathematically predicting the resistance of an RTD is to simply look up its predicted resistance versus temperature in a *table* of values published for that RTD type. Tables capture the nuances of an RTD’s non-linearity without adding any mathematical complexity: simply look up the resistance corresponding to a given temperature, or vice-versa. If a value falls between two nearest entries in the table, you may *interpolate* to find the desired value, regarding the two nearest table entries as end-points defining a line segment, calculating the point you desire along that line.

100 Ω is a very common reference resistance (R_{ref} at 0 degrees Celsius) for industrial RTDs. 1000 Ω is another common reference resistance, and some industrial RTDs have reference resistances as low as 10 Ω . Compared to thermistors with their tens or even hundreds of thousands of ohms’ resistance, an RTD’s resistance is comparatively small. This can cause problems with measurement, since the wires connecting an RTD to its ohmmeter possess their own resistance, which will be a more substantial percentage of the total circuit resistance than for a thermistor.

21.3.2 Two-wire RTD circuits

The following schematic diagrams show the relative effects of 2 ohms total wire resistance on a thermistor circuit and on an RTD circuit:



As you can see, wire resistance *adds* to the sensing element's resistance to create a larger total circuit resistance which will be interpreted by the receiving instrument (ohmmeter) as a *falsely high* temperature reading, assuming a positive temperature coefficient of resistance for the sensing element.


Clearly, wire resistance is more problematic for low-resistance RTDs than for high-resistance thermistors. In the RTD circuit, wire resistance constitutes 1.96% of the total circuit resistance. In the thermistor circuit, the same 2 ohms of wire resistance constitutes only 0.004% of the total circuit resistance. The thermistor's huge reference resistance value "swamps"² the wire resistance to the point that the latter becomes insignificant by comparison.

In HVAC (Heating, Ventilation, and Air Conditioning) systems, where the temperature measurement range is relatively narrow, the nonlinearity of thermistors is not a serious concern and their relative immunity to wire resistance error is a definite advantage over RTDs. In industrial temperature measurement applications where the temperature ranges are usually much wider, the nonlinearity of thermistors becomes a significant problem, so we must find a way to use low-resistance RTDs and deal with the (lesser) problem of wire resistance.

²"Swamping" is the term given to the overshadowing of one effect by another. Here, the normal resistance of the thermistor greatly overshadows ("swamps") any wire resistance in the circuit, such that wire resistance becomes negligible.

21.3.3 Four-wire RTD circuits

A very old electrical measurement technique known as the *Kelvin* or *four-wire* method is a practical solution to the problem of wire resistance. Commonly employed to make precise resistance measurements for scientific experiments in laboratory conditions, the four-wire technique uses four wires to connect the resistance under test (in this case, the RTD) to the measuring instrument, which consists of a voltmeter and a precision current source. Two wires carry “excitation” current to the RTD from the current source while the other two wires merely “sense” voltage drop across the RTD resistor element and carry that voltage signal to the voltmeter. RTD resistance is calculated using Ohm’s Law: taking the measured voltage displayed by the voltmeter and dividing that figure by the regulated current value of the current source. A simple 4-wire RTD circuit is shown here for illustration:



Wire resistances are completely inconsequential in this circuit. The two “excitation” wires carrying current to the RTD will drop some voltage along their length, but this voltage drop is only “seen” by the current source and not the voltmeter. The two “sense” wires connecting the voltmeter to the RTD also possess resistance, but they drop negligible voltage because the voltmeter draws so little current through them³. Thus, the resistances of the current-carrying wires are of no effect because the voltmeter never senses their voltage drops, and the resistances of the voltmeter’s sensing wires are of no effect because they carry practically zero current.


Note how wire colors (*white* and *red*) indicate which wires are common pairs at the RTD. The RTD is polarity-insensitive because it is nothing more than a resistor, which is why it doesn’t matter which color is positive and which color is negative.

The only disadvantage of the four-wire method is the sheer number of wires necessary. Four wires *per RTD* can add up to a sizeable wire count when many different RTDs are installed in a process area.

³Remember that an ideal voltmeters has *infinite* input impedance, and modern semiconductor-amplified voltmeters have impedances of several mega-ohms or more.

21.3.4 Three-wire RTD circuits

A compromise between two-wire and four-wire RTD connections is the *three-wire* connection, which looks like this:



In a three-wire RTD circuit, voltmeter “A” measures the voltage dropped across the RTD plus the voltage dropped across the bottom current-carrying wire. Voltmeter “B” measures just the voltage dropped across the top current-carrying wire. Assuming both current-carrying wires will have (very nearly) the same resistance, subtracting the indication of voltmeter “B” from the indication given by voltmeter “A” yields the voltage dropped across the RTD:

$$V_{RTD} = V_{\text{meter(A)}} - V_{\text{meter(B)}}$$

Once again, RTD resistance is calculated from the RTD voltage and the known current source value using Ohm’s Law, just as it is in a 4-wire circuit.


If the resistances of the two current-carrying wires are precisely identical (and this includes the electrical resistance of any connections within those current-carrying paths, such as terminal blocks), the calculated RTD voltage will be the same as the true RTD voltage, and no wire-resistance error will appear. If, however, one of those current-carrying wires happens to exhibit more resistance than the other, the calculated RTD voltage will not be the same as the actual RTD voltage, and a measurement error will result.

Thus, we see that the three-wire RTD circuit saves us wire cost over a four-wire circuit, but at the “expense” of a potential measurement error. The beauty of the four-wire design was that wire resistances were completely irrelevant: a true determination of RTD voltage (and therefore RTD resistance) could be made regardless of how much resistance each wire had, or even if the wire resistances were different from each other. The error-canceling property of the three-wire circuit, by contrast, hinges on the assumption that the two current-carrying wires have exactly the same resistance, which may or may not actually be true.

It should be understood that real three-wire RTD instruments do not employ direct-indicating voltmeters as shown in these simplified examples. Actual RTD instruments use either analog or digital “conditioning” circuits to measure the voltage drops and perform the necessary calculations

to compensate for wire resistance. The voltmeters shown in the four-wire and three-wire diagrams serve only to illustrate the basic concepts, not to showcase practical instrument designs.

A practical electronic circuit for a 3-wire RTD sensor is shown here (differential voltages shown in blue, ground-referenced voltages shown in red):



Note that the voltage appearing at point B with reference to ground is the RTD's voltage plus the voltage dropped by the lower current-carrying wire: $V_{RTD} + V_{error}$. It is this "error" voltage we must eliminate in order to achieve an accurate measurement of RTD voltage drop, essential for accurately inferring RTD temperature. The voltage appearing at point A is greater by the upper wire's voltage drop ($V_{error} + V_{RTD} + V_{error}$) because that point spans one more wire resistance in the circuit than point B⁴.

Like all negative-feedback operational amplifier circuits, the amplifier does its best to maintain the two input terminals at (nearly) the same voltage. Thus, the voltage at point B is duplicated at the inverting input terminal by the amplifier's feedback action. From this we may see that the voltage drop across the left-hand 100 kΩ resistor is simply V_{error} : the potential difference between point A and point B. The feedback current driving through this resistor goes through the other 100 kΩ feedback resistor equally, causing the same voltage drop to appear there (V_{error}). We may see that the polarity of this second resistor's voltage drop ends up subtracting that quantity from the voltage appearing at the inverting input terminal. The inverting terminal voltage ($V_{RTD} + V_{error}$) minus the right-hand 100 kΩ resistor's voltage drop (V_{error}) is simply V_{RTD} , and so the voltmeter registers the true RTD voltage drop without any wire resistance error.

Like the dual-voltmeter circuit shown previously, this amplified 3-wire RTD sensing circuit "assumes" the two current-carrying wires will have the exact same resistance and therefore drop

⁴Note that the middle wire resistance is of no effect because it does not carry the RTD's current. The amount of current entering or exiting an operational amplifier is assumed to be zero for all practical purposes.

the same amount of voltage. If this is not the case, and one of these wires drops more voltage than the other, the circuit will fail to yield the exact RTD voltage (V_{RTD}) at the amplifier output. This is the fundamental limitation of any 3-wire RTD circuit: the cancellation of wire resistance is only as good as the wires' resistances are precisely equal to each other.

A photograph of a modern temperature transmitter capable of receiving input from 2-wire, 3-wire, or 4-wire RTDs (as well as thermocouples, another type of temperature sensor entirely) shows the connection points and the labeling describing how the sensor is to be connected to the appropriate terminals:

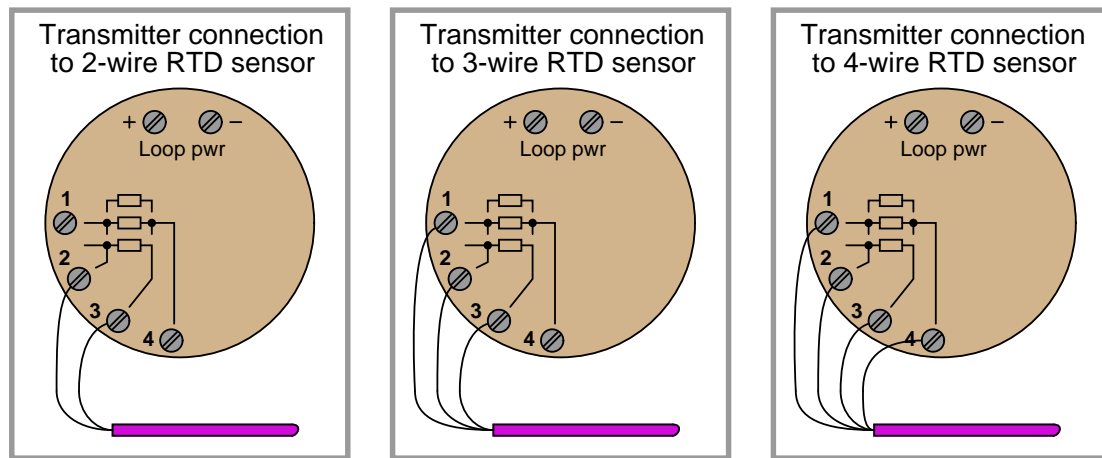


The rectangle symbol shown on the label represents the resistive element of the RTD. The symbol with the “+” and “-” marks represents a thermocouple junction, and may be ignored for the purposes of this discussion. As shown by the diagram, a two-wire RTD would connect between terminals 2 and 3. Likewise, a three-wire RTD would connect to terminals 1, 2, and 3 (with terminals 1 and 2 being the points of connection for the two common wires of the RTD). Finally, a four-wire RTD would connect to terminals 1, 2, 3, and 4 (terminals 1 and 2 being common, and terminals 3 and 4 being common, at the RTD).

Once the RTD has been connected to the appropriate terminals of the temperature transmitter, the transmitter needs to be electronically configured for that type of RTD. In the case of this particular temperature transmitter, the configuration is performed using a “smart” communicator device using the HART digital protocol to access the transmitter’s microprocessor-based settings. Here, the technician would configure the transmitter for 2-wire, 3-wire, or 4-wire RTD connection.

21.3.5 Proper RTD sensor connections

Proper connections for all three types of RTD sensor (2-wire, 3-wire, and 4-wire) to a user-configurable transmitter are shown in the following illustrations:







It is critically important to note that the common connections shown by the symbols for 3- and 4-wire RTD sensors represent junction points *at the sensor*; not terminals jumpered by the technician at the time of installation, and not internal jumpers inside the transmitter. The whole purpose of having 3-wire and 4-wire RTD circuits is to eliminate errors due to voltage drop along the current-carrying wires, and this can only be realized if the “sensing” wire(s) extend out to the RTD itself and connect there. If the transmitter’s sensing terminal(s) are only jumpered to a current-carrying terminal, the transmitter will sense voltage dropped by the RTD *plus* voltage dropped by the current-carrying wire(s), leading to falsely high temperature indications.

Misconceptions surrounding proper RTD connections unfortunately abound both in students and in working industry professionals. With any luck, the following presentation will help you avoid such mistakes, and more importantly help you *understand why* the correct connections are best.


Always bear in mind the purpose of a 3-wire or a 4-wire RTD connection: *to avoid inaccuracies caused by voltage drops along the current-carrying wires*. The only way to do this is to ensure the sensing (non-current-carrying) wire(s) extend from the transmitter terminal(s) all the way to the sensor itself. This way, the transmitter is able to “look past” the voltage drops of the current-carrying wires to “see” the voltage dropped only by the RTD itself.

The following illustrations show both correct and incorrect ways to connect a 2-wire RTD to a 3- or 4-wire transmitter:



Jumpers placed at the transmitter terminals defeat the purpose of the transmitter’s 3-wire or 4-wire capabilities, downgrading its performance to that of a 2-wire system.


A similar problem occurs when someone tries to connect a 3-wire RTD to a 3-wire transmitter using a conveniently available 4-wire cable:



3-wire RTD measurement is based on the assumption that both current-carrying wires have exactly the same electrical resistance. By paralleling two of the four wires in the 4-wire cable, you will create unequal resistances in the current path, thus leading to measurement errors at the transmitter⁵.

⁵These errors will result only if the paralleled wires carry current. If the two wires you paralleled happen to join the transmitter's sensing terminal to the RTD (the one carrying no current), no errors will result. However, many RTD transmitters do not document which of the terminals *sense* (carry no current) versus which of them *excite* (carry current to the RTD), and so there is a probability of getting it wrong if you simply guess. Given that there is no real benefit to having paralleled wires connecting the transmitter's sensing terminal to the RTD, my advice is to either use all four wires and configure the transmitter for 4-wire mode, or don't use the fourth wire at all.

Better solutions for the 3-wire RTD and 4-wire cable scenario include configuring the transmitter for 4-wire RTD input and actually using all four terminals (shown on left), or keeping the transmitter configured for 3-wire RTD input and not using the fourth wire in the cable at all (shown on right):



21.3.6 Self-heating error

One problem inherent to both thermistors and RTDs is *self-heating*. In order to measure the resistance of either device, we must pass an electric current through it. Unfortunately, this results in the generation of heat at the resistance according to Joule's Law:

$$P = I^2 R$$

This dissipated power causes the thermistor or RTD to increase in temperature beyond its surrounding environment, introducing a positive measurement error. The effect may be minimized by limiting excitation current to a bare minimum, but this results in less voltage dropped across the device. The smaller the developed voltage, the more sensitive the voltage-measuring instrument must be to accurately sense the condition of the resistive element. Furthermore, a decreased signal voltage means we will have a decreased signal-to-noise ratio, for any given amount of noise induced in the circuit from external sources.

One clever way to circumvent the self-heating problem without diminishing excitation current to the point of uselessness is to *pulse* current through the resistive sensor and digitally sample the voltage only during those brief time periods while the thermistor or RTD is powered. This technique works well when we are able to tolerate slow sample rates from our temperature instrument, which is often the case because most temperature measurement applications are slow-changing by nature. The pulsed-current technique enjoys the further advantage of reducing power consumption for the instrument, an important factor in battery-powered temperature measurement applications.

21.4 Thermocouples


RTDs are completely passive sensing elements, requiring the application of an externally-sourced electric current in order to function as temperature sensors. Thermocouples, however, generate their own electric potential. In some ways, this makes thermocouple systems simpler because the device receiving the thermocouple's signal does not have to supply electric power to the thermocouple. It also makes thermocouple systems potentially safer than RTDs in applications where explosive compounds may exist in the atmosphere, because the power levels generated by a thermocouple tend to be less than the power levels dissipated by an RTD. The self-powering nature of thermocouples also means they do not suffer from the same "self-heating" effect as RTDs.

In other ways, however, thermocouple circuits are more complex and troublesome than RTD circuits because the generation of voltage actually occurs in *two* different locations within the circuit, not simply at the sensing point. This means the receiving circuit must "compensate" for temperature in another location in order to accurately measure temperature in the desired location.


Though typically not as accurate as RTDs, thermocouples are more rugged, have greater temperature measurement spans, and are easier to manufacture in different physical forms.

21.4.1 Dissimilar metal junctions

When two dissimilar metal wires are joined together at one end, a voltage is produced at the other end that is approximately proportional to temperature. That is to say, the junction of two different metals behaves like a temperature-sensitive battery. This form of electrical temperature sensor is called a *thermocouple*:



This phenomenon provides us with a simple way to electrically infer temperature: simply measure the voltage produced by the junction, and you can tell the temperature of that junction. And it would be that simple, if it were not for an unavoidable consequence of electric circuits: when we connect any kind of electrical instrument to the thermocouple wires, we inevitably produce another junction of dissimilar metals. The following schematic shows this fact, where the iron-copper junction J_1 is necessarily complemented by a second iron-copper junction J_2 of opposing polarity:



Junction J_1 is a junction of iron and copper – two dissimilar metals – which will generate a voltage related to temperature. Note that junction J_2 , which is necessary for the simple fact that we must somehow connect our copper-wired voltmeter to the iron wire, is also a dissimilar-metal junction which will also generate a voltage related to temperature. Further note how the polarity of junction J_2 stands opposed to the polarity of junction J_1 (iron = positive ; copper = negative). A third junction (J_3) also exists between wires, but it is of no consequence because it is a junction of two identical metals which does not generate a temperature-dependent voltage at all.

The presence of this second voltage-generating junction (J_2) helps explain why the voltmeter registers 0 volts when the entire system is at room temperature: any voltage generated by the iron-copper junctions will be equal in magnitude and opposite in polarity, resulting in a net (series-total) voltage of zero. Only when the two junctions J_1 and J_2 are at *different* temperatures will the voltmeter register any voltage at all.

We may express this relationship mathematically as follows:

$$V_{meter} = V_{J1} - V_{J2}$$

With the measurement (J_1) and reference (J_2) junction voltages opposed to each other, the voltmeter only “sees” the difference between these two voltages.

Thus, thermocouple systems are fundamentally *differential* temperature sensors. That is, they provide an electrical output proportional to the difference in temperature between two different points. For this reason, the wire junction we use to measure the temperature of interest is called the *measurement junction* while the other junction (which we cannot eliminate from the circuit) is called the *reference junction* (or the *cold junction*, because it is typically at a cooler temperature than the process measurement junction).

Much of the complexity of thermocouples is related to the reference junction voltage and how we must deal with that (unwanted) potential when using a thermocouple as a measuring device. For most practical applications, we just want to measure the temperature at one location, not the difference in temperature between two locations which is what a thermocouple naturally does. A number of different techniques exist to deal with this problem – forcing a differential temperature sensor to act like a single-point temperature sensor – and we will explore the most common techniques in this section.

Students and working professionals alike often find this concept of a reference junction and its effects endlessly confusing. My advice to the confused is to return to the simple iron-copper wire circuit shown previously as a “starting point,” and then deduce its behavior from first principles⁶. We know that a dissimilar-metal junction creates a voltage with temperature. We also know that in order to make a complete circuit with iron and copper wire *there must be a second iron-copper junction somewhere else in that same circuit, the polarity of which is necessarily opposed to the first*. If we call the first iron-copper junction J_1 and the second junction J_2 , we absolutely must conclude that the net voltage registered by the voltmeter in this circuit will be $V_{J1} - V_{J2}$.

All thermocouple circuits – no matter how simple or complex – exhibit this fundamental property. Mentally constructing a simple circuit of two dissimilar-metal wires and then performing “thought experiments” to see how that circuit will behave with those junctions at the same temperature and also at different temperatures is the best way I can suggest for any person to comprehend thermocouples. Students especially tend to cope with complexity through memorization: committing to memory catch-phrases and formulae such as $V_{meter} = V_{J1} - V_{J2}$. This is a poor coping mechanism, as it grants the illusion of understanding with none of the substance. The real secret is to know *why* a thermocouple circuit acts as it does, and that only comes through practiced reasoning. Throughout the rest of this section, as we explore reference junction compensation, how to interpret voltage measurements in thermocouple circuits, and how to simulate thermocouples at temperature, we will keep returning to this simple iron-copper wire circuit to refresh our understanding of how and why thermocouple circuits behave. If you understand this one fundamental concept, the rest will make sense to you. If you continually find yourself confused by thermocouple circuits, it means you do not yet fully understand this basic circuit, and you need to return to it and think it through until you do.

⁶By “first principles,” I mean the basic laws of electric circuits. In this case, the most important law to apply is Kirchhoff’s Voltage Law: the algebraic sum of voltages in any loop must be equal to zero.

21.4.2 Thermocouple types

Thermocouples exist in many different types, each with its own color codes for the dissimilar-metal wires. Here is a table showing the more common thermocouple types and their standardized colors⁷, along with some distinguishing characteristics of the metal types to aid in polarity identification when the wire colors are not clearly visible:

Type	Positive wire characteristic	Negative wire characteristic	Plug	Temp. range
T	Copper (blue) <i>yellow colored</i>	Constantan (red) <i>silver colored</i>	Blue	−300 to 700 °F
J	Iron (white) <i>magnetic, rusty?</i>	Constantan (red) <i>non-magnetic</i>	Black	32 to 1400 °F
E	Chromel (violet) <i>shiny finish</i>	Constantan (red) <i>dull finish</i>	Violet	32 to 1600 °F
K	Chromel (yellow) <i>non-magnetic</i>	Alumel (red) <i>magnetic</i>	Yellow	32 to 2300 °F
N	Nicrosil (orange)	Nisil (red)	Orange	32 to 2300 °F
S	Pt90% - Rh10% (black)	Platinum (red)	Green	32 to 2700 °F
B	Pt70% - Rh30% (grey)	Pt94% - Rh6% (red)	Grey	32 to 3380 °F

Types S and B use platinum or platinum-rhodium alloy wire, with different alloying distinguishing the positive from the negative wires. Sometimes type B is colored green and red rather than grey and red.

Note how the negative (−) wire of every thermocouple type is color-coded *red*. While this may seem backward to those familiar with modern electronics (where red and black usually represent the positive and negative poles of a DC power supply, respectively), bear in mind that thermocouple color codes actually pre-date electronic power supply wire coloring!

Aside from having different usable temperature ranges, these thermocouple types also differ in terms of the atmospheres they may withstand at elevated temperatures. Type J thermocouples, for instance, by virtue of the fact that one of the wire types is *iron*, will rapidly corrode in any oxidizing⁸ atmosphere. Type K thermocouples are attacked by reducing⁹ atmospheres as well as sulfur and cyanide. Type T thermocouples are limited in upper temperature by the oxidation of copper (a very reactive metal when hot), but stand up to both oxidizing and reducing atmospheres quite well at lower temperatures, even when wet.

One final note on the thermocouple types shown in this table is that the temperature ranges given are approximate, and vary with the intended measurement accuracy. One may have to stay

⁷The colors in this table apply only to the United States and Canada. A stunning diversity of colors has been “standardized” for each thermocouple type per nationality. The British and Czechs use their own color code, as do the Dutch and Germans. France has its own unique color code as well. Just for fun, an “international” color code also exists which doesn’t match any of the others. There are other deviations as well: the wire colors for type R and S thermocouples, for example, are standardized for *extension-grade wire* but not for *thermocouple-grade wire*.

⁸By “oxidizing,” what is meant is any atmosphere containing sufficient oxygen molecules or molecules of a similar element such as chlorine or fluorine.


⁹“Reducing” refers to atmospheres rich in elements that readily oxidize. Practically any fuel gas (hydrogen, methane, etc.) will create a reducing atmosphere in sufficient concentration.

within a more limited range of temperature than what is shown in this table if a certain minimum level of accuracy is desired from the thermocouple. Consult manufacturers' data for details!

21.4.3 Connector and tip styles

In its simplest form, a thermocouple is nothing more than a pair of dissimilar-metal wires joined together. However, in industrial practice, we often must package thermocouples in a more rugged form than a bare metal junction. For instance, most industrial thermocouples are manufactured in such a way that the dissimilar-metal wires are protected from physical damage by a stainless steel or ceramic *sheath*, and they are often equipped with molded-plastic plugs for quick connection to and disconnection from a thermocouple-based instrument.

A photograph of a type K industrial thermocouple (approximately 20 inches in length) reveals this “sheathed” and “connectorized” construction:



The stainless steel sheath of this particular thermocouple shows signs of discoloration from previous service in a hot process. Note the different diameters of the plug terminals. This “polarized” design makes it difficult¹⁰ to insert backward into a matching socket.

¹⁰It should be noted that no amount of engineering or design is able to *completely* prevent people from doing the wrong thing. I have seen this style of thermocouple plug forcibly mated the wrong way to a socket. The amount of insertion force necessary to make the plug fit backward into the socket was quite extraordinary, yet this apparently was not enough of a clue for this wayward individual to give them pause.

A miniature version of this same plug (designed to attach to thermocouple wire by screw terminals, rather than be molded onto the end of a sheathed assembly) is shown here, situated next to a ballpoint pen for size comparison:



Industrial-grade thermocouples are available with this miniature style of molded plug end as an alternative to the larger (standard) plug. Miniature plug-ends are often the preferred choice for laboratory applications, while standard-sized plugs are often the preferred choice for field applications.

Some industrial thermocouples have no molded plug at all, but terminate simply in a pair of open wire ends. The following photograph shows a type J thermocouple of this construction:



If the electronic measuring instrument (e.g. temperature transmitter) is located near enough for the thermocouple's wires to reach the connection terminals, no plug or socket is needed at all in the circuit. If, however, the distance between the thermocouple and measuring instrument is too far to span with the thermocouple's own wires, a common termination technique is to attach a special terminal block and connection "head" to the top of the thermocouple allowing a pair of thermocouple extension wires to join and carry the millivoltage signal to the measuring instrument.


This next photograph shows a close-up view of such a thermocouple “head”:




As you can see from this photograph, the screws directly press against the solid metal thermocouple wires to make a firm connection between each wire and the brass terminal block. Since the “head” attaches directly to one end of the thermocouple, the thermocouple’s wires will be trimmed just long enough to engage with the terminal screws inside the head. Both brass terminal blocks are mounted on a ceramic base, the purpose of the ceramic being to help equalize the temperatures between the two brass blocks while still maintaining electrical isolation. This assembly is sometimes referred to as an *isothermal terminal block* because it acts to keep all connection points at a common temperature (“iso-thermal” = “same-temperature”). A threaded cover on the head provides easy access to these connection points for installation and maintenance, while ensuring the connections are covered and protected from ambient weather conditions the rest of the time.

Thermocouple wires are most often manufactured in solid form rather than stranded form. A common mistake made with thermocouple wires is for technicians to crimp compression-style terminals (“lugs”) onto the solid wires. While this may form a usable connection at first, compression-style terminals are simply unable to maintain adequate compression when applied to solid wire of any type, thermocouple wire included. Over time, solid wires *will* loosen inside compression terminals leading to circuit problems. In the case of a thermocouple circuit, bad wire connections lead to a situation where the receiving instrument “thinks” the thermocouple has failed open. This situation is commonly called *burnout*, referring to the phenomenon where a thermocouple junction fails open from being “burned out” by excessive temperature.

Fork terminal




Ring terminal



You will most often find compression terminals (improperly) applied to solid thermocouple wire tips where those wires must terminate under the head of a screw. Compression terminals are correct to use in applications where stranded wire terminates at a screw head, but not solid wire. The proper termination technique for solid wire under a screw head is to wrap the solid wire in a semi-circle and directly clamp it under the screw head.

At the other end of the thermocouple, we have a choice of tip styles. For maximum sensitivity and fastest response, the dissimilar-metal junction may be unsheathed (bare). This design, however, makes the thermocouple more fragile. Sheathed tips are typical for industrial applications, available in either *grounded* or *ungrounded* forms:




Grounded-tip thermocouples exhibit faster response times¹¹ and greater sensitivity than ungrounded-tip thermocouples, but they are vulnerable to *ground loops*: circuitous paths for electric current between the conductive sheath of the thermocouple and some other point in the thermocouple circuit. In order to avoid this potentially troublesome effect, most industrial thermocouples are of the ungrounded design.

¹¹Grounded thermocouples often have thermal time constant values less than *half* that of comparable ungrounded thermocouples. Exposed-tip thermocouples are even faster than grounded-tip, typically by even larger ratios than grounded-tip thermocouples are to ungrounded thermocouples.

21.4.4 Manually interpreting thermocouple voltages

Recall that the amount of voltage indicated by a voltmeter connected to a thermocouple is the *difference* between the voltage produced by the measurement junction (the point where the two dissimilar metals join at the location we desire to sense temperature at) and the voltage produced by the reference junction (the point where the thermocouple wires join to the voltmeter wires):



$$V_{meter} = V_{J1} - V_{J2}$$

This makes thermocouples inherently *differential* sensing devices: they generate a measurable voltage in proportion to the *difference* in temperature between two locations. This inescapable fact of thermocouple circuits complicates the task of interpreting any voltage measurement obtained from a thermocouple.

In order to translate a voltage measurement produced by a voltmeter connected to a thermocouple, we must *add* the voltage produced by the measurement junction (V_{J2}) to the voltage indicated by the voltmeter to find the voltage being produced by the measurement junction (V_{J1}). In other words, we manipulate the previous equation into the following form:

$$V_{J1} = V_{J2} + V_{meter}$$

We may ascertain the reference junction voltage by placing a thermometer near that junction (where the thermocouple wire attaches to the voltmeter test leads) and referencing a *thermocouple table* showing temperatures and corresponding voltages for that thermocouple type. Then, we may take the voltage sum for V_{J1} and re-reference that same table, finding the temperature value corresponding to the calculated measurement junction voltage. The National Institute of Standards and Technology (NIST) in the United States publishes tables showing junction voltages


and temperatures for standardized thermocouple types. While it is possible to mathematically model a thermocouple junction's voltage in the same way we may model an RTD's resistance, the functions for thermocouples are less linear than for RTDs, and so tables are greatly preferred for practical use.

To illustrate, suppose we connected a voltmeter to a type K thermocouple and measured 14.30 millivolts. A thermometer situated near the thermocouple wire / voltmeter junction point shows an ambient temperature of 73 degrees Fahrenheit. Referencing a table of voltages for type K thermocouples (in this case, the NIST "ITS-90" reference standard), we see that a type K junction at 73 degrees Fahrenheit corresponds to 0.910 millivolts. Adding this figure to our meter measurement of 14.30 millivolts, we arrive at a sum of 15.21 millivolts for the measurement ("hot") junction. Going back to the same table of values, we see 15.21 millivolts falls between 701 and 702 degrees Fahrenheit. Linearly interpolating between the table values (15.203 mV at 701 °F and 15.226 mV at 702 °F), we may more precisely determine the measurement junction to be 701.3 degrees Fahrenheit.

The process of manually taking voltage measurements, referencing a table of millivoltage values, performing addition, then re-referencing the same table is rather tedious. Compensation for the reference junction's inevitable presence in the thermocouple circuit is something we must do, but it is not something that must always be done by a human being. The next subsection discusses ways to automatically compensate for the effect of the reference junction, which is the only practical alternative for continuous thermocouple-based temperature instruments.

21.4.5 Reference junction compensation

Multiple techniques exist to deal with the influence of the reference junction's temperature¹². One technique is to physically fix the temperature of that junction at some constant value so it is always stable. This way, any changes in measured voltage *must* be due to changes in temperature at the measurement junction, since the reference junction has been rendered incapable of changing temperature. This may be accomplished by immersing the reference junction in a bath of ice and water, the ice/water mixture ensuring a stable temperature by means of water's *latent heat of fusion*¹³:



In fact, this is how thermocouple temperature/voltage tables are referenced: describing the amount of voltage produced for given temperatures at the measurement junction with the reference junction held at the freezing point of water ($0\text{ }^{\circ}\text{C} = 32\text{ }^{\circ}\text{F}$). With the reference junction maintained at the freezing point of water, and thermocouple tables referenced to that specific cold junction temperature, the voltmeter's indication will simply and directly correspond to the temperature of measurement junction J_1 at all times.


However, fixing the reference junction at the temperature of freezing water is impractical for any real thermocouple application outside of a laboratory. Instead, we need to find some other way to compensate for changes in reference junction temperature, so that we may accurately interpret the temperature of the measurement junction despite random changes in reference junction temperature.

¹²Early texts on thermocouple use describe multiple techniques for automatic compensation of the reference (“cold”) junction. One design placed a mercury bulb thermometer at the reference junction, with a loop of thin platinum wire dipped into the mercury. As junction temperature rose, the mercury column would rise and short past a greater length of the platinum wire loop, causing its resistance to decrease which in turn would electrically bias the measurement circuit to offset the effects of the reference junction's voltage. Another design used a bi-metallic spring to offset the pointer of the meter movement, so that changes in temperature at the indicating instrument (where the reference junction was located) would result in the analog meter's needle becoming offset from its normal “zero” point, thus compensating for the offset in voltage created by the reference junction.

¹³For any two-phase mixture of any single substance (in this case, H_2O) the temperature of that mixture will be a strict function of pressure, the mixture possessing only one thermodynamic degree of freedom. Any addition or removal of heat from the ice/water mix results in a phase change (e.g. either more ice melts to become water, or more water freezes to become ice) rather than a temperature change. If even more precision is desired, a *triple point cell* may be used to fix the reference junction's temperature, since any substance at its triple point will possess zero degrees of thermodynamic freedom (i.e. neither its pressure nor temperature can change).

A practical way to compensate for the reference junction voltage is to include an additional voltage source within the thermocouple circuit equal in magnitude and opposite in polarity to the reference junction voltage. If this additional voltage is made continually equal to the reference junction's potential, it will precisely counter the reference junction voltage, resulting in the full (measurement junction) voltage appearing at the measuring instrument terminals. This is called a *reference junction compensation* or *cold junction*¹⁴ compensation circuit:

**Compensating for the effects of J_2
using a “reference junction compensation”
source to generate a counter-voltage**



In order for such a compensation strategy to work, the compensating voltage must continuously track the voltage produced by the reference junction. To do this, the compensating voltage source (V_{rjc} in the above schematic) uses some other temperature-sensing device such as a thermistor or RTD to sense the local temperature at the terminal block where junction J_2 is formed and produce a counter-voltage that is precisely equal and opposite to J_2 's voltage ($V_{rjc} = V_{J2}$) at all times. Having canceled the effect of the reference junction, the voltmeter now only registers the voltage produced by the measurement junction J_1 :


$$V_{meter} = V_{J1} - V_{J2} + V_{rjc}$$

$$V_{meter} = V_{J1} + 0 \quad (\text{If } V_{rjc} = -V_{J2})$$

$$V_{meter} = V_{J1}$$

¹⁴Please note that “cold junction” is just a synonymous label for “reference junction.” In fact the “cold” reference junction may very well be at a warmer temperature than the so-called “hot” measurement junction! Nothing prevents anyone from using a thermocouple to measure temperatures below the freezing point of water.

Some instrument manufacturers sell *electronic ice point* modules designed to provide reference junction compensation for un-compensated instruments such as standard voltmeters. The “ice point” circuit performs the function shown by V_{rjc} in the previous diagram: it inserts a counter-acting voltage to cancel the voltage generated by the reference junction, so that the voltmeter only “sees” the measurement junction’s voltage. This compensating voltage is maintained at the proper value according to the terminal temperature where the thermocouple wires connect to the ice point module, sensed by a thermistor or RTD:



Example values:

$$\begin{aligned}
 T_{J1} &= 570 \text{ }^\circ\text{F (type J)} & V_{J1} &= 16.266 \text{ mV} \\
 T_{J2} &= 69 \text{ }^\circ\text{F (type J)} & V_{J2} &= 1.048 \text{ mV} & V_{rjc} &= 1.048 \text{ mV} \\
 V_{\text{meter}} &= V_{J1} - V_{J2} + V_{rjc} \\
 V_{\text{meter}} &= 16.266 \text{ mV} - 1.048 \text{ mV} + 1.048 \text{ mV} \\
 V_{\text{meter}} &= \mathbf{16.266 \text{ mV (equivalent to } 570 \text{ }^\circ\text{F)}}
 \end{aligned}$$

In this example, we see the measurement junction (J1) at a temperature of 570 degrees Fahrenheit, generating a voltage of 16.266 millivolts. If this thermocouple were directly connected to the meter, the meter would only register 15.218 millivolts, because the reference junction (J2, at 69 degrees Fahrenheit) opposes with its own voltage of 1.048 millivolts. With the ice point compensation circuit installed, however, the 1.048 millivolts of the reference junction is canceled by the ice point circuit’s equal-and-opposite 1.048 millivolt source. This allows the full 16.266 millivolt signal from the measurement junction reach the voltmeter where it may be read and correlated to temperature by a type J thermocouple table.


At first it may seem pointless to go through the trouble of building a reference junction compensation (ice point) circuit, when doing so requires the use of some other temperature-sensing element such as a thermistor or RTD. After all, why bother to do this just to be able to use a thermocouple to accurately measure temperature, when we could simply use this “other” device to directly measure the process temperature? In other words, isn’t the usefulness of a thermocouple invalidated if we must rely on some other type of electrical temperature sensor just to compensate

for an idiosyncrasy of the thermocouple?


The answer to this very good question is that thermocouples enjoy certain advantages over these other sensor types. Thermocouples are extremely rugged and have far greater temperature-measurement ranges than thermistors, RTDs, and other primary sensing elements. However, if the application does not demand extreme ruggedness or large measurement ranges, a thermistor or RTD *is* most likely the better choice!

21.4.6 Law of Intermediate Metals

It is crucial to realize that the phenomenon of a “reference junction” is an inevitable effect of having to close the electric circuit loop in a circuit made of dissimilar metals. This is true regardless of the number of metals involved. In the last example, only two metals were involved: iron and copper. This formed one iron-copper junction (J_1) at the measurement end and one iron-copper junction (J_2) at the indicator end. Recall that the copper-copper junction J_3 was of no consequence because its identical metallic composition generates no thermal voltage:




But what about more complex thermocouple circuits, involving more than two wire types? How do we define what a “reference junction” is, or how it behaves, when we have more than two dissimilar-metal junctions in the same circuit? Take for instance this example of a type J thermocouple:




Here we have *three* voltage-generating junctions: J_1 of iron and constantan, J_2 of iron and copper, and J_3 of copper and constantan¹⁵. Upon first inspection it would seem we have a much more complex situation than we did with just two metals (iron and copper), but fortunately the situation is just as simple as it was before provided the temperatures of J_2 and J_3 are equal, which will be true if those two junctions are located very near each other (at the voltmeter).

¹⁵A junction of copper and constantan just happens to be a *type T* thermocouple junction.

A principle of thermo-electric circuits called the *Law of Intermediate Metals* helps us see this clearly. According to this law, intermediate metals in a series of junctions are of no consequence to the overall (net) voltage so long as those intermediate junctions are all at the same temperature. Representing this pictorially, the net effect of having four different metals (A, B, C, and D) joined together in series is the same as just having the first and last metal in that series (A and D) joined together in one junction, if all intermediate junctions are at the same temperature:




A simple proof of the Law of Intermediate Metals may be built upon the *Law of Energy Conservation*, one of the most fundamental principles in all of physics. Consider what would happen if we were to join the series of dissimilar metal wires shown above into a continuous loop:



In this diagram we see that the wire made of metal “A” connects to a string of metal junctions formed by metals “B”, “C”, and “D”. If all these dissimilar metal junctions are at the same temperature, there will be no difference of temperature anywhere in the circuit to drive a current, and we would therefore expect the current in this circuit to be zero. This is in accordance with the Law of Energy Conservation, which forbids the passage of electric current through resistive wire without some motive power source driving it. Thus, based on the premise that energy must be conserved (i.e. that an electric current cannot flow through any resistance without a power source), we must conclude that *the net effect of all those series-connected metal junctions at the same temperature must be zero*. In other words, junctions A-B, B-C, C-D, and D-A all at the same temperature and connected in series *must* generate zero voltage, as if those junctions were all reduced to a single A-A junction which of course cannot produce any electromotive force (voltage) because it is not comprised of dissimilar metals. If the Law of Intermediate Metals were untrue, it would mean that the junctions A-B-C-D were *not* equivalent to the single junction A-D, which would mean they would produce a different voltage than the D-A junction at the right-hand end of this circuit (while at the same temperature), and therefore this circuit would produce some net voltage to drive a current continuously through resistive wire in violation of the Law of Energy Conservation. Since we know the Law of Energy Conservation to be well-founded (and we can also *build* such dissimilar metal


loop circuits and empirically determine their currents to be zero), we may rest assured that the Law of Intermediate Metals is true.

In our type J thermocouple circuit where iron and constantan both join to copper, we see copper as an intermediate metal between junctions J_2 and J_3 . Being located next to each other on the indicating instrument, identical temperature is a reasonable assumption for J_2 and J_3 , so we may invoke the Law of Intermediate Metals and simply treat junctions J_2 and J_3 as a *single* iron-constantan reference junction. In other words, the Law of Intermediate Metals tells us we can treat the following two circuits identically:



The practical importance of this Law is that we can always treat the reference junction(s) as a single junction made from the same two metal types as the measurement junction, so long as all dissimilar metal junctions at the reference location are at the same temperature.

This fact is extremely important in the age of semiconductor circuitry, where the connection of a thermocouple to an electronic amplifier involves a long series of dissimilar-material junctions. Here we see a multitude of reference junctions, formed by the necessary connections from thermocouple wire to the silicon substrate inside the amplifier chip:




Here we see the metals of the thermocouple wire (type J – iron and constantan) joining to a pair of brass terminal screws, which in turn join to copper traces on a printed circuit board, which join to lead/tin solder, which join to thin wires made of Kovar, which terminate at gold traces on the silicon chip, which are bonded to the silicon itself.

It should be obvious that each complementary junction pair in this series loop cancel each other if each pair is at the same temperature (e.g. gold-silicon junction J_{12} cancels with silicon-gold junction J_{13} because they generate the exact same amount of voltage with opposing polarities; Kovar-gold junction J_{10} cancels with gold-Kovar junction J_{11} for the same reason; etc.). The Law of Intermediate Metals goes one step further by telling us junctions J_2 through J_{13} taken together in series are of the same effect as a single reference junction of iron and constantan. Automatic reference junction compensation is as simple as counter-acting the voltage produced by this equivalent iron-constantan junction at whatever temperature junctions J_2 through J_{13} happen to be at.

21.4.7 Software compensation

Previously, it was suggested that automatic compensation could be accomplished by intentionally inserting a temperature-dependent voltage source in series with the circuit, oriented in such a way as to oppose the reference junction's voltage:

*Compensating for the effects of J_2
using a "reference junction compensation"
source to generate a counter-voltage*



$$V_{meter} = V_{J1} - V_{J2} + V_{rjc}$$


If the series voltage source V_{rjc} is exactly equal in magnitude to the reference junction's voltage (V_{J2}), those two terms cancel out of the equation and lead to the voltmeter measuring only the voltage of the measurement junction J_1 :

$$V_{meter} = V_{J1} + 0$$

$$V_{meter} = V_{J1}$$

This technique is known as *hardware* compensation, and is employed in analog thermocouple temperature transmitter designs. Previously we saw an example of this called an *ice point*, the purpose of which was to electrically counter the reference junction voltage to render that junction's voltage inconsequential as though that junction were immersed in a bath of ice-water.

A modern technique for reference junction compensation more suitable to digital transmitter designs is called *software* compensation:



Instead of canceling the effect of the reference junction *electrically*, we cancel the effect *arithmetically* inside the microprocessor-based transmitter. In other words, we let the receiving analog-digital converter circuit see the difference in voltage between the measurement and reference junctions ($V_{input} = V_{J1} - V_{J2}$), but then after digitizing this voltage measurement we have the microprocessor add the equivalent voltage value corresponding to the ambient temperature sensed by the RTD or thermistor (V_{rjc}):

$$\text{Compensated total} = V_{input} + V_{rjc}$$

$$\text{Compensated total} = (V_{J1} - V_{J2}) + V_{rjc}$$


Since we know the calculated value of V_{rjc} should be equal to the real reference junction voltage (V_{J2}), the result of this digital addition should be a compensated total equal only to the measurement junction voltage V_{J1} :

$$\text{Compensated total} = V_{J1} - V_{J2} + V_{rjc}$$

$$\text{Compensated total} = V_{J1} + 0$$

$$\text{Compensated total} = V_{J1}$$

A block diagram of a thermocouple temperature transmitter with software compensation appears here:




Perhaps the greatest advantage of software compensation is the flexibility to easily switch between different thermocouple types with no hardware modification. So long as the microprocessor memory is programmed with look-up tables relating voltage values to temperature values, it may accurately measure (and compensate for the reference junction of) *any* thermocouple type. Hardware-based compensation schemes (e.g. an analog "ice point" circuit) require re-wiring or replacement to accommodate different thermocouple types, since each ice-point circuit is built to generate a compensating voltage for a specific type of thermocouple.

21.4.8 Extension wire

In every thermocouple circuit there must be both a measurement junction and a reference junction: this is an inevitable consequence of forming a complete circuit (loop) using dissimilar-metal wires. As we already know, the voltage received by the measuring instrument from a thermocouple will be the *difference* between the voltages produced by the measurement and reference junctions. Since the purpose of most temperature instruments is to accurately measure temperature *at a specific location*, the effects of the reference junction's voltage must be "compensated" for by some means, either a special circuit designed to add an additional canceling voltage or by a software algorithm to digitally cancel the reference junction's effect.


In order for reference junction compensation to be effective, the compensation mechanism must "know" the temperature of the reference junction. This fact is so obvious, it hardly requires mentioning. However, what is not so obvious is how easily this compensation may be unintentionally defeated simply by installing a different type of wire in a thermocouple circuit.

To illustrate, let us examine a simple type K thermocouple installation, where the thermocouple connects directly to a panel-mounted temperature indicator by long wires:



Like all modern thermocouple instruments, the panel-mounted indicator contains its own internal reference junction compensation, so that it is able to compensate for the temperature of the reference junction formed at its connection terminals, where the internal (copper) wires of the indicator join to the chromel and alumel wires of the thermocouple. The indicator senses this junction temperature using a small thermistor thermally bonded to the connection terminals.


Now let us consider the same thermocouple installation with a length of copper cable (two wires) joining the field-mounted thermocouple to the panel-mounted indicator:



Even though nothing has changed in the thermocouple circuit except for the type of wires joining the thermocouple to the indicator, the reference junction has completely shifted position. What used to be a reference junction (at the indicator's terminals) is no longer, because now we have copper wires joining to copper wires. Where there is no dissimilarity of metals, there can be no thermoelectric potential. At the thermocouple's connection "head," however we now have a joining of chromel and alumel wires to copper wires, thus forming a reference junction *in a new location* at the thermocouple head. What is worse, this new location is likely to be at a different temperature than the panel-mounted indicator, which means the indicator's reference junction compensation will be compensating for the wrong temperature.

The only practical way to avoid this problem is to keep the reference junction where it belongs: at the terminals of the panel-mounted instrument where the ambient temperature is measured and the reference junction's effects accurately compensated. If we must install "extension" wire to join a thermocouple to a remotely-located instrument, that wire must be of a type that does not form another dissimilar-metal junction at the thermocouple head, but will form one at the receiving instrument.

An obvious approach is to simply use thermocouple wire of the same type as the installed thermocouple to join the thermocouple to the indicator. For our hypothetical type K thermocouple, this means a type K cable installed between the thermocouple head and the panel-mounted indicator:




With chromel joining to chromel and alumel joining to alumel at the head, no dissimilar-metal junctions are created at the thermocouple. However, with chromel and alumel joining to copper at the indicator (again), the reference junction has been re-located to its rightful place. This means the thermocouple head's temperature will have no effect on the performance of this measurement system, and the indicator will be able to properly compensate for any ambient temperature changes at the panel as it was designed to do. The only problem with this approach is the potential expense of thermocouple-grade cable. This is especially true with some types of thermocouples, where the metals used are somewhat exotic (e.g. types R, S, and B).

A more economical alternative, however, is to use something called *extension-grade* wire to make the connection between the thermocouple and the receiving instrument. "Extension-grade" thermocouple wire is made less expensive than full "thermocouple-grade" wire by choosing metal alloys similar in thermo-electrical characteristics to the real thermocouple wires within modest temperature ranges. So long as the temperatures at the thermocouple head and receiving instrument terminals don't get too hot or too cold, the extension wire metals joining to the thermocouple wires and joining to the instrument's copper wires need not be *precisely identical* to the true thermocouple wire alloys. This allows for a wider selection of metal types, some of which are substantially less expensive than the measurement-grade thermocouple alloys. Also, extension-grade wire may use insulation with a narrower temperature rating than thermocouple-grade wire, reducing cost even

further.

An interesting historical reference to the use of extension-grade wire appears in Charles Robert Darling's 1911 text *Pyrometry – A Practical Treatise on the Measurement of High Temperatures*. On page 61, Darling describes “compensating leads” marketed under the brand-name of *Peake* designed to be used with platinum-alloy thermocouples. These “compensating” wires were made of two different copper-nickel alloys, each copper-nickel alloy matched with the respective thermocouple metal (in this case, pure platinum and a 90%-10% platinum-iridium alloy) to generate an equal and opposite millivoltage at any reasonable temperature found at the thermocouple head. Thus, the only reference junction in the thermocouple circuit is where these copper-nickel extension wires joined with the indicating instrument, rather than being located at the thermocouple head as it would be if simple copper extension wires were employed. With platinum being such an expensive metal (both then and now!), the cost savings realized by being able to use cheaper extension wire to connect the platinum thermocouple to a distant receiving instrument is significant.

Extension-grade cable is denoted by a letter “X” following the thermocouple letter. For our hypothetical type K thermocouple system, this would mean type “KX” extension cable:



Thermocouple extension cable also differs from thermocouple-grade (measurement) cable in the coloring of its outer jacket. Whereas thermocouple-grade cable is typically¹⁶ brown in exterior color, extension-grade cable is usually colored¹⁷ to match the thermocouple plug (yellow for type K, black for type J, blue for type T, etc.).


¹⁶No coloring standard exists in the United States for platinum thermocouple-grade wire (e.g. types R, S, etc.).

¹⁷The colors I list here are for thermocouples in the United States.

21.4.9 Side-effects of reference junction compensation


Reference junction compensation is a necessary part of any precision thermocouple circuit, due to the inescapable fact of the reference junction's existence. When you form a complete circuit of dissimilar metals, you *will* form both a measurement junction and a reference junction, with those two junctions' polarities opposed to one another. This is why reference junction compensation – whether it takes the form of a hardware circuit or an algorithm in software – must exist within every precision thermocouple instrument.

The presence of reference junction compensation in every precision thermocouple instrument results in an interesting phenomenon: *if you directly short-circuit the thermocouple input terminals of such an instrument, it will always register ambient temperature, regardless of the thermocouple type the instrument is built or configured for.* This behavior may be illustrated by example, first showing a normal operating temperature measurement system and then with that same system short-circuited. Here we see a temperature indicator receiving a 4-20 mA current signal from a temperature transmitter, which is receiving a millivoltage signal from a type “K” thermocouple sensing a process temperature of 780 degrees Fahrenheit:



The transmitter's internal reference junction compensation feature compensates for the ambient temperature of 68 degrees Fahrenheit. If the ambient temperature rises or falls, the compensation will automatically adjust for the change in reference junction potential, such that the output will still register the process (measurement junction) temperature of 780 degrees F. This is what the reference junction compensation is designed to do.

Now, we disconnect the thermocouple from the temperature transmitter and short-circuit the transmitter's input:



With the input short-circuited, the transmitter “sees” no voltage at all from the thermocouple circuit. There is no measurement junction nor a reference junction to compensate for, just a piece of wire making both input terminals electrically common. This means the reference junction compensation inside the transmitter no longer performs a useful function. However, the transmitter does not “know” it is no longer connected to the thermocouple, so the compensation keeps on working even though it has nothing to compensate for. Recall the voltage equation relating measurement, reference, and compensation voltages in a hardware-compensated thermocouple instrument:

$$V_{meter} = V_{J1} - V_{J2} + V_{rjc}$$

Disconnecting the thermocouple wire and connecting a shorting jumper to the instrument eliminates the V_{J1} and V_{J2} terms, leaving only the compensation voltage to be read by the meter¹⁸:

$$V_{meter} = 0 + V_{rjc}$$


$$V_{meter} = V_{rjc}$$

¹⁸The effect will be exactly the same for an instrument with software compensation rather than hardware compensation. With software compensation, there is no literal V_{rjc} voltage source, but the equivalent millivolt value is digitally added to the zero input measured at the thermocouple connection terminals, resulting in the same effect of measuring ambient temperature.

This is why the instrument registers the equivalent temperature created by the reference junction compensation feature: this is the only signal it “sees” with its input short-circuited. This phenomenon is true regardless of which thermocouple type the instrument is configured for, which makes it a convenient “quick test” of instrument function in the field. If a technician short-circuits the input terminals of any thermocouple instrument, it should respond as though it is sensing ambient temperature.


While this interesting trait is a somewhat useful side-effect of reference junction compensation in thermocouple instruments, there are other effects that are not quite so useful. The presence of reference junction compensation becomes quite troublesome, for example, if one tries to simulate a thermocouple using a precision millivoltage source. Simply setting the millivoltage source to the value corresponding to the desired (simulation) temperature given in a thermocouple table will yield an incorrect result for any ambient temperature other than the freezing point of water!

Suppose, for example, a technician wished to simulate a type K thermocouple at 300 degrees Fahrenheit by setting a millivolt source to 6.094 millivolts (the voltage corresponding to 300 °F for type K thermocouples according to the ITS-90 standard). Connecting the millivolt source to the instrument will *not* result in an instrument response appropriate for 300 degrees F:



Instead, the instrument registers 339 degrees because its internal reference junction compensation feature is still active, compensating for a reference junction voltage that no longer exists. The millivolt source’s output of 6.094 mV gets *added* to the compensation voltage (inside the transmitter) of 0.865 mV – the necessary millivolt value to compensate for a type K reference junction at 71 °F – with the result being a larger millivoltage (6.959 mV) interpreted by the transmitter as a temperature of 339 °F.


One way to use a millivoltage source to simulate a desired temperature is for the instrument technician to “out-think” the transmitter’s compensation feature by specifying a millivolt signal that is offset by the amount of equivalent voltage generated by the transmitter’s compensation. In other words, instead of setting the millivolt source to a value of 6.094 mV, the technician should set the source to only 5.229 mV so the transmitter’s compensation will add 0.865 mV to this value to arrive at 6.094 mV and properly register as 300 degrees Fahrenheit:



Years ago, the only suitable piece of test equipment available for generating the precise millivoltage signals necessary to calibrate thermocouple instruments was a device called a *precision potentiometer*. These “potentiometers” used a stable *mercury cell* battery (sometimes called a *standard cell*) as a voltage reference and a potentiometer with a calibrated knob to output low-voltage signals. Photographs of two vintage precision potentiometers are shown here:



Of course, modern thermocouple calibrators also provide direct entry of temperature and automatic compensation to “un-compensate” the transmitter such that any desired temperature may be easily simulated:



In this example, when the technician sets the calibrator for 300 °F (type K), it measures the ambient temperature and automatically subtracts 0.865 mV from the output signal, so only 5.229 mV is sent to the transmitter terminals instead of the full 6.094 mV. The transmitter’s internal reference junction compensation adds the 0.865 mV offset value (thinking it must compensate for a reference junction that in reality is not there) and “sees” a total signal voltage of 6.094 mV, interpreting this properly as 300 degrees Fahrenheit.

The following photograph shows the display of a modern thermocouple calibration device (a Fluke model 744 documenting process calibrator) being used to generate a thermocouple signal. In this particular example, the thermocouple type is set to type “S” (Platinum-Rhodium/Platinum) at a temperature of 2650 degrees Fahrenheit:




The ITS-90 thermocouple standard declares a millivoltage signal value of 15.032 mV for a type S thermocouple junction at 2650 degrees F (with a reference junction temperature of 32 degrees F). Note how the calibrator does *not* output 15.032 mV even though the simulated temperature has been set to 2650 degrees F. Instead, it outputs 14.910 mV, which is 0.122 mV less than 15.032 mV. This offset of 0.122 mV corresponds to the calibrator’s local temperature of 70.8 degrees F (according to the ITS-90 standard for type S thermocouples).

When the calibrator’s 14.910 mV signal reaches the thermocouple instrument being calibrated (be it an indicator, transmitter, or even a controller equipped with a type S thermocouple input), the instrument’s own internal reference junction compensation will add 0.122 mV to the received signal of 14.910 mV, “thinking” it needs to compensate for a real reference junction. The result will be a perceived measurement junction signal of 15.032 mV, which is exactly what we want the instrument to “think” it sees if our goal is to simulate connection to a real type S thermocouple at a temperature of 2650 degrees F.

21.4.10 Burnout detection

Another consideration for thermocouples is *burnout detection*. The most common failure mode for thermocouples is to fail open, otherwise known as “burning out.” An open thermocouple is problematic for any voltage-measuring instrument with high input impedance because the lack of a complete circuit on the input makes it possible for electrical noise from surrounding sources (power lines, electric motors, variable-frequency motor drives) to be detected by the instrument and falsely interpreted as a wildly varying temperature.

For this reason it is prudent to design into the thermocouple instrument some provision for generating a consistent state in the absence of a complete circuit. This is called the *burnout mode* of a thermocouple instrument. A simple thermocouple circuit equipped with burnout detection is shown in this diagram:



The resistor in this circuit provides a path for current in the event of an open thermocouple. It is sized in the mega-ohm range to minimize its effect during normal operation when the thermocouple circuit is complete. Only when the thermocouple fails open will the miniscule current through the resistor have any substantial effect on the voltmeter’s indication. The SPDT switch provides a selectable burnout mode: in the event of a burnt-out thermocouple, we can configure the meter to either read high temperature (sourced by the instrument’s internal milli-voltage source) or low temperature (grounded), depending on what failure mode we deem safest for the application.

21.5 Non-contact temperature sensors

Virtually any mass above absolute zero temperature emits electromagnetic radiation (photons, or light) as a function of that temperature. This basic fact makes possible the measurement of temperature by analyzing the light emitted by an object. The *Stefan-Boltzmann Law* of radiated energy quantifies this fact, declaring that the rate of heat lost by radiant emission from a hot object is proportional to the fourth power of the absolute temperature:

$$\frac{dQ}{dt} = e\sigma AT^4$$

Where,

$\frac{dQ}{dt}$ = Radiant heat loss rate (watts)

e = Emissivity factor (unitless)

σ = Stefan-Boltzmann constant (5.67×10^{-8} W / m² · K⁴)

A = Surface area (square meters)

T = Absolute temperature (Kelvin)

The primary advantage of non-contact thermometry (or *pyrometry* as high-temperature measurement is often referred) is rather obvious: with no need to place a sensor in direct contact with the process, a wide variety of temperature measurements may be made that are either impractical or impossible to make using any other technology.

A major disadvantage of non-contact thermometry is that it only reveals the *surface temperature* of an object. Sensing the thermal radiation emanated from a pipe, for instance, only tells you the surface temperature of that pipe and not the true temperature of the fluid within the pipe. Another example is when doctors use non-contact thermometry to assess irregularities in body temperature: what they detect is just skin temperature. While it may be true that “hot spots” beneath the surface of an object may be detectable this way, it is only because the surface temperature of that object differs as a consequence of the hot spot(s) beneath. If a hotter-than-normal region inside of an object fails to transfer enough thermal energy to the surface to manifest as a hotter surface temperature, that region will be invisible to non-contact thermometry.


It may surprise some readers to discover that non-contact pyrometry is nearly as old as thermocouple technology¹⁹, the first non-contact pyrometer being constructed in 1892.

¹⁹Although Seebeck discovered thermo-electricity in 1822, the technique of measuring temperature by sensing the voltage produced at a dissimilar-metal junction was delayed in practical development until 1886 when rugged and accurate electrical meters became available for industrial use.

21.5.1 Concentrating pyrometers

A time-honored design for non-contact pyrometers is to concentrate²⁰ incident light from the surface of a heated object onto a small temperature-sensing element. A rise in temperature at the sensor reveals the intensity of the infrared optical energy falling upon it, which as discussed previously is a function of the target object's surface temperature (absolute temperature to the fourth power):

Two designs of non-contact pyrometer




The fourth-power characteristic of Stefan-Boltzmann's law means that a doubling of absolute temperature at the hot object results in sixteen times as much radiant energy falling on the sensor, and therefore a sixteen-fold increase in the sensor's temperature rise over ambient. A tripling of target temperature (absolute) yields *eighty one times as much radiant energy*, and therefore an 81-fold increase in sensor temperature rise. This extreme nonlinearity limits the practical application of non-contact pyrometry to relatively narrow ranges of target temperature wherever good accuracy is required.

Thermocouples were the first type of sensor used in non-contact pyrometers, and they still find application in modern versions of the same technology. Since the sensor does not become nearly as hot as the target object, the output of any single thermocouple junction at the sensor area will be quite small. For this reason, instrument manufacturers often employ a series-connected array of thermocouples called a *thermopile* to generate a stronger electrical signal.

²⁰ Anyone who has ever used a magnifying glass (a concentrating lens) to concentrate sunlight knows how this works. If you were to use a magnifying glass to concentrate sunlight onto a thermocouple-type sensor, you could (at least in principle) infer the temperature of the sun in this manner.

The basic concept of a thermopile is to connect multiple thermocouple junctions in series so their voltages will add:



Examining the polarity marks of each junction (type E thermocouple wires are assumed in this example: chromel and constantan), we see that all the “hot” junctions’ voltages aid each other, as do all the “cold” junctions’ voltages. Like all thermocouple circuits, though, the each “cold” junction voltage opposes each the “hot” junction voltage. The example thermopile shown in this diagram, with four hot junctions and four cold junctions, will generate four times the potential difference that a single type E thermocouple hot/cold junction pair would generate, assuming all the hot junctions are at the same temperature and all the cold junctions are at the same temperature.

When used as the detector for a non-contact pyrometer, the thermopile is oriented so all the concentrated light falls on the hot junctions (the “focal point” where the light focuses to a small spot), while the cold junctions face away from the focal point to a region of ambient temperature. Thus, the thermopile acts like a multiplied thermocouple, generating more voltage than a single thermocouple junction could under the same temperature conditions.

A popular design of non-contact pyrometer manufactured for years by Honeywell was the *Radiamatic*²¹, using ten thermocouple junction pairs arrayed in a circle. All the “hot” junctions were placed at the center of this circle where the focal point of the concentrated light fell, while all the “cold” junctions were situated around the circumference of the circle away from the heat of the focal point. A table of values showing the approximate relationship between target temperature and millivolt output for one model of Radiamatic sensing unit reveals the fourth-power function:

Target temperature (K)	Millivolt output
4144 K	34.8 mV
3866 K	26.6 mV
3589 K	19.7 mV
3311 K	14.0 mV
3033 K	9.9 mV
2755 K	6.6 mV
2478 K	4.2 mV
2200 K	2.5 mV
1922 K	1.4 mV
1644 K	0.7 mV

We may test the basic²² validity of the Stefan-Boltzmann law by finding the ratio of temperatures for any two temperature values in this table, raising that ratio to the fourth power, and seeing if the millivolt output signals for those same two temperatures match the new ratio. The operating theory here is that increases in target temperature will produce fourth-power increases in sensor temperature rise, since the sensor’s temperature rise should be a direct function of radiation *power* impinging on it.

For example, if we were to take 4144 K and 3033 K as our two test temperatures, we find that the ratio of these two temperature values is 1.3663. Raising this ratio to the fourth power gives us 3.485 for a predicted ratio of millivolt values. Multiplying the 3033 K millivoltage value of 9.9 mV by 3.485 gives us 34.5 mV, which is quite close to the value of 34.8 mV advertised by Honeywell:

$$\frac{4144 \text{ K}}{3033 \text{ K}} = 1.3663$$

$$\left(\frac{4144 \text{ K}}{3033 \text{ K}}\right)^4 = 1.3663^4 = 3.485$$

$$(3.485)(9.9 \text{ mV}) \approx 34.8 \text{ mV}$$

If accuracy is not terribly important, and if the range of measured temperatures for the process is modest, we may take the millivoltage output of such a sensor and interpret it *linearly*. When

²¹Later versions of the Radiamatic (dubbed the *Radiamatic II*) were more than just a bare thermopile and optical concentrator, containing electronic circuitry to output a linearized 4-20 mA signal representing target temperature.

²²Comparing temperature ratios versus thermopile millivoltage ratios assumes linear thermocouple behavior, which we know is not exactly true. Even if the thermopile focal point temperatures precisely followed the ratios predicted by the Stefan-Boltzmann law, we would still expect some inconsistencies due to the non-linearities of thermocouple voltages. There will also be variations from predicted values due to shifts in radiated light frequencies, changes in emissivity factor, thermal losses within the sensing head, and other factors that refuse to remain constant over wide ranges of received radiation intensity. The lesson here is to not expect perfect agreement with theory!

used in this fashion, a non-contact pyrometer is often referred to as an *infrared thermocouple*, with the output voltage intended to connect directly to a thermocouple-input instrument such as an indicator, transmitter, recorder, or controller. An example of this usage is the OS-36 line of infrared thermocouples manufactured by Omega.

Infrared thermocouples are manufactured for a narrow range of temperature (most OS-36 models limited to a calibration span of 100 °F or less), their thermopiles designed to produce millivolt signals corresponding to a standard thermocouple type (T, J, K, etc.) over that narrow range.


21.5.2 Distance considerations

An interesting and useful characteristic of non-contact pyrometers is that their calibration does *not* depend on the distance separating the sensor from the target object's surface²³. This is counter-intuitive to anyone who has ever stood near an intense radiative heat source: standing in close proximity to a bonfire, for example, results in much hotter skin temperature than standing far away from it. Why wouldn't a non-contact pyrometer register cooler target temperatures when it was far away, given the fact that infrared radiation from the object spreads out with increased separation distance? The fact that an infrared pyrometer does not suffer from this limitation is good for our purposes in measuring temperature, but it doesn't seem to make sense at first.

One key to understanding this paradox is to quantify the bonfire experience, where perceived temperature falls off with increased distance. In physics, this is known as the *inverse square law*: the intensity of radiation falling on an object from a point-source decreases with the *square* of the distance separating the radiation source from the object. Backing away to twice the distance from a bonfire results in a *four-fold* decrease in received infrared radiation; backing way to three times the distance results in a *nine-fold* decrease in received radiation.


²³An important caveat to this rule is *so long as the target object completely fills the sensor's field of view (FOV)*. The reason for this caveat will become clear at the conclusion of the explanation.

Placing a sensor at three integer distances (x , $2x$, and $3x$) from a radiation point-source results in relative power levels of 100%, 25% (one-quarter), and 11.1% (one-ninth) falling upon the sensor at those locations, respectively:



This is a basic physical principle for all kinds of radiation, grounded in simple geometry. If we examine the radiation flux emanating from a point-source, we find that it must spread out as it travels in straight lines, and that the spreading-out happens at a rate defined by the square of the distance. An analogy for this phenomenon is to imagine a spherical latex balloon expanding as air is blown into it. The surface area of the balloon is proportional to the square of its radius. Likewise, the radiation flux emanating from a point-source spreads out in straight lines, in all directions, reaching a total area proportional to the square of the distance from the point (center). The total flux measured as a sphere will be the same no matter what the distance from the point-source, but the area it is divided over increases with the square of the distance, and so any object of fixed area backing away from a point-source of radiation encounters a smaller and smaller fraction of that flux.

If non-contact pyrometers really were “looking” at a point-source of infrared radiation, their signals *would* indeed decrease with distance. The saving grace here is that non-contact pyrometers are focused-optic devices, with a definite *field of view*, and that field of view should always be completely filled by the target object (assumed to be at a uniform temperature). As distance between the pyrometer and the target object changes, the cone-shaped field of view covers a surface area on that object proportional to the square of the distance²⁴. Backing the pyrometer away to twice the distance increases the viewing area on the target object by a factor of four; backing away to three times the distance increases the viewing area nine times:




So, even though the inverse square law correctly declares that radiation emanating from the hot wall (which may be thought of as a collection of point-sources) decreases in intensity with the square of the distance, this attenuation is perfectly balanced by an increased viewing area of the pyrometer. Doubling the separation distance *does* result in the flux from any given point on the wall spreading out by a factor of four, but the pyrometer now sees *four times as many similar points on the wall* as it did previously. So long as all the points within the field of view are uniform in temperature, the result is a perfect cancellation with the pyrometer providing the exact same temperature measurement at *any* distance from the target.

If the sensor’s field of view expands far enough to capture objects *other* than the one whose temperature we intend to measure, measurement errors will result. The sensor will now yield a weighted average of all objects within its field of view, and so it is important to ensure that field is limited to cover just the object we intend to measure.

²⁴The field of view (a circle where the viewing “cone” intercepts the flat surface of the object) increases linearly in diameter with increases in distance between the sensor and the object. However, since the area of a circle is proportional to the square of its diameter ($A = \frac{\pi D^2}{4}$ or $A = \pi r^2$), we may say that the viewing area increases with the square of the distance between the sensor and object.

Non-contact sensor fields-of-view are typically specified either as an angle, as a distance ratio, or both. For example, the following illustration shows a non-contact temperature sensor with a 5:1 (approximately 11°) field of view:




The mathematical relationship between viewing angle (θ) and distance ratio (D/d) follows the tangent function:

$$\frac{D}{d} = \frac{1}{2 \tan\left(\frac{\theta}{2}\right)} \quad \theta = 2 \tan^{-1}\left(\frac{d}{2D}\right)$$

A sampling of common field-of-view distance ratios and approximate viewing angles appears in this table:

Distance ratio	Angle (approximate)
1:1	53°
2:1	30°
3:1	19°
5:1	11°
7:1	8°
10:1	6°

A trigonometric explanation for these equations is shown in the following diagram, where the isosceles field-of-vision triangle is split into two “right” triangles, each one having an adjacent side length of D and an opposite side length of $d/2$ for angle $\theta/2$:



Since we know the tangent function is the ratio of opposite to adjacent side lengths for a right triangle, this means the tangent of the half-angle ($\theta/2$) will be equal to the ratio of the opposite side length ($d/2$) to the adjacent side length (D):

$$\tan\left(\frac{\theta}{2}\right) = \frac{d/2}{D} = \frac{d}{2D}$$

Solving for the length ratio D/d is now just a matter of algebraically manipulating this equation:

$$\tan\left(\frac{\theta}{2}\right) = \frac{d}{2D}$$

$$\frac{2D}{d} = \frac{1}{\tan\left(\frac{\theta}{2}\right)}$$

$$\frac{D}{d} = \frac{1}{2 \tan\left(\frac{\theta}{2}\right)}$$

Solving for the viewing angle (θ) requires another form of manipulation on the basic tangent equation, where we “un-do” the tangent function by using the *inverse* tangent (or “arctangent”) function:

$$\tan\left(\frac{\theta}{2}\right) = \frac{d}{2D}$$

$$\tan^{-1}\left[\tan\left(\frac{\theta}{2}\right)\right] = \tan^{-1}\left(\frac{d}{2D}\right)$$

$$\frac{\theta}{2} = \tan^{-1}\left(\frac{d}{2D}\right)$$

$$\theta = 2 \tan^{-1}\left(\frac{d}{2D}\right)$$

21.5.3 Emissivity

Aside from their inherent nonlinearity, perhaps the main disadvantage of non-contact temperature sensors is their inaccuracy. The emissivity factor (e) in the Stefan-Boltzmann equation varies with the composition of a substance, but beyond that there are several other factors (surface finish, shape, etc.) that affect the amount of radiation a sensor will receive from an object. For this reason, emissivity is not a very practical way to gauge the effectiveness of a non-contact pyrometer. Instead, a more comprehensive measure of an object's "thermal-optical measureability" is *emittance*.

A perfect emitter of thermal radiation is known as a *blackbody*. Emittance for a blackbody is unity (1), while emittance figures for any real object is a value between 1 and 0. The only certain way to know the emittance of an object is to test that object's thermal radiation at a known temperature. This assumes we have the ability to measure that object's temperature by direct contact, which of course renders void one of the major purposes of non-contact thermometry: to be able to measure an object's temperature without having to touch it. Not all hope is lost, though: all we have to do is obtain an emittance value for that object *one time*, and then we may calibrate any non-contact pyrometer for that object's particular emittance so as to measure its temperature in the future without contact.

Beyond the issue of emittance, other idiosyncrasies plague non-contact pyrometers. Objects also have the ability to *reflect* and *transmit* radiation from other bodies, which taints the accuracy of any non-contact device sensing the radiation from that body. An example of the former is trying to measure the temperature of a silver mirror using an optical pyrometer: the radiation received by the pyrometer is mostly from other objects, merely *reflected* by the mirror. An example of the latter is trying to measure the temperature of a gas or a clear liquid, and instead primarily measuring the temperature of a solid object in the background (*through* the gas or liquid).


Nevertheless, non-contact pyrometers have been and will continue to be useful in specific applications where other, contact-based temperature measurement techniques are impractical.

21.5.4 Thermal imaging

A very useful application of non-contact sensor technology is *thermal imaging*, where a dense array of infrared radiation sensors provides a graphic display of objects in its view according to their temperatures. Each object shown on the digital display of a thermal imager is artificially colored in the display on a chromatic scale that varies with temperature, hot objects typically registering as red tones and cold objects typically registering as blue tones. Thermal imaging is very useful in the electric power distribution industry, where technicians may inspect power line insulators and other objects at elevated potential for "hot spots" without having to make physical contact with those objects. Thermal imaging is also useful in performing "energy audits" of buildings and other heated structures, providing a means of revealing points of heat escape through walls, windows, and roofs. In such applications, relative differences in temperature are often more important to detect than specific temperature values. "Hot spots" readily appear on a thermal imager display, and may give useful data on the test subject even in the absence of accurate temperature measurement at any one spot.

Again, it is important to stress that thermal imaging only provides an assessment of the object's *surface* temperature, and not the temperature *within* that object. Variations in surface temperature detectable by thermal imaging are a secondary effect of temperature variations within the object, and as such may not accurately depict the true thermal gradient(s) within the object.

A digital image taken with a thermal imaging instrument by maintenance personnel at a municipal water treatment facility shows “hot spots” on an electric motor. A color scale on the right-hand side of the image serves as a “legend” for interpreting color as temperature. In this particular shot, dark blue is 68.1 °F and white is 152 °F:



This particular electric motor is in a vertical orientation, with the electrical connection box in the upper-left corner and two prominent hot spots on both the near and the left-facing sides of the case. A manual valve handle appears in the foreground, silhouetted in dark blue against a lighter blue (warmer) background. A lifting “eye” on the motor case appears as a green (cooler) shape in the middle of a white (warmer) area. The two “hot spots” correspond directly to stator windings and magnetic pole faces inside the motor, which are close enough to the motor’s casing to cause variations in surface temperature.

Thermal imaging is particularly useful for detecting hot spots on equipment unsafe to directly touch, as is the case with many “live” electrical components. This next thermal image shows an operating three-phase motor starter (contactor and overload block):



The bright spot in the center of the contactor is the higher temperature of the electromagnetic coil, providing magnetic force to actuate the contactor mechanism. Perhaps the most telling detail of this thermal image, however, is the difference in temperature between the overload heater connections (the six screws located near the bottom of the starter assembly). Note how the middle heater’s screws register slightly higher temperatures than the screws on either of the other two heater elements. Large temperature differences may indicate poor electrical connections (i.e. greater resistance) at the hot spots, or imbalances in phase current. It is also possible that the elevated temperature of this particular overload heater is simply due to it having less open surface area for it to radiate heat, since the two overload heaters flanking it enjoy the advantage of having more air cooling. If three people pack themselves into a narrow bench seat, the middle person is going to be warmer than either of the outer two!

Another noteworthy detail in this image is the “Spot Difference” measurement provided by the

thermal imager. A cross-hair cursor on the display serves to locate a particular spot in the image, which in this case is contrasted against a reference spot chosen in an earlier step. The instrument automatically subtracts the current and reference spot temperatures to give a ΔT indication, in this particular case 84.7 °F.

A thermal image of a three-phase circuit breaker shows a much more even distribution of temperature:




The hottest objects in this image are the three load screw terminals, appearing as white/red against a blue/green background. Note the range of the temperature scale on the right of the image: this particular image only spans a temperature range of 61.3 °F to 68.6 °F. This narrow temperature range tells us the differences in temperature shown by colors in this image are nothing to worry about.

Here, the instrument provides a single-point temperature measurement of 68.4 °F at the cursor (“Spot”) location rather than a differential temperature measurement between two points.

21.6 Temperature sensor accessories


One of the most important accessories for any temperature-sensing element is a pressure-tight sheath known as a *thermowell*. This may be thought of as a thermally conductive protrusion into a process vessel or pipe²⁵ allowing a temperature-sensitive instrument to detect process temperature without opening a hole in the vessel or pipe. Thermowells are critically important for installations where the temperature element (RTD, thermocouple, thermometer, etc.) must be replaceable without de-pressurizing the process.

Thermowells may be made out of any material that is thermally conductive, pressure-tight, and not chemically reactive with the process. Most thermowells are formed out of either metal (stainless steel or other alloy) or ceramic materials. A simple diagram showing a thermowell in use with a temperature gauge is shown here:



²⁵In general, it is better to install a thermowell in a pipe rather than in a vessel because the greater fluid turbulence of flow in a pipe expedites heat transfer by convection as well as helps to clean solid fouling off of the thermowell's surface.

If the temperature gauge is removed for maintenance or replacement, the thermowell maintains pressure integrity of the pipe (no process fluid leaking out, and no air leaking in):



Photographs of a real (stainless steel) thermowell are shown here, the left-hand photo showing the entire length of the thermowell, and the right-hand photo showing the end where the temperature-sensing device is inserted:



A photo of a complete RTD assembly (connection head, RTD, and thermowell) appears in the next photograph:



Another photo shows an RTD installed in a thermowell on the side of a commercial freezer, using a Rosemount model 3044C temperature transmitter to output a 4-20 mA signal to an operator display:



As useful as thermowells are, they are not without their caveats. All thermowells, no matter how well they may be installed, increase the first-order time lag of the temperature sensor by virtue of their mass and specific heat value. It should be intuitively obvious that a few pounds of metal will not heat up and cool down as fast as a few ounces' worth of RTD or thermocouple, and therefore the addition of a thermowell to the sensing element will decrease the responsiveness of any temperature-sensing element. What is not so obvious is that such time lags, if severe enough, may compromise the stability of feedback control. A control system receiving a “delayed” temperature measurement will not see the live temperature of the process in real time due to this lag.

A potential problem with thermowells is incorrect installation of the temperature-sensing element. The element *must* be inserted with full contact at the bottom of the thermowell's blind hole. If any air gap is allowed to exist between the end of the temperature element and the bottom of the thermowell's hole, this will add a *second* time lag to the measurement system²⁶. Some thermowells include a spring clip in the bottom of the blind hole to help maintain constant metal-to-metal contact between the sensing element and the thermowell wall.

²⁶The air gap acts as a thermal *resistance* while the mass of the element itself acts as a thermal *capacitance*. Thus, the inclusion of an air gap forms a thermal “RC time constant” delay network secondary to the thermal delay incurred by the thermowell. This adds another “order” of lag to the system, not just an increase in its thermal time constant. Generally speaking, multiple orders of lag are detrimental to process control because they increase phase shift in a feedback loop and may lead to oscillation.

21.7 Process/instrument suitability

The primary consideration for selecting a proper temperature sensing element for any application is the expected temperature range. Mechanical (bi-metal) and filled-system temperature sensors are limited to relatively low process temperatures, and cannot relay signals very far from the point of measurement.

Thermocouples are by far the most rugged and wide-ranging of the contact-type temperature sensors. Accuracies vary with thermocouple type and installation quality.

RTDs are more fragile than thermocouples, but they require no reference compensation and are inherently more linear.

Optical sensors lack the ability to measure temperature of fluids inside vessels unless a transparent window is provided in the vessel for light emissions to reach the sensor. Otherwise, the best an optical sensor can do is report the skin temperature of a vessel. For monitoring surface temperatures of solid objects, especially objects that would be impractical or even dangerous to contact (e.g. electrical insulators on high-voltage power lines), optical sensors are the only appropriate solution.

Chemical reactivity is a concern for contact-type sensors. If the sensing element is held inside a thermowell, that thermowell must be selected for minimum reaction with the process fluid(s). Bare thermocouples are particularly vulnerable to chemical reactions given the nature of most thermocouple metals (iron, nickel, copper, etc.), and must be carefully chosen for the particular process chemistry to avoid reliability problems later.

21.8 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Kirchhoff's Voltage Law:** the algebraic sum of all voltages in a loop is equal to zero. Relevant to thermocouple circuit calculations, where the reference junction voltage always opposes the measurement junction voltage. Also relevant to RTD circuits when determining the amount of voltage sensed by the instrument compared to the amount of voltage dropped by the RTD resistance element.
- **Ideal Gas Law:** $PV = nRT$, describing the relationship between gas pressure, chamber volume, gas quantity (in moles), and gas temperature. Relevant to Class III filled-bulb temperature sensors, where the increased pressure of an enclosed gas corresponds to the temperature of that gas.
- **Self-balancing opamp circuits:** all self-balancing operational amplifier circuits work on the principle of negative feedback maintaining a nearly zero differential input voltage to the opamp. Making the "simplifying assumption" that the opamp's differential input voltage is exactly zero assists in circuit analysis, as does the assumption that the input terminals draw negligible current.
- **Stefan-Boltzmann Law:** $\frac{dQ}{dt} = e\sigma AT^4$, that all objects warmer than absolute zero radiate thermal energy (photons). Relevant to non-contact pyrometry, where the intensity of the received radiation is proportional to the fourth power of the object's absolute temperature.
- **Time constant:** (τ), defined as the amount of time it takes a system to change 63.2% of the way from where it began to where it will eventually stabilize. The system will be within 1% of its final value after 5 time constants' worth of time has passed (5τ). Relevant to temperature lags caused by sensor mass and thermowells.
- **Inverse square law:** the strength of a field radiating away from a point-source diminishes proportionately to the square of the distance from the source. Relevant to determining the amount of radiant energy intercepted by a sensor when monitoring a point-source of heat.

References

Baker, Bonnie C., “Precision Temperature-Sensing With RTD Circuits”, application note AN687, Microchip Technology Incorporated, 2003.

Beckerath, Alexander von; Eberlein, Anselm; Julien, Hermann; Kersten, Peter; and Kreutzer, Jochem, *WIKA-Handbook, Pressure and Temperature Measurement*, WIKA Alexander Wiegand GmbH & Co., Klingenberg, Germany, 1995.

Darling, Charles Robert, *Pyrometry – A Practical Treatise on the Measurement of High Temperatures*, E. & F.N. Spon, Ltd, London, 1911.

Fribance, Austin E., *Industrial Instrumentation Fundamentals*, McGraw-Hill Book Company, New York, NY, 1962.

Irwin, J. David, *The Industrial Electronics Handbook*, CRC Press, Boca Raton, FL, 1997.

Kallen, Howard P., *Handbook of Instrumentation and Controls*, McGraw-Hill Book Company, Inc., New York, NY, 1961.

Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.

“Model 444 Alphaline Temperature Transmitters”, Document 00809-0100-4263, Revision AA, Rosemount, Inc., 1998

“Radiamatic Detectors and Accessories”, Specification document 23-75-03-03, Honeywell, Inc., Fort Washington, PA, 1992.

“Temperature - Electromotive Force (EMF) Tables for Standardized Thermocouples”, Pyromation, Inc.

“Temperature Measurement – Thermocouples”, ISA-MC96.1-1982, Instrument Society of America, Research Triangle Park, NC, 1982.

Chapter 22

Continuous fluid flow measurement

The measurement of fluid flow is arguably the single most complex type of process variable measurement in all of industrial instrumentation¹. Not only is there a bewildering array of technologies one might use to measure fluid flow – each one with its own limitations and idiosyncrasies – but the very nature of the variable itself lacks a singular definition. “Flow” may refer to volumetric flow (the number of fluid *volumes* passing by per unit time), mass flow (the number of fluid mass units passing by per unit time), or even *standardized* volumetric flow (the number of gas volumes flowing, supposing different pressure and temperature values than what the actual process line operates at). Flowmeters configured to work with gas or vapor flows often are unusable on liquid flows. The dynamic properties of the fluids themselves change with flow rates. Most flow measurement technologies cannot achieve respectable measurement linearity from the maximum rated flow all the way to zero flow, no matter how well matched they might be to the process application.

Furthermore, the performance of most flowmeter technologies critically depends on proper installation. One cannot simply hang a flowmeter at any location in a piping system and expect it to function as designed. This is a constant source of friction between piping (mechanical) engineers and instrumentation (controls) engineers on large industrial projects. What might be considered excellent piping layout from the perspective of process equipment function and economy is often poor (at best) for good flow measurement, and vice-versa. In many cases the flowmeter equipment gets installed improperly and the instrument technicians have to deal with the resulting measurement problems during process unit start-up.

Even after a flowmeter has been properly selected for the process application and properly installed in the piping, problems may arise due to changes in process fluid properties (density, viscosity, conductivity), or the presence of impurities in the process fluid. Flowmeters are also subject to far more “wear and tear” than most other primary sensing elements, given the fact that a flowmeter’s sensing element(s) must lie directly in the path of potentially abrasive fluid streams.


Given all these complications, it is imperative for instrumentation professionals to understand the complexities of flow measurement. What matters most is that you thoroughly understand the *physical principles* upon which each flowmeter depends. If the “first principles” of each technology are understood, the appropriate applications and potential problems become much easier

¹Analytical (chemical composition) measurement is undeniably more complex and diverse than flow measurement, but analytical measurement covers a great deal of specific measurement types. As a *single* process variable, flow measurement is probably the most complex.

to recognize.

22.1 Pressure-based flowmeters


All masses require force to accelerate (we can also think of this in terms of the mass generating a reaction force as a result of being accelerated). This is quantitatively expressed by Newton's Second Law of Motion:



Newton's Second Law formula

$$F = ma$$

All fluids possess mass, and therefore require force to accelerate just like solid masses. If we consider a quantity of fluid confined inside a pipe², with that fluid quantity having a mass equal to its volume multiplied by its mass density ($m = \rho V$, where ρ is the fluid's mass per unit volume), the force required to accelerate that fluid "plug" would be calculated just the same as for a solid mass:



Newton's Second Law formula

$$F = ma \quad F = \rho Va$$

²Sometimes referred to as a *plug* of fluid.

Since this accelerating force is applied on the cross-sectional area of the fluid plug, we may express it as a *pressure*, the definition of pressure being force per unit area:

$$F = \rho V a$$

$$\frac{F}{A} = \rho \frac{V}{A} a$$


$$P = \rho \frac{V}{A} a$$

Since the rules of algebra required we divide *both* sides of the force equation by area, it left us with a fraction of volume over area ($\frac{V}{A}$) on the right-hand side of the equation. This fraction has a physical meaning, since we know the volume of a cylinder divided by the area of its circular face is simply the length of that cylinder:

$$P = \rho \frac{V}{A} a$$

$$P = \rho l a$$

When we apply this to the illustration of the fluid mass, it makes sense: the pressure described by the equation is actually a *differential*³ pressure drop from one side of the fluid mass to the other, with the length variable (l) describing the spacing between the differential pressure ports:




This tells us we can accelerate a “plug” of fluid by applying a difference of pressure across its length. The amount of pressure we apply will be in direct proportion to the density of the fluid and its rate of acceleration. Conversely, we may measure a fluid’s rate of acceleration by measuring the pressure developed across a distance over which it accelerates.

We may easily force a fluid to accelerate by altering its natural flow path. The difference of pressure generated by this acceleration will indirectly indicate the rate of acceleration. Since the acceleration we see from a change in flow path is a direct function of how fast the fluid was originally moving, the acceleration (and therefore the pressure drop) indirectly indicates fluid flow rate.


³What really matters in Newton’s Second Law equation is the *resultant* force causing the acceleration. This is the vector sum of all forces acting on the mass. Likewise, what really matters in this scenario is the *resultant* pressure acting on the fluid plug, and this resultant pressure is the difference of pressure between one face of the plug and the other, since those two pressures impart two forces on the fluid mass in direct opposition to each other.

A very common way to cause linear acceleration in a moving fluid is to pass the fluid through a constriction in the pipe, thereby increasing its velocity (remember that the definition of acceleration is a change in velocity). The following illustrations show several devices used to linearly accelerate moving fluids when placed in pipes, with differential pressure transmitters connected to measure the pressure drop resulting from this acceleration:




Another way we may accelerate a fluid is to force it to turn a corner through a pipe fitting called an *elbow*. This will generate radial acceleration, causing a pressure difference between the outside and inside of the elbow which may be measured by a differential pressure transmitter:

Pipe elbow



The pressure tap located on the outside of the elbow's turn registers a greater pressure than the tap located on the inside of the elbow's turn, due to the inertial force of the fluid's mass being "flung" to the outside of the turn as it rounds the corner.

Yet another way to cause a change in fluid velocity is to force it to *decelerate* by bringing a portion of it to a full stop. The pressure generated by this deceleration (called the *stagnation pressure*) tells us how fast it was originally flowing. A few devices working on this principle are shown here:




The following subsections in this flow measurement chapter explore different primary sensing elements (PSE's) used to generate differential pressure in a moving fluid stream. Despite their very different designs, they all operate on the same fundamental principle: causing a fluid to accelerate or decelerate by forcing a change in its flow path, and thus generating a measurable pressure difference. The following subsection will introduce a device called a *venturi tube* used to measure fluid flow

rates, and derive mathematical relationships between fluid pressure and flow rate starting from basic physical conservation laws.

22.1.1 Venturi tubes and basic principles

The standard “textbook example” flow element used to create a pressure change by accelerating a fluid stream is the *venturi tube*: a pipe purposefully narrowed to create a region of low pressure. As shown previously, venturi tubes are not the only structure capable of producing a flow-dependent pressure drop. You should keep this in mind as we proceed to derive equations relating flow rate with pressure change: although the venturi tube is the canonical form, the exact same mathematical relationship applies to all flow elements generating a pressure drop by accelerating fluid, including orifice plates, flow nozzles, V-cones, segmental wedges, pipe elbows, pitot tubes, etc.

If the fluid going through the venturi tube is a liquid under relatively low pressure, we may vividly show the pressure at different points in the tube by means of *piezometers*⁴, which are transparent tubes allowing us to view liquid column heights. The greater the height of liquid column in the piezometer, the greater the pressure at that point in the flowstream:



As indicated by the piezometer liquid heights, pressure at the constriction (point 2) is the least, while pressures at the wide portions of the venturi tube (points 1 and 3) are the greatest. This is a counter-intuitive result, but it has a firm grounding in the physics of mass and energy conservation. If we assume no energy is added (by a pump) or lost (due to friction) as fluid travels through this pipe, then the Law of Energy Conservation describes a situation where the fluid’s energy must remain constant at all points in the pipe as it travels through. If we assume no fluid joins this flowstream from another pipe, or is lost from this pipe through any leaks, then the Law of Mass Conservation describes a situation where the fluid’s mass flow rate must remain constant at all points in the pipe as it travels through.

⁴Think of a piezometer tube as nothing more than a manometer tube: the greater the fluid pressure at the bottom of the tube, the higher the liquid will rise inside the tube.

So long as fluid density remains fairly constant⁵, fluid velocity must increase as the cross-sectional area of the pipe decreases, as described by the Law of Continuity (see section 2.11.11 beginning on page 205 for more details on this concept):

$$A_1 \bar{v}_1 = A_2 \bar{v}_2$$

Rearranging variables in this equation to place velocities in terms of areas, we get the following result:

$$\frac{\bar{v}_2}{\bar{v}_1} = \frac{A_1}{A_2}$$

This equation tells us that the ratio of fluid velocity between the narrow throat (point 2) and the wide mouth (point 1) of the pipe is the same ratio as the mouth's area to the throat's area. So, if the mouth of the pipe had an area 5 times as great as the area of the throat, then we would expect the fluid velocity in the throat to be 5 times as great as the velocity at the mouth. Simply put, the narrow throat causes the fluid to accelerate from a lower velocity to a higher velocity.

⁵This is a very sound assumption for liquids, and a fair assumption for gases when pressure changes through the venturi tube are modest.

We know from our study of energy in physics that kinetic energy is proportional to the square of a mass's velocity ($E_k = \frac{1}{2}mv^2$). If we know the fluid molecules increase velocity as they travel through the venturi tube's throat, we may safely conclude that those molecules' kinetic energies must increase as well. However, we also know that the total energy at any point in the fluid stream must remain constant, because no energy is added to or taken away from the stream in this simple fluid system⁶. Therefore, if kinetic energy increases at the throat, potential energy must correspondingly decrease to keep the total amount of energy constant at any point in the fluid.

Potential energy may be manifest as height above ground, and/or as pressure in a fluid system. Since this venturi tube is level with the ground, there cannot be a height change to account for a change in potential energy. Therefore, there *must* be a change of pressure (P) as the fluid travels through the venturi throat. The Laws of Mass and Energy Conservation invariably lead us to this conclusion: fluid pressure must decrease as it travels through the narrow throat of the venturi tube⁷.

Conservation of energy at different points in a fluid stream is neatly expressed in *Bernoulli's Equation* as a constant sum of elevation, pressure, and velocity "heads" (see section 2.11.13 beginning on page 208 for more details on this concept):

$$z_1\rho g + \frac{v_1^2\rho}{2} + P_1 = z_2\rho g + \frac{v_2^2\rho}{2} + P_2$$

Where,

z = Height of fluid (from a common reference point, usually ground level)

ρ = Mass density of fluid

g = Acceleration of gravity

v = Velocity of fluid

P = Pressure of fluid

⁶One of the simplifying assumptions we make in this derivation is that friction plays no significant role in the fluid's behavior as it moves through the venturi tube. In truth, no industrial fluid flow is totally frictionless (especially through more primitive flow elements such as orifice plates), and so our "theoretical" equations must be adjusted a bit to match real life.

⁷To see a graphical relationship between fluid acceleration and fluid pressures in a venturi tube, examine the illustration found in section B beginning on page 3163.

We will use Bernoulli's equation to develop a precise mathematical relationship between pressure and flow rate in a venturi tube. To simplify our task, we will hold to the following assumptions for our venturi tube system:

- No energy lost or gained in the venturi tube (all energy is conserved)
- No mass lost or gained in the venturi tube (all mass is conserved)
- Fluid is incompressible
- Venturi tube centerline is level (no height changes to consider)

Applying the last two assumptions to Bernoulli's equation, we see that the "elevation head" term drops out of both sides, since z , ρ , and g are equal at all points in the system:

$$\frac{v_1^2 \rho}{2} + P_1 = \frac{v_2^2 \rho}{2} + P_2$$

Now we will algebraically re-arrange this equation to show pressures at points 1 and 2 in terms of velocities at points 1 and 2:

$$\frac{v_2^2 \rho}{2} - \frac{v_1^2 \rho}{2} = P_1 - P_2$$

Factoring $\frac{\rho}{2}$ out of the velocity head terms:

$$\frac{\rho}{2}(v_2^2 - v_1^2) = P_1 - P_2$$

The Continuity equation shows us the relationship between velocities v_1 and v_2 and the areas at those points in the venturi tube, assuming constant density (ρ):

$$A_1 v_1 = A_2 v_2$$

Specifically, we need to re-arrange this equation to define v_1 in terms of v_2 so we may substitute into Bernoulli's equation:

$$v_1 = \left(\frac{A_2}{A_1} \right) v_2$$

Performing the algebraic substitution:

$$\frac{\rho}{2} \left(v_2^2 - \left[\left(\frac{A_2}{A_1} \right) v_2 \right]^2 \right) = P_1 - P_2$$

Distributing the "square" power:

$$\frac{\rho}{2} \left(v_2^2 - \left(\frac{A_2}{A_1} \right)^2 v_2^2 \right) = P_1 - P_2$$

Factoring v_2^2 out of the outer parentheses set:

$$\frac{\rho v_2^2}{2} \left(1 - \left(\frac{A_2}{A_1} \right)^2 \right) = P_1 - P_2$$

Solving for v_2 , step by step:

$$\frac{\rho v_2^2}{2} = \left(\frac{1}{1 - \left(\frac{A_2}{A_1} \right)^2} \right) (P_1 - P_2)$$

$$\rho v_2^2 = 2 \left(\frac{1}{1 - \left(\frac{A_2}{A_1} \right)^2} \right) (P_1 - P_2)$$

$$v_2^2 = 2 \left(\frac{1}{1 - \left(\frac{A_2}{A_1} \right)^2} \right) \left(\frac{P_1 - P_2}{\rho} \right)$$

$$v_2 = \sqrt{2} \frac{1}{\sqrt{1 - \left(\frac{A_2}{A_1} \right)^2}} \sqrt{\frac{P_1 - P_2}{\rho}}$$

The result shows us how to solve for fluid velocity at the venturi throat (v_2) based on a difference of pressure measured between the mouth and the throat ($P_1 - P_2$). We are only one step away from a volumetric flow equation here, and that is to convert velocity (v) into flow rate (Q). Velocity is expressed in units of length per time (feet or meters per second or minute), while volumetric flow is expressed in units of volume per time (cubic feet or cubic meters per second or minute). Simply multiplying throat velocity (v_2) by throat area (A_2) will give us the result we seek:

General flow/area/velocity relationship:

$$Q = Av$$

Equation for throat velocity:

$$v_2 = \sqrt{2} \frac{1}{\sqrt{1 - \left(\frac{A_2}{A_1} \right)^2}} \sqrt{\frac{P_1 - P_2}{\rho}}$$

Multiplying both sides of the equation by throat area:

$$A_2 v_2 = \sqrt{2} A_2 \frac{1}{\sqrt{1 - \left(\frac{A_2}{A_1} \right)^2}} \sqrt{\frac{P_1 - P_2}{\rho}}$$

Now we have an equation solving for volumetric flow (Q) in terms of pressures and areas:

$$Q = \sqrt{2}A_2 \frac{1}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} \sqrt{\frac{P_1 - P_2}{\rho}}$$

Please note how many constants we have in this equation. For any given venturi tube, the mouth and throat areas (A_1 and A_2) will be fixed. This means nearly half the variables found within this rather long equation are actually constant for any given venturi tube, and therefore do not change with pressure, density, or flow rate. Knowing this, we may re-write the equation as a simple proportionality:

$$Q \propto \sqrt{\frac{P_1 - P_2}{\rho}}$$

To make this a more precise mathematical statement, we may insert a *constant of proportionality* (k) and once more have a true equation to work with:

$$Q = k \sqrt{\frac{P_1 - P_2}{\rho}}$$

22.1.2 Volumetric flow calculations

As we saw in the previous subsection, we may derive a relatively simple equation for predicting flow through a fluid-accelerating element given the pressure drop generated by that element and the density of the fluid flowing through it:

$$Q = k \sqrt{\frac{P_1 - P_2}{\rho}}$$

This equation is a simplified version of the one derived from the physical construction of a venturi tube:

$$Q = \sqrt{2} A_2 \frac{1}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} \sqrt{\frac{P_1 - P_2}{\rho}}$$

As you can see, the constant of proportionality (k) shown in the simpler equation is nothing more than a condensation of the first half of the longer equation: k represents the geometry of the venturi tube. If we define k by the mouth and throat areas (A_1 , A_2) of any particular venturi tube, we must be very careful to express the pressures and densities in compatible units of measurement. For example, with k strictly defined by flow element geometry (tube areas measured in square *feet*), the calculated flow rate (Q) must be in units of cubic *feet* per second, the pressure values P_1 and P_2 must be in units of pounds per square *foot*, and mass density must be in units of *slugs* per cubic *foot*. We cannot arbitrarily choose different units of measurement for these variables, because the units must “agree” with one another. If we wish to use more convenient units of measurement such as inches of water column for pressure and specific gravity (unitless) for density, the original (longer) equation simply will not work.

However, if we happen to know the differential pressure produced by any particular flow element tube with any particular fluid density at a specified flow rate (real-life conditions), we may *calculate* a value for k in the short equation that makes all those measurements “agree” with one another. In other words, we may use the constant of proportionality (k) as a *unit-of-measurement correction factor* as well as a definition of element geometry. This is a useful property of all proportionalities: simply insert values (expressed in any unit of measurement) determined by physical experiment and solve for the proportionality constant’s value to satisfy the expression as an equation. If we do this, the value we arrive at for k will automatically compensate for whatever units of measurement we arbitrarily choose for pressure and density.

For example, if we know a particular orifice plate develops 45 inches of water column differential pressure at a flow rate of 180 gallons per minute of water (specific gravity = 1), we may insert these values into the equation and solve for k :

$$Q = k\sqrt{\frac{P_1 - P_2}{\rho}}$$

$$180 = k\sqrt{\frac{45}{1}}$$

$$k = \frac{180}{\sqrt{\frac{45}{1}}} = 26.83$$

Now we possess a value for k (26.83) that yields a flow rate in units of “gallons per minute” given differential pressure in units of “inches of water column” and density expressed as a specific gravity for this particular orifice plate. From the known fact of all accelerating flow elements’ behavior (flow rate proportional to the square root of pressure divided by density) and from a set of values experimentally determined for this particular orifice plate, we now have an equation useful for calculating flow rate given any set of pressure and density values we may happen to encounter with this particular orifice plate:

$$\left[\frac{\text{gal}}{\text{min}}\right] = 26.83\sqrt{\frac{[\text{” W.C.}]}{\text{Specific gravity}}}$$

This k value lets us predict flow for any given pressure difference – and vice-versa – for this particular orifice plate. For example, if we wished to know the water flow rate corresponding to a pressure difference of 60 inches water column, we could use this equation to calculate a flow rate of 207.8 gallons per minute:

$$Q = 26.83\sqrt{\frac{60}{1}}$$

$$Q = 207.8 \text{ GPM}$$

As another example, a measured differential pressure of 110 inches water column across this orifice plate generated by a flow of gasoline (specific gravity = 0.657) would correspond to a gasoline flow rate of 347 gallons per minute:

$$Q = 26.83\sqrt{\frac{110}{0.657}}$$

$$Q = 347 \text{ GPM}$$

Suppose, though, we wished to have an equation for calculating the flow rate through this same orifice plate given pressure and density data in different units (say, kPa instead of inches water column, and kilograms per cubic meter instead of specific gravity). In order to do this, we would need to re-calculate the constant of proportionality (k) to accommodate those new units of measurement. To do this, all we would need is a single set of experimental data for the orifice plate relating flow in GPM, pressure in kPa, and density in kg/m^3 .

Applying this to our original data where a water flow rate of 180 GPM resulted in a pressure drop of 45 inches water column, we could convert the pressure drop of 45 "W.C. into 11.21 kPa and express the density as $1000 \text{ kg}/\text{m}^3$ to solve for a new value of k :

$$Q = k \sqrt{\frac{P_1 - P_2}{\rho}}$$

$$180 = k \sqrt{\frac{11.21}{1000}}$$

$$k = \frac{180}{\sqrt{\frac{11.21}{1000}}} = 1700$$

Nothing about the orifice plate's geometry has changed from before, only the units of measurement we have chosen to work with. Now we possess a value for k (1700) for the same orifice plate yielding a flow rate in units of "gallons per minute" given differential pressure in units of "kilopascals" and density in units of "kilograms per cubic meter."

$$\left[\frac{\text{gal}}{\text{min}} \right] = 1700 \sqrt{\frac{[\text{kPa}]}{\text{kg}/\text{m}^3}}$$

If we were to be given a pressure drop in kPa and a fluid density in kg/m^3 for this orifice plate, we could calculate the corresponding flow rate (in GPM) with our new value of k (1700) just as easily as we could with the old value of k (26.83) given pressure in "W.C. and specific gravity.

22.1.3 Mass flow calculations

Measurements of *mass* flow are preferred over measurements of *volumetric* flow in process applications where mass balance (monitoring the rates of mass entry and exit for a process) is important. Whereas volumetric flow measurements express the fluid flow rate in such terms as *gallons per minute* or *cubic meters per second*, mass flow measurements always express fluid flow rate in terms of actual mass units over time, such *pounds (mass) per second* or *kilograms per minute*. Applications for mass flow measurement include custody transfer (where a fluid product is bought or sold by its mass), chemical reaction processes (where the mass flow rates of reactants must be maintained in precise proportion in order for the desired chemical reactions to occur), and steam boiler control systems (where the out-flow of vaporous steam must be balanced by an equivalent in-flow of liquid water to the boiler – here, volumetric comparisons of steam and water flow would be useless because one cubic foot of steam is certainly not the same number of H₂O molecules as one cubic foot of water).

If we wish to calculate *mass* flow instead of volumetric flow, the equation does not change much. The relationship between volume (V) and mass (m) for a sample of fluid is its mass density (ρ):

$$\rho = \frac{m}{V}$$

Similarly, the relationship between a volumetric *flow rate* (Q) and a mass *flow rate* (W) is also the fluid's mass density (ρ):

$$\rho = \frac{W}{Q}$$

Solving for W in this equation leads us to a product of volumetric flow rate and mass density:

$$W = \rho Q$$

A quick dimensional analysis check using common metric units confirms this fact. A mass flow rate in kilograms per second will be obtained by multiplying a mass density in kilograms per cubic meter by a volumetric flow rate in cubic meters per second:

$$\left[\frac{\text{kg}}{\text{s}} \right] = \left[\frac{\text{kg}}{\text{m}^3} \right] \left[\frac{\text{m}^3}{\text{s}} \right]$$

Therefore, all we have to do to turn our general volumetric flow equation into a mass flow equation is multiply both sides by fluid density (ρ):

$$Q = k\sqrt{\frac{P_1 - P_2}{\rho}}$$

$$\rho Q = k\rho\sqrt{\frac{P_1 - P_2}{\rho}}$$

$$W = k\rho\sqrt{\frac{P_1 - P_2}{\rho}}$$

It is generally considered “inelegant” to show the same variable more than once in an equation if it is not necessary, so let’s try to consolidate the two densities (ρ) using algebra. First, we may write ρ as the product of two square-roots:

$$W = k\sqrt{\rho}\sqrt{\rho}\sqrt{\frac{P_1 - P_2}{\rho}}$$

Next, we will break up the last radical into a quotient of two separate square roots:

$$W = k\sqrt{\rho}\sqrt{\rho}\frac{\sqrt{P_1 - P_2}}{\sqrt{\rho}}$$

Now we see how one of the square-rooted ρ terms cancels out the one in the denominator of the fraction:

$$W = k\sqrt{\rho}\sqrt{P_1 - P_2}$$

It is also considered “inelegant” to have multiple radicands in an equation where one will suffice, so we will re-write our equation for esthetic improvement⁸:

$$W = k\sqrt{\rho(P_1 - P_2)}$$

As with the volumetric flow equation, all we need in order to arrive at a suitable k value for any particular flow element is a set of values taken from that real element in service, expressed in whatever units of measurement we desire.

⁸This re-write is solidly grounded in the rules of algebra. We know that $\sqrt{a}\sqrt{b} = \sqrt{ab}$, which is what allows us to do the re-write.

For example, if we had a venturi tube generating a differential pressure of 2.30 kilopascals (kPa) at a mass flow rate of 500 kilograms per minute of naphtha (a petroleum product having a density of 0.665 kilograms per liter), we could solve for the k value of this venturi tube as such:

$$W = k\sqrt{\rho(P_1 - P_2)}$$

$$500 = k\sqrt{(0.665)(2.3)}$$

$$k = \frac{500}{\sqrt{(0.665)(2.3)}}$$

$$k = 404.3$$

Now that we know a value of 404.3 for k will yield kilograms per minute of liquid flow through this venturi tube given pressure in kPa and density in kilograms per liter, we may readily predict the mass flow rate through this tube for any other pressure drop and fluid density we might happen to encounter. The value of 404.3 for k relates the disparate units of measurement for us:

$$\left[\frac{\text{kg}}{\text{min}} \right] = 404.3 \sqrt{\left[\frac{\text{kg}}{1} \right] [\text{kPa}]}$$

As with volumetric flow calculations, the calculated value for k neatly accounts for any set of measurement units we may arbitrarily choose. The key is first knowing the proportional relationship between flow rate, pressure drop, and density. Once we combine that proportionality with a specific set of data experimentally gathered from a particular flow element, we have a true equation properly relating all the variables together in our chosen units of measurement.

If we happened to measure 6.1 kPa of differential pressure across this same venturi tube as it flowed sea water (density = 1.03 kilograms per liter), we could calculate the mass flow rate quite easily using the same equation (with the k factor of 404.3):


$$W = 404.3\sqrt{(1.03)(6.1)}$$

$$W = 1013.4 \frac{\text{kg}}{\text{min}}$$

22.1.4 Square-root characterization

It should be apparent by now that the relationship between flow rate (whether it be volumetric or mass) and differential pressure for any fluid-accelerating flow element is non-linear: a doubling of flow rate will *not* result in a doubling of differential pressure. Rather, a doubling of flow rate will result in a *quadrupling* of differential pressure. Likewise, a tripling of flow rate results in *nine times* as much differential pressure developed by the fluid-accelerating flow element.

When plotted on a graph, the relationship between flow rate (Q) and differential pressure (ΔP) is quadratic, like one-half of a parabola. Differential pressure developed by a venturi, orifice plate, pitot tube, or any other acceleration-based flow element is proportional to the *square* of the flow rate:




An unfortunate consequence of this quadratic relationship is that a pressure-sensing instrument connected to such a flow element will *not* directly sense flow rate. Instead, the pressure instrument will be sensing what is essentially the square of the flow rate. The instrument may register correctly at the 0% and 100% range points if correctly calibrated for the flow element it connects to, but it will fail to register linearly in between. Any indicator, recorder, or controller connected to the pressure-sensing instrument will likewise register incorrectly at any point between 0% and 100% of range, because the pressure signal is not a direct representation of flow rate.

In order that we may have indicators, recorders, and controllers that actually do register linearly with flow rate, we must mathematically “condition” or “characterize” the pressure signal sensed by the differential pressure instrument. Since the mathematical function inherent to the flow element is quadratic (square), the proper conditioning for the signal must be the inverse of that: *square root*. Just as taking the square-root of the square of a number yields the original number⁹, taking the square-root of the differential pressure signal – which is itself a function of flow squared – yields a signal directly representing flow.

⁹For positive numbers only!


The traditional means of implementing the necessary signal characterization was to install a “square root” function relay between the transmitter and the flow indicator, as shown in the following diagram:



The modern solution to this problem is to incorporate square-root signal characterization either inside the transmitter or inside the receiving instrument (e.g. indicator, recorder, or controller). Either way, the square-root function must be implemented *somewhere*¹⁰ in the loop in order that flow may be accurately measured throughout the operating range.

¹⁰With so many modern instruments being capable of digitally implementing this square-root function, one must be careful to ensure it is only done *once* in the loop. I have personally witnessed flow-measurement installations where both the pressure transmitter and the indicating device were configured for square-root characterization. This essentially performed a *fourth* root characterization on the signal, which is just as bad as no characterization at all! Like anything else technical, the key to successful implementation is a correct understanding of how the system is supposed to work. Simply memorizing that “the instrument must be set up with square-root to measure flow” and blindly applying that mantra is a recipe for failure.

In the days of pneumatic instrumentation, this square-root function was performed in a separate device called a *square root extractor*. The Foxboro model 557 (left) and Moore Products model 65 (right) pneumatic square root extractors are classic examples of this technology¹¹:



Pneumatic square root extraction relays approximated the square-root function by means of triangulated force or motion. In essence, they were *trigonometric* function relays, not square-root relays per se. However, for small angular motions, certain trigonometric functions were close enough to a square-root function that the relays were able to serve their purpose in characterizing the output signal of a pressure sensor to yield a signal representing flow rate.


¹¹Despite the impressive craftsmanship and engineering that went into the design of pneumatic square root extractors, their obsolescence is mourned by no one. These devices were notoriously difficult to set up and calibrate accurately, especially as they aged.

The following table shows the ideal response of a pneumatic square root relay:

Input signal	Input %	Output %	Output signal
3 PSI	0%	0%	3 PSI
4 PSI	8.33%	28.87%	6.464 PSI
5 PSI	16.67%	40.82%	7.899 PSI
6 PSI	25%	50%	9 PSI
7 PSI	33.33%	57.74%	9.928 PSI
8 PSI	41.67%	64.55%	10.75 PSI
9 PSI	50%	70.71%	11.49 PSI
10 PSI	58.33%	76.38%	12.17 PSI
11 PSI	66.67%	81.65%	12.80 PSI
12 PSI	75%	86.60%	13.39 PSI
13 PSI	83.33%	91.29%	13.95 PSI
14 PSI	91.67%	95.74%	14.49 PSI
15 PSI	100%	100%	15 PSI


As you can see from the table, the square-root relationship is most evident in comparing the input and output *percentage* values. For example, at an input signal pressure of 6 PSI (25%), the output signal percentage will be the square root of 25%, which is 50% ($0.5 = \sqrt{0.25}$) or 9 PSI as a pneumatic signal. At an input signal pressure of 10 PSI (58.33%), the output signal percentage will be 76.38%, because $0.7638 = \sqrt{0.5833}$, yielding an output signal pressure of 12.17 PSI.

When graphed, the function of a square-root extractor is precisely opposite (inverted) of the quadratic function of a flow-sensing element such as an orifice plate, venturi, or pitot tube:



When cascaded – the square-root function placed immediately after the flow element’s “square” function – the result is an output signal that tracks linearly with flow rate (Q). An instrument connected to the square root relay’s signal will therefore register flow rate as it should.

Although analog electronic square-root relays have been built and used in industry for characterizing the output of 4-20 mA electronic transmitters, a far more common implementation of electronic square-root characterization occurs in DP transmitters designed with the square-root function built in. This way, no external relay device is necessary to characterize the DP transmitter's signal into a flow rate signal:




Using a characterized DP transmitter, any 4-20 mA sensing instrument connected to the transmitter's output wires will directly interpret the signal as flow rate rather than as pressure. A calibration table for such a DP transmitter (with an input range of 0 to 150 inches water column) is shown here:

ΔP	Input %	Output % = $\sqrt{\text{Input \%}}$	Output signal
0 "W.C.	0%	0%	4 mA
37.5 "W.C.	25%	50%	12 mA
75 "W.C.	50%	70.71%	15.31 mA
112.5 "W.C.	75%	86.60%	17.86 mA
150 "W.C.	100%	100%	20 mA

Once again, we see how the square-root relationship is most evident in comparing the input and output *percentages*. Note how the percentages in this table precisely match the percentages in the

pneumatic relay table: 0% input gives 0% output; 25% input gives 50% output, 50% input gives 70.71% output, etc.

An ingenious solution to the problem of square-root characterization, commonly seen in pneumatic flow-measurement systems where the DP transmitter lacks square-root characterization, is to use an indicating device with a square-root indicating scale. For example, the following photograph shows a 3-15 PSI “receiver gauge” designed to directly sense the output of a pneumatic DP transmitter:



Note how the gauge mechanism responds directly and linearly to a 3-15 PSI input signal range (note the “3 PSI” and “15 PSI” labels in small print at the extremes of the scale, and the linearly-spaced marks around the outside of the scale arc representing 1 PSI each), but how the flow markings (0 through 10 on the inside of the scale arc) are spaced in a non-linear fashion.


An electronic variation on this theme is to draw a square-root scale on the face of a meter movement driven by the 4-20 mA output signal of an electronic DP transmitter:



As with the square-root receiver gauge, the meter movement's response to the transmitter signal is linear. Note the linear scale (drawn in black text labeled "LINEAR") on the bottom and the corresponding square-root scale (in green text labeled "FLOW") on the top. This makes it possible for a human operator to read the scale in terms of (characterized) flow units. Instead of using complicated mechanisms or circuitry to characterize the transmitter's signal, a non-linear scale "performs the math" necessary to interpret flow.

A major disadvantage to the use of these non-linear indicator scales is that the transmitter signal itself remains un-characterized. Any other instrument receiving this un-characterized signal will either require its own square-root characterization or simply not interpret the signal in terms of flow at all. An un-characterized flow signal input to a process controller can cause loop instability at high flow rates, where small changes in actual flow rate result in huge changes in differential pressure sensed by the transmitter. A fair number of flow control loops operating without characterization have been installed in industrial applications (usually with square-root scales drawn on the face of the indicators, and square-root paper installed in chart recorders), but these loops are notorious for achieving good flow control at only one setpoint value. If the operator raises or lowers the setpoint value, the "gain" of the control loop changes thanks to the nonlinearities of the flow element, resulting in either under-responsive or over-responsive action from the controller.

Despite the limited practicality of non-linear indicating scales, they hold significant value as teaching tools. Closely examine the scales of both the receiver gauge and the 4-20 mA indicating meter, comparing the linear and square-root values at common points on each scale. A couple of examples are highlighted on the electric meter's scale:



A few correlations between the linear and square-root scales of either the pneumatic receiver gauge or the electric indicating meter verify the fact that the square-root function is *encoded* in the spacing of the numbers on each instrument's non-linear scale.

Another valuable lesson we may take from the faces of these indicating instruments is how uncertain the flow measurement becomes at the low end of the scale. Note how for each indicating instrument (both the receiver gauge and the meter movement), the square-root scale is “compressed” at the low end, to the point where it becomes impossible to interpret fine increments of flow at that end of the scale. At the high end of each scale, it's a different situation entirely: the numbers are spaced so far apart that it's easy to read fine distinctions in flow values (e.g. 94% flow versus 95% flow). However, the scale is so crowded at the low end that it's really impossible to clearly distinguish two different flow values such as 4% from 5%.


This “crowding” is not just an artifact of a visual scale; it is a reflection of a fundamental limitation in measurement certainty with this type of flow measurement. The amount of differential pressure separating different low-range values of flow for a flow element is so little, even small amounts of pressure-measurement error equate to large amounts of flow-measurement error. In other words, it becomes more and more difficult to precisely interpret flow rate as the flow rate decreases toward the low end of the scale. The “crowding” we see on these indicator's square-root scales is a visual reflection of this fundamental problem: even a small error in interpreting the pointer's position at the low end of the scale can yield major errors in flow interpretation.

A technical term used to quantify this problem is *turndown*. “Turndown” refers to the ratio of high-range measurement to low-range measurement possible for an instrument while maintaining reasonable accuracy. For pressure-based flowmeters, which must deal with the non-linearities of Bernoulli's Equation, the practical turndown is often no more than 3 to 1 (3:1). This means a flowmeter ranged for 0 to 300 GPM might only read with reasonable accuracy down to a flow

of 100 GPM. Below that, the accuracy becomes so poor that the measurement is almost useless. Advances in DP transmitter technology have pushed this ratio further, perhaps as far as 10:1 for certain installations. However, the fundamental problem is not transmitter resolution, but rather the nonlinearity of the flow element itself. This means *any* source of pressure-measurement error – whether originating in the transmitter’s pressure sensor or not – compromises our ability to accurately measure flow at low rates. Even with a *perfectly* calibrated transmitter, errors resulting from wear of the flow element (e.g. a dulled edge on an orifice plate) or from uneven liquid columns in the impulse tubes connecting the transmitter to the element, will cause large flow-measurement errors at the low end of the instrument’s range where the flow element produces only small differential pressures. Everyone involved with the technical details of flow measurement needs to understand this fact: the fundamental problem of limited turndown is grounded in the physics of turbulent flow and potential/kinetic energy exchange for these flow elements. Technological improvements will help, but they cannot overcome the limitations imposed by physics. If better turndown is required for a particular flow-measurement application, an entirely different flowmeter technology should be considered.

22.1.5 Orifice plates


Of all the pressure-based flow elements in existence, the most common is the *orifice plate*. This is simply a metal plate with a hole in the middle for fluid to flow through. Orifice plates are typically sandwiched between two flanges of a pipe joint, allowing for easy installation and removal:



The point where the fluid flow profile constricts to a minimum cross-sectional area after flowing through the orifice is called the *vena contracta*, and it is the area of minimum fluid pressure. The vena contracta corresponds to the narrow throat of a venturi tube. The precise location of the vena contracta for an orifice plate installation will vary with flow rate, and also with the *beta ratio* (β) of the orifice plate, defined as the ratio of bore diameter (d) to inside pipe diameter (D):


$$\beta = \frac{d}{D}$$

The simplest design of orifice plate is the *square-edged, concentric* orifice. This type of orifice plate is manufactured by machining a precise, straight hole in the middle of a thin metal plate. Looking at a side view of a square-edged concentric orifice plate reveals sharp edges (90° corners) at the hole:




Square-edged orifice plates may be installed in either direction, since the orifice plate “appears” exactly the same from either direction of fluid approach. In fact, this allows square-edged orifice plates to be used for measuring bidirectional flow rates (where the fluid flow direction reverses itself from time to time). A text label printed on the “paddle” of any orifice plate customarily identifies the upstream side of that plate, but in the case of the square-edged orifice plate it does not matter.

The purpose of having a square edge on the hole in an orifice plate is to minimize contact with the fast-moving moving fluid stream going through the hole. Ideally, this edge will be knife-sharp. If the orifice plate is relatively thick (1/8 or an inch or more), it may be necessary to bevel the downstream side of the hole to further minimize contact with the fluid stream:




Looking at the side-view of this orifice plate, the intended direction of flow is left-to-right, with the sharp edge facing the incoming fluid stream and the bevel providing a non-contact outlet for the fluid. Beveled orifice plates are obviously uni-directional, and *must* be installed with the paddle text facing upstream.

Other square-edged orifice plates exist to address conditions where gas bubbles or solid particles may be present in liquid flows, or where liquid droplets or solid particles may be present in gas flows. The first of this type is called the *eccentric* orifice plate, where the hole is located off-center to allow the undesired portions of the fluid to pass through the orifice rather than build up on the upstream face:




For gas flows, the hole should be offset downward, so any liquid droplets or solid particles may easily pass through. For liquid flows, the hole should be offset upward to allow gas bubbles to pass through and offset downward to allow heavy solids to pass through.

The second off-center orifice plate type is called the *segmental orifice plate*, where the hole is not circular but rather just a segment of a concentric circle:



As with the eccentric orifice plate design, the segmental hole should be offset downward in gas flow applications and either upward or downward in liquid flow applications depending on the type of undesired material(s) in the flowstream.

An alternative to offsetting or re-shaping the bore hole of an orifice plate is to simply drill a small hole near the edge of the plate, flush with the inside diameter of the pipe, allowing undesired substances to pass through the plate rather than collect on the upstream side. If such a hole is oriented upward to pass vapor bubbles, it is called a *vent hole*. If the hole is oriented downward to pass liquid droplets or solids, it is called a *drain hole*. Vent and drain holes are useful when the concentration of these undesirable substances is not significant enough to warrant either an eccentric or segmental orifice:



The addition of a vent or drain hole should have negligible impact on the performance of an orifice plate due to its small size relative to the main bore. If the quantity of undesirable material in the flowstream (bubbles, droplets, or solids) is excessive, an eccentric or segmental orifice plate might be a better choice¹².


¹²L.K. Spink, in his book *Principles and Practice of Flow Meter Engineering*, notes that drain holes intended to pass solid objects may be useless in small pipe sizes, where the hole is so small it will probably become plugged with solid debris and cease to provide benefit. In such installations he recommends re-orienting the pipe vertically instead

Some orifice plates employ non-square-edged holes for the purpose of improving performance at low Reynolds number¹³ values, where the effects of fluid viscosity are more apparent. These orifice plate types employ rounded- or conical-entrance holes in an effort to minimize the effects of fluid viscosity. Experiments have shown that decreased Reynolds number causes the flowstream to not contract as much when traveling through an orifice, thus limiting fluid acceleration and decreasing the amount of differential pressure produced by the orifice plate. However, experiments have also shown that decreased Reynolds number in a venturi-type flow element causes an *increase* in differential pressure due to the effects of friction against the entrance cone walls. By manufacturing an orifice plate in such a way that the hole exhibits “venturi-like” properties (i.e. a dull edge where the fast-moving fluid stream has more contact with the plate), these two effects tend to cancel each other, resulting in an orifice plate that maintains consistent accuracy at lower flow rates and/or higher viscosities than the simple square-edged orifice.


of horizontally. This allows solids to pass through the main bore of the orifice without “damming” on the upstream side of the orifice plate. I would add the suggestion to consider a different primary element entirely, such as a venturi tube. The small size of the line will limit the cost of such an element, and the performance is likely to be far better than an orifice plate anyway.

¹³To read more about the concept of Reynolds number, refer to section 2.11.10 beginning on page 201.

Two common non-square-edge orifice plate designs are the *quadrant-edge* and *conical-entrance* orifices. The quadrant-edge is shown first:




The conical-entrance orifice plate looks like a beveled square-edge orifice plate installed backwards, with flow entering the conical side and exiting the square-edged side:



Here, it is vitally important to pay attention to the paddle's text label. This is the only sure indication of which direction an orifice plate needs to be installed. One can easily imagine an instrument technician mistaking a conical-entrance orifice plate for a square-edged, beveled orifice plate and installing it backward!

Several standards exist for pressure tap locations. Ideally, the upstream pressure tap will detect fluid pressure at a point of minimum velocity, and the downstream tap will detect pressure at the vena contracta (maximum velocity). In reality, this ideal is never perfectly achieved. An overview of the most popular tap locations for orifice plates is shown in the following illustration:



Flange taps are the most popular tap location for orifice meter runs on large pipes in the United States. Flanges may be manufactured with tap holes pre-drilled and finished before the flange is even welded to the pipe, making this a very convenient pressure tap configuration. Most of the other tap configurations require drilling into the pipe after installation, which is not only labor-intensive, but may possibly weaken the pipe at the locations of the tap holes.

Vena contracta taps offer the greatest differential pressure for any given flow rate, but require precise calculations to properly locate the downstream tap position. *Radius taps* are an approximation of vena contracta taps for large pipe sizes (one-half pipe diameter downstream for the low-pressure tap location). An unfortunate characteristic of both these taps is the requirement of drilling through the pipe wall. Not only does this weaken the pipe, but the practical necessity of drilling the tap holes in the installed location rather than in a controlled manufacturing environment means there is considerable room for installation error¹⁴.

Corner taps must be used on small pipe diameters where the vena contracta is so close to the downstream face of the orifice plate that a downstream flange tap would sense pressure in the highly turbulent region (too far downstream). Corner taps obviously require special (i.e. expensive) flange fittings, which is why they tend to be used only when necessary.

Care should be taken to avoid measuring downstream pressure in the highly turbulent region following the vena contracta. This is why the *pipe tap* (also known as *full-flow tap*) standard calls for a downstream tap location eight pipe diameters away from the orifice: to give the flow stream room to stabilize for more consistent pressure readings¹⁵.

Wherever the taps are located, it is vitally important that the tap holes be completely flush with the inside wall of the pipe or flange. Even the smallest recess or burr left from drilling will cause measurement errors, which is why tap holes are best drilled in a controlled manufacturing environment rather than at the installation site where the task will likely be performed by non-experts.

¹⁴One significant source of error for customer-drilled tap holes is the interior finish of the holes. Even a small “burr” of metal left where the hole penetrates the inner surface of the pipe wall will cause substantial flow measurement errors!

¹⁵What this means is that a “pipe tap” installation is actually measuring permanent pressure loss, which also happens to scale with the square of flow rate because the primary mechanism for energy loss in turbulent flow conditions is the translation of linear velocity to angular (swirling) velocity in the form of eddies. This kinetic energy is eventually dissipated in the form of heat as the eddies eventually succumb to viscosity.

A photograph of an orifice plate used to measure the flow of natural gas to a large turbine engine is shown here, with a Rosemount model 3051 differential pressure transmitter sensing the pressure drop generated by the orifice:



Flange taps are used in this orifice installation, with the taps and pressure transmitter located above the pipe centerline to avoid picking up any liquid droplets that may pass through the pipe. The direction of gas flow in this particular installation is from left to right, making the left-hand pressure tap the “high pressure” side and the right-hand pressure tap the “low pressure” side.

As you can see by the pressure gauge’s indication in the photograph, the static line pressure of the natural gas inside the pipe is over 300 PSI. The amount of pressure drop generated by the orifice plate at full flow, however, is likely only a few PSI (100 inches water column is typical for many orifice plate installations). This is why we must use a *differential* pressure transmitter to measure the orifice plate’s pressure drop: only a DP transmitter will sense the difference in pressure across the orifice while rejecting the static (common-mode) line pressure inside the pipe.

A photograph of another orifice plate with flange taps appears here, shown on a vertical pipe. In this example, the pipe and flanges are formed of acrylic (transparent plastic):



As is customary with orifice plates mounted in vertical pipes, the direction of flow is upward (from bottom to top), making the bottom tap the “high pressure” side and the top tap the “low pressure” side. Since this is a liquid application, the transmitter is located below¹⁶ the taps in order to avoid collecting bubbles of air or other gases in the impulse lines.



This particular orifice and flow transmitter (Rosemount model 1151) is used on a process “trainer” unit at Brazosport College in Lake Jackson, Texas. The transparent flanges, pipes, and process vessels make it easier for students to visualize the fluid motion.

¹⁶One installation error seen in this photograph is a green plastic impulse tube with a bend extending above the upper flange tap. Any elevated portion of the impulse tube system will tend to collect gas bubbles over time, possibly causing measurement errors. A better installation would ensure the impulse tubes never extend above the flange tap they connect to on the liquid-bearing pipe.

For relatively low flow rates, an alternative arrangement is the *integral orifice plate*. This is where a small orifice plate directly attaches to the differential pressure-sensing element, eliminating the need for impulse lines. A photograph of an integral orifice plate and transmitter is shown here, in an application measuring the flow of purified oxygen gas through a copper pipe:





Even smaller integral orifices exist for the measurement of very low flow rates, another type of integral orifice is available. This style uses the pass-through nature of a typical differential pressure transmitter's capsule flanges to advantage. Note the paths of fluid flow, as well as the unusual orientation of the three-valve manifold, in these two example illustrations:



In both examples, the third hand valve serves the purpose of bypassing process fluid flow rather than equalizing both ports of the differential pressure transmitter. Process fluid actually flows *through* at least one chamber of the DP transmitter's body. Clearly, this kind of integral orifice plate is practical only for very low rates of flow, typically where the process flow line is no larger than the size of the pressure ports on the transmitter's body.


The task of properly sizing an orifice plate for any given application is complex enough to recommend the use of special orifice sizing computer software provided by orifice plate manufacturers. There are a number of factors to consider in orifice plate sizing, and these software packages account for all of them. Best of all, the software provided by manufacturers is often linked to data for that manufacturer's product line, helping to assure installed results in close agreement with predictions.

In the days before ubiquitous personal computers and the Internet, some orifice plate manufacturers provided customers with paper "slide rule" calculators to help them select appropriate orifice plate sizes from known process parameters. The following photographs show the front and back sides of one such slide rule:




22.1.6 Other differential producers

Other pressure-based flow elements exist as alternatives to the orifice plate. The *Pitot tube*, for example, senses pressure as the fluid stagnates (comes to a complete stop) against the open end of a forward-facing tube. A shortcoming of the classic single-tube Pitot assembly is sensitivity to fluid velocity at just one point in the pipe, so a more common form of Pitot tube seen in industry is the *averaging* Pitot tube consisting of several stagnation holes sensing velocity at multiple points across the width of the flow:




A variation on the latter theme is the *Annubar* flow element, a trade name of the Dieterich Standard corporation. An “Annubar” is an averaging pitot tube consolidating high and low pressure-sensing ports in a single probe assembly:




What appears at first glance to be a single, square-shaped tube inserted into the pipe is actually a double-ported tube with holes on both the upstream and downstream edges:



A section of Annubar tube clearly shows the porting and dual chambers, designed to bring upstream (stagnation) and downstream pressures out of the pipe to a differential pressure-sensing instrument:



A less sophisticated realization of the stagnation principle is the *target* flow sensor, consisting of a blunt “paddle” (or “drag disk”) inserted into the flowstream. The force exerted on this paddle by the moving fluid is sensed by a special transmitter mechanism, which then outputs a signal corresponding to flow rate (proportional to the square of fluid velocity, just like an orifice plate):




The following photograph shows a Foxboro model 18 target flow transmitter installed on a line conveying liquid wood pulp:




The classic venturi tube pioneered by Clemens Herschel in 1887 has been adapted in a variety of forms broadly classified as *flow tubes*. All flow tubes work on the same principle: developing a differential pressure by channeling fluid flow from a wide tube to a narrow tube. They differ from the classic venturi only in construction details, the most significant detail being a significantly shorter length than the classic venturi tube. Examples of flow tube designs include the *Dall* tube, *Lo-Loss* flow tube, *Gentile* or *Bethlehem* flow tube, and the *B.I.F. Universal Venturi*.

Another variation on the venturi theme is called a *flow nozzle*, designed to be clamped between the faces of two pipe flanges in a manner similar to an orifice plate. The goal here is to achieve simplicity of installation approximating that of an orifice plate while improving performance (less permanent pressure loss) over orifice plates:



Two more variations on the venturi theme are the *V-cone* and *Segmental wedge* flow elements. The V-cone (or “venturi cone,” a trade name of the McCrometer division of the Danaher corporation) may be thought of as a venturi tube or orifice plate in reverse¹⁷: instead of narrowing the tube’s diameter to cause fluid acceleration, fluid must flow around a cone-shaped obstruction placed in the middle of the tube. The tube’s effective area will be reduced by the presence of this cone, causing fluid to accelerate through the restriction just as it would through the throat of a classic venturi tube:




This cone is hollow, with a pressure-sensing port on the downstream side allowing for easy detection of fluid pressure near the vena contracta. Upstream pressure is sensed by another port in the pipe wall upstream of the cone. The following photograph shows a V-cone flow tube, cut away for demonstration purposes:




¹⁷If an orifice plate is a “donut,” the V-cone is a “donut hole.”

Segmental wedge elements are special pipe sections with wedge-shaped restrictions built in. These devices, albeit crude, are useful for measuring the flow rates of *slurries*¹⁸, especially when pressure is sensed by the transmitter through remote-seal diaphragms (to eliminate the possibility of impulse tube plugging):



Finally, the lowly pipe elbow may be pressed into service as a flow-measuring element, since fluid turning a corner¹⁹ in the elbow experiences radial acceleration and therefore generates a differential pressure along the axis of acceleration:



Pipe elbows should be considered for flow measurement only as a last resort. Their inaccuracies tend to be extreme, owing to the non-precise construction of most pipe elbows and the relatively

¹⁸A “slurry” is a suspension of solid particles within a liquid. *Mud* is a common example of a slurry.

¹⁹This phenomenon may be observed when watching the flow of water through a turn in a river, especially if the river is fast-moving. Water level at the far (outside) bank of the turn will be higher than the water level at the near (inside) bank of the turn, due to radial acceleration of the water and the pressure difference that acceleration generates. In fact, that difference in water height may even be used to estimate the river’s flow rate!

weak differential pressures generated²⁰.

A final point should be mentioned on the subject of differential-producing elements, and that is their energy dissipation. Orifice plates are simple and relatively inexpensive to install, but their permanent pressure loss is high compared with other primary elements such as venturi tubes. Permanent pressure loss is permanent energy loss from the flowstream, which usually represents a loss in energy invested into the process by pumps, compressors, and/or blowers. Fluid energy dissipated by an orifice plate thus (usually) translates into a requirement of greater energy input to that process²¹.

With the financial and ecological costs of energy being non-trivial in our modern world, it is important to consider energy loss as a significant factor in choosing the appropriate primary element for a pressure-based flowmeter. It might very well be that an “expensive” venturi tube saves more money in the long term than a “cheap” orifice plate, while delivering greater measurement accuracy as an added benefit.

²⁰The fact that a pipe elbow generates small differential pressure is an accuracy concern because other sources of pressure become larger by comparison. Noise generated by fluid turbulence in the elbow, for example, becomes a significant portion of the pressure sensed by the transmitter when the differential pressure is so low (i.e. the signal-to-noise ratio becomes smaller). Errors caused by differences in elbow tap elevation and different impulse line fill fluids, for example, become more significant as well.


²¹This is not always the case, as primary elements are often found on throttled process lines. In such cases where a control valve normally throttles the flow rate, any energy dissipated by the orifice plate is simply less energy that the valve would otherwise be required to dissipate. Therefore, the presence or absence of an orifice plate has no net impact on energy dissipation when used on a process flow throttled by a control valve.

22.1.7 Proper installation

Perhaps the most common way in which the flow measurement accuracy of any flowmeter becomes compromised is incorrect installation, and pressure-based flowmeters are no exception to this rule. The following list shows some of the details one must consider in installing a pressure-based flowmeter element:

- Necessary upstream and downstream straight-pipe lengths
- Beta ratio (ratio of orifice bore diameter to pipe diameter: $\beta = \frac{d}{D}$)
- Impulse tube tap locations
- Tap finish
- Transmitter location in relation to the pipe


Sharp turns in piping networks introduce large-scale turbulence²² into the flowstream. Elbows, tees, valves, fans, and pumps are some of the most common causes of large-scale turbulence in piping systems. Successive pipe elbows in different planes are some of the worst offenders in this regard. When the natural flow path of a fluid is disturbed by such piping arrangements, the velocity profile of that fluid becomes distorted; e.g. the velocity gradient from one wall boundary of the pipe to the other will not be orderly. Large eddies in the flowstream (called *swirl*) will appear. This may cause problems for pressure-based flow elements which rely on linear acceleration (change in velocity in one dimension) to measure fluid flow rate. If the flow profile is distorted enough, the acceleration detected at the element may be too great or too little, and therefore not properly represent the full fluid flowstream²³. For this reason, pressure-based flowmeters should always be located *upstream* of major disturbances such as control valves and pipe elbows where possible.



²²This is not to be confused with micro-turbulence in the fluid, which cannot be eliminated at high Reynolds number values. In fact, “fully-developed turbulent flow” is desirable for head-based meter elements such as orifice plates because it means the flow profile will be relatively flat (even velocities across the pipe’s diameter) and frictional forces (viscosity) will be negligible. The thing we are trying to avoid is *large-scale* turbulent effects such as eddies, swirl, and asymmetrical flow profiles, which compromise the ability of most flowmeters to accurately measure flow rate.


²³L.K. Spink mentions in his book *Principles and Practice of Flow Meter Engineering* that certain tests have shown flow measurement errors induced from severe disturbances as far as 60 to 100 pipe diameters upstream of the primary flow element. The April 2000 update of API standard 14.3 (for custody-transfer measurement of natural gas using orifice plates) calls for upward of *145 pipe diameters* of straight-length pipe upstream of the orifice plate!

Even disturbances located *downstream* of the flow element may affect measurement accuracy if the disturbances are severe enough and/or close enough to the flow element. Unfortunately, both upstream and downstream flow disturbances are unavoidable on all but the simplest fluid systems. This means we must devise ways to stabilize a flowstream's velocity profile near the flow element in order to achieve accurate measurements of flow rate. A very simple and effective way to stabilize a flow profile is to provide adequate lengths of straight pipe ahead of (and behind) the flow element. Given enough time, even the most chaotic flowstream will “settle down” to a symmetrical profile all on its own. The following illustration shows the effect of a pipe elbow on a flowstream, and how the velocity profile returns to a normal (symmetrical) form after traveling through a sufficient length of straight pipe:



Recommendations for minimum upstream and downstream straight-pipe lengths vary significantly with the nature of the turbulent disturbance, piping geometry, and flow element. As a general rule, elements having a smaller beta ratio (ratio of throat diameter d to pipe diameter D) are more tolerant of disturbances, with profiled flow devices (e.g. venturi tubes, flow tubes, V-cones) having the greatest tolerance²⁴. Ultimately, you should consult the flow element manufacturer's documentation for a more detailed recommendation appropriate to any specific application.

In applications where sufficient straight-run pipe lengths are impractical, another option exists for “taming” turbulence generated by piping disturbances. Devices called *flow conditioners* may be installed upstream of the flow element to help the flow profile achieve symmetry in a far shorter distance than simple straight pipe could do alone. Flow conditioners take the form of a series of tubes or vanes installed inside the pipe, parallel to the direction of flow. These tubes or vanes force the fluid molecules to travel in straighter paths, thus stabilizing the flowstream prior to entering a flow element:



²⁴Flow elements with low beta ratio values tolerate greater disturbance in the flow pattern because they accelerate the flowstream to a greater degree. This may be best visualized by a thought experiment where we imagine an orifice plate with a very large beta ratio (i.e. one where the bore size is nearly as large as the pipe diameter): such an orifice plate would hardly accelerate the fluid at all, which would mean a mis-shapen flow profile entering the bore would probably remain mis-shapen exiting it. The acceleration imparted to a flowstream by a low-beta element tends to overshadow any asymmetries in the flow profile. However, there are disadvantages to using low-beta elements, one of them being increased permanent pressure loss which may translate to increased operating costs due to energy loss.

This next photograph shows a *very poor* orifice plate installation, where the straight-run pipe requirements were completely ignored:




Not only is the orifice plate placed much too close to an elbow, the elbow happens to be on the upstream side of the orifice plate, where disturbances have the greatest effect! The saving grace of this installation is that it is not used for critical monitoring or control: it is simply a manual indication of air flow rate where accuracy is not terribly important. Nevertheless, it is sad to see how an orifice meter installation could have been so easily improved with just a simple re-location of the orifice plate along the piping length.

Poor installations such as this are surprisingly common, owing to the ignorance many piping designers have of flowmeter design and operating principles. Of all the criteria which must be balanced when designing a piping layout, optimum flowmeter location is often placed low in the order of importance (if it appears at all!). In applications where accuracy is important, flowmeter location needs to be a very high priority even if it means a more expensive, cumbersome, and/or unattractive²⁵ piping design.


²⁵Beauty is truly in the eye of the beholder. While a piping designer might see straight-run lengths of pipe in awkward locations – necessitating more pipe and/or more bends elsewhere in the system to accommodate – as wasteful and ugly, the instrument engineer sees it as a thing of beauty.

Another common source of trouble for pressure-based flowmeters is improper transmitter location. Here, the type of process fluid flow being measured dictates how the pressure-sensing instrument should be located in relation to the pipe. For gas and vapor flows, it is important that no stray liquid droplets collect in the impulse lines leading to the transmitter, lest a vertical liquid column begin to collect in those lines and generate an error-producing pressure. For liquid flows, it is important that no gas bubbles collect in the impulse lines, or else those bubbles may displace liquid from the lines and thereby cause unequal vertical liquid columns, which would (again) generate an error-producing differential pressure.

In order to let gravity do the work of preventing these problems, we must locate the transmitter *above* the pipe for gas flow applications and *below* the pipe for liquid flow applications. This illustration shows both installations for a horizontal pipe:



This next illustration shows both installations on a vertical pipe:



Condensable vapor applications (such as steam flow measurement) have traditionally been treated similarly to liquid measurement applications. Here, condensed liquid will collect in the transmitter's impulse lines so long as the impulse lines are cooler than the vapor flowing through the pipe (which is typically the case). Placing the transmitter below the pipe allows vapors to condense and fill the impulse lines with liquid (condensate), which then acts as a natural seal protecting the transmitter from exposure to hot process vapors.

In such applications it is important for the technician to pre-fill both impulse lines with condensed liquid prior to placing the flowmeter into service. "Tee" fittings with removable plugs or fill valves are provided to do this. Failure to pre-fill the impulse lines will likely result in measurement errors during initial operation, as condensed vapors will inevitably fill the impulse lines at slightly different rates and cause a difference in vertical liquid column heights within those lines.

It should be noted that some steam flow element installations, however, will work well if the impulse lines are above the pipe. If such an installation is possible, the advantage of not having to deal with pre-filling impulse lines (or waiting for steam to condense to equal levels in both lines) are significant. For more information, I recommend consulting the Rosemount whitepaper entitled "Top Mount Installation for DP Flowmeters in Steam Service" (document 00870-0200-4809 first published August 2009).

If tap holes must be drilled into the pipe (or flanges) at the process site, great care must be taken to properly drill and de-burr the holes. A pressure-sensing tap hole should be flush with the inner pipe wall, with no rough edges or burrs to create turbulence. Also, there should be no reliefs or countersinking near the hole on the inside of the pipe. Even small irregularities at the tap holes may generate surprisingly large flow-measurement errors.

22.1.8 High-accuracy flow measurement

When we derived a formula for predicting flow rate from pressure dropped by a venturi tube (or orifice), we had to make many assumptions, chief among them being a total lack of friction (i.e. no energy dissipated due to friction) within the moving fluid and perfect stream-line flow (i.e. complete lack of turbulence). Suffice it to say, the flow formulae you have seen so far in this chapter are only approximations of reality. Orifice plates are some of the worst offenders in this regard, since the fluid encounters such abrupt changes in geometry passing through the orifice. Venturi tubes are nearly ideal, since the machined contours of the tube ensure gradual changes in fluid pressure and minimize turbulence.

However, in the real world we must often do the best we can with imperfect technologies. Orifice plates, despite being less than perfect as flow-sensing elements, are convenient and economical to install in flanged pipes. Orifice plates are also the easiest type of flow element to replace in the event of damage or routine servicing. In applications such as custody transfer (also called “fiscal” measurement), where the flow of fluid represents product being bought and sold, flow measurement accuracy is paramount. It is therefore important to figure out how to coax the most accuracy from the common orifice plate in order that we may measure fluid flows both accurately and economically.

If we compare the true flow rate through a pressure-generating primary sensing element against the theoretical flow rate predicted by an idealized equation, we may notice a substantial discrepancy²⁶. Causes of this discrepancy include, but are not limited to:

- Energy losses due to turbulence and viscosity
- Energy losses due to friction against the pipe and element surfaces
- Unstable location of *vena contracta* with changes in flow
- Uneven velocity profiles caused by irregularities in the pipe
- Fluid compressibility
- Thermal expansion (or contraction) of the element and piping
- Non-ideal pressure tap location(s)
- Excessive turbulence caused by rough internal pipe surfaces

The ratio between true flow rate and theoretical flow rate for any measured amount of differential pressure is known as the *discharge coefficient* of the flow-sensing element, symbolized by the variable C . Since a value of 1 represents a theoretical ideal, the actual value of C for any real pressure-generating flow element will be less than 1:

$$C = \frac{\text{True flow}}{\text{Theoretical flow}}$$

²⁶Richard W. Miller, in his outstanding book *Flow Measurement Engineering Handbook*, states that venturi tubes may come within 1 to 3 percent of ideal, while a square-edged orifice plate may perform as poorly as only 60 percent of theoretical!

For gas and vapor flows, true flow rate deviates even more from the theoretical (ideal) flow value than liquids do, for reasons that have to do with the compressible nature of gases and vapors. A *gas expansion factor* (Y) may be calculated for any flow element by comparing its discharge coefficient for gases against its discharge coefficient for liquids. As with the discharge coefficient, values of Y for any real pressure-generating element will be less than 1:

$$Y = \frac{C_{gas}}{C_{liquid}}$$

$$Y = \frac{\left(\frac{\text{True gas flow}}{\text{Theoretical gas flow}}\right)}{\left(\frac{\text{True liquid flow}}{\text{Theoretical liquid flow}}\right)}$$

Incorporating these factors into the ideal volumetric flow equation developed on in section 22.1.1, we arrive at the following formulation:

$$Q = \sqrt{2} \frac{CY A_2}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} \sqrt{\frac{P_1 - P_2}{\rho}}$$

If we wished, we could even add another factor to account for any necessary unit conversions (N), getting rid of the constant $\sqrt{2}$ in the process:

$$Q = N \frac{CY A_2}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} \sqrt{\frac{P_1 - P_2}{\rho}}$$

Sadly, neither the discharge coefficient (C) nor the gas expansion factor (Y) will remain constant across the entire measurement range of any given flow element. These variables are subject to some change with flow rate, which further complicates the task of accurately inferring flow rate from differential pressure measurement. However, if we know the values of C and Y for typical flow conditions, we may achieve good accuracy most of the time.

Likewise, the fact that C and Y change with flow places limits on the accuracy obtainable with the “proportionality constant” formulae seen earlier. Whether we are measuring volumetric or mass flow rate, the k factor calculated at one particular flow condition will not hold constant for *all* flow conditions:

$$Q = k \sqrt{\frac{P_1 - P_2}{\rho}}$$

$$W = k \sqrt{\rho(P_1 - P_2)}$$

This means after we have calculated a value for k based on a particular flow condition, we can only trust the results of the equation for flow conditions not too different from the one we used to calculate k .

As you can see in both flow equations, the density of the fluid (ρ) is an important factor. If fluid density is relatively stable, we may treat ρ as a constant, incorporating its value into the proportionality factor (k) to make the two formulae even simpler:

$$Q = k_Q \sqrt{P_1 - P_2}$$

$$W = k_W \sqrt{P_1 - P_2}$$

However, if fluid density is subject to change over time, we will need some means to continually calculate ρ so our inferred flow measurement will remain accurate. Variable fluid density is a typical state of affairs in gas flow measurement, since all gases are compressible by definition. A simple change in static gas pressure within the pipe is all that is needed to make ρ change, which in turn affects the relationship between flow rate and differential pressure drop.

The American Gas Association (AGA) provides a formula for calculating volumetric flow of any gas using orifice plates in their #3 Report, compensating for changes in gas pressure and temperature. A variation of that formula is shown here (consistent with previous formulae in this section):

$$Q = N \frac{CY A_2}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} \sqrt{\frac{Z_s P_1 (P_1 - P_2)}{G_f Z_{f1} T}}$$

Where,

Q = Volumetric flow rate (SCFM = standard cubic feet per minute)

N = Unit conversion factor

C = Discharge coefficient (accounts for energy losses, Reynolds number corrections, pressure tap locations, etc.)

Y = Gas expansion factor

A_1 = Cross-sectional area of mouth

A_2 = Cross-sectional area of throat

Z_s = Compressibility factor of gas under standard conditions

Z_{f1} = Compressibility factor of gas under flowing conditions, upstream


G_f = Specific gravity of gas (density compared to ambient air)

T = Absolute temperature of gas

P_1 = Upstream pressure (absolute)

P_2 = Downstream pressure (absolute)

This equation implies the continuous measurement of absolute gas pressure (P_1) and absolute gas temperature (T) inside the pipe, in addition to the differential pressure produced by the orifice plate ($P_1 - P_2$). These measurements may be taken by three separate devices, their signals routed to a gas flow computer:



Note the location of the RTD (thermowell), positioned downstream of the orifice plate so the turbulence it generates will not create additional turbulence at the orifice plate. The American Gas Association (AGA) allows for upstream placement of the thermowell, but only if located at least three feet upstream of a flow conditioner²⁷.

In order to best control all the physical parameters necessary for good orifice metering accuracy, it is standard practice for custody transfer flowmeter installations to use *honed meter runs* rather than standard pipe and pipe fittings. A “honed run” is a complete piping assembly consisting of a manufactured fitting to hold the orifice plate and sufficient straight lengths of pipe upstream and downstream, the interior surfaces of that pipe machined (“honed”) to have a glass-smooth surface with precise and symmetrical dimensions. Honed runs ensure minimum disruption to the flowing gas or liquid, thus improving measurement accuracy by avoiding unnecessary turbulence and/or distorted flow profiles. Such piping “runs” are quite expensive, but necessary to achieve flow measurement accuracy worthy of custody transfer.

²⁷Specified in Part 2 of the AGA Report #3, section 2.6.5, page 22. A major reason for this is von Kármán vortex shedding caused by the gas having to flow around the width of the thermowell. The “street” of vortices shed by the thermowell will cause serious pressure fluctuations at the orifice plate unless mitigated by a flow conditioner, or by locating the thermowell downstream so that the vortices do not reach the orifice.

This photograph shows a set of AGA3-compliant orifice meter runs measuring the flow of natural gas:




Note the special transmitter manifolds, built to accept both the differential pressure and absolute pressure (Rosemount model 3051) transmitters. Also note the quick-change fittings (the ribbed metal housings) holding the orifice plates, to facilitate convenient change-out of the orifice plates which is periodically necessary due to wear. It is not unheard of to replace orifice plates on a daily basis in some industries to ensure the sharp orifice edges necessary for accurate measurement²⁸.

Although not visible in this photograph, these meter runs are connected together by a network of shut-off valves directing the flow of natural gas through as few meter runs as desired. When the total gas flow rate is great, all meter runs are placed into service and their respective flow rates summed to yield a total flow measurement. When the total flow rate decreases, individual meter runs are shut off, resulting in increased flow rates through the remaining meter runs. This “staging” of meter runs expands the effective *turndown* or *rangeability* of the orifice plate as a flow-sensing element, resulting in much more accurate flow measurement over a wide range of flow rates than if a single (large) orifice meter run were used.

²⁸This is especially true in the gas exploration industry, where natural gas coming out of the well is laden with mineral debris.

An alternative to multiple instruments (differential pressure, absolute pressure, and temperature) installed on each meter run is to use a single *multi-variable* transmitter capable of measuring gas temperature as well as both static and differential pressures. This approach enjoys the advantage of simpler installation over the multi-instrument approach:



The Rosemount model 3095MV and Yokogawa model EJX910 are examples of multi-variable transmitters designed to perform compensated gas flow measurement, equipped with multiple pressure sensors, a connection port for an RTD temperature sensor, and sufficient digital computing power to continuously calculate flow rate based on the AGA equation. Such multi-variable transmitters may provide an analog output for computed flow rate, or a digital output where all three primary variables *and* the computed flow rate may be transmitted to a host system (as shown in the previous illustration). The Yokogawa EJX910A provides an interesting signal output option: a digital *pulse* signal, where each pulse represents a specific quantity (either volume or mass) of fluid. The frequency of this pulse train represents flow rate, while the total number of pulses counted over a period of time represents the total amount of fluid that has passed through the orifice plate over that amount of time.

This photograph shows a Rosemount 3095MV transmitter used to measure mass flow on a pure oxygen (gas) line. The orifice plate is an “integral” unit immediately below the transmitter body, sandwiched between two flange plates on the copper line. A three-valve manifold interfaces the model 3095MV transmitter to the integral orifice plate structure:



The temperature-compensation RTD may be clearly seen on the left-hand side of the photograph, installed at the elbow fitting in the copper pipe.

Liquid flow measurement applications may also benefit from compensation, because liquid density changes with temperature. Static pressure is not a concern here, because liquids are considered incompressible for all practical purposes²⁹. Thus, the formula for compensated liquid flow measurement does not include any terms for static pressure, just differential pressure and temperature:

$$Q = N \frac{CYA_2}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} \sqrt{(P_1 - P_2)[1 + k_T(T - T_{ref})]}$$

The constant k_T shown in the above equation is the proportionality factor for liquid expansion with increasing temperature. The difference in temperature between the measured condition (T) and the reference condition (T_{ref}) multiplied by this factor determines how much less dense the liquid is compared to its density at the reference temperature. It should be noted that some liquids – notably hydrocarbons – have thermal expansion factors significantly greater than water. This makes temperature compensation for hydrocarbon liquid flow measurement very important if the measurement principle is volumetric rather than mass-based.

²⁹Liquids can and do compress, the measurement of their “compressibility” being what is called the *bulk modulus*. However, this compressibility is too slight to be of any consequence in most flow measurement applications. A notable exception is the metering of diesel fuel through a high-pressure injection pump, where liquid pressures range in the *tens of thousands* of PSI, and the compressibility of the liquid diesel fuel may affect the precise timing of individual injections into the engine cylinders.

22.1.9 Equation summary

Volumetric flow rate (Q) full equation:

$$Q = N \frac{CY A_2}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} \sqrt{\frac{P_1 - P_2}{\rho_f}}$$

Volumetric flow rate (Q) simplified equation:

$$Q = k \sqrt{\frac{P_1 - P_2}{\rho_f}}$$

Mass flow rate (W):

$$W = N \frac{CY A_2}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} \sqrt{\rho_f (P_1 - P_2)}$$

Mass flow rate (W) simplified equation:

$$W = k \sqrt{\rho_f (P_1 - P_2)}$$

Where,

Q = Volumetric flow rate (e.g. gallons per minute, flowing cubic feet per second)

W = Mass flow rate (e.g. kilograms per second, slugs per minute)

N = Unit conversion factor

C = Discharge coefficient (accounts for energy losses, Reynolds number corrections, pressure tap locations, etc.)

Y = Gas expansion factor ($Y = 1$ for liquids)

A_1 = Cross-sectional area of mouth

A_2 = Cross-sectional area of throat

ρ_f = Fluid density at flowing conditions (actual temperature and pressure at the element)

k = Constant of proportionality (determined by experimental measurements of flow rate, pressure, and density)

The beta ratio (β) of a differential-producing element is the ratio of throat diameter to mouth diameter ($\beta = \frac{d}{D}$). This is the primary factor determining acceleration as the fluid increases velocity entering the constricted throat of a flow element (venturi tube, orifice plate, wedge, etc.). The following expression is often called the *velocity of approach factor* (commonly symbolized as E_v), because it relates the velocity of the fluid through the constriction to the velocity of the fluid as it approaches the flow element:

$$E_v = \frac{1}{\sqrt{1 - \beta^4}} = \text{Velocity of approach factor}$$

This same velocity approach factor may be expressed in terms of mouth and throat areas (A_1 and A_2 , respectively):

$$E_v = \frac{1}{\sqrt{1 - \left(\frac{A_2}{A_1}\right)^2}} = \text{Velocity of approach factor}$$

Beta ratio has a significant impact on the number of straight-run pipe lengths needed to condition the flow profile upstream and downstream of the flow element. Large beta ratios (where the bore diameter approaches the flowtube's inside diameter) are more sensitive to piping disturbances, since there is less acceleration of the flowstream through the element, and therefore flow profile asymmetries caused by piping disturbances are significant in comparison to the fluid's through-bore velocity. Small beta ratio values correspond to larger acceleration factors, where disturbances in the flow profile become "swamped"³⁰ by the high throat velocities created by the element's constriction. A disadvantage of small beta ratio values is that the flow element exhibits a greater permanent pressure loss, which is an operational cost if the flow is provided by a machine such as an engine- or motor-driven pump (more energy required to turn the pump, equating to a greater operating cost to run the process).

³⁰ "Swamping" is a term commonly used in electrical engineering, where a bad effect is overshadowed by some other effect much larger in magnitude, to the point where the undesirable effect is negligible in comparison.

When computing the volumetric flow of a gas in *standard* volume units (e.g. SCFM), the equation becomes much more complex than the simple (flowing) volumetric rate equation. Any equation computing flow in standard units must predict the effective expansion of the gas if it were to transition from flowing conditions (the actual pressure and temperature it experiences flowing through the pipe) to standard conditions (one atmosphere pressure at 60 degrees Fahrenheit). The compensated gas flow measurement equation published by the American Gas Association (AGA Report #3) in 1992 for orifice plates with flange taps calculates this expansion to standard conditions with a series of factors accounting for flowing and standard (“base”) conditions, in addition to the more common factors such as velocity of approach and gas expansion. Most of these factors are represented in the AGA3 equation by different variables beginning with the letter *F*:

$$Q = F_n(F_c + F_{sl})YF_{pb}F_{tb}F_{tf}F_{gr}F_{pv}\sqrt{h_W P_{f1}}$$

Where,

Q = Volumetric flow rate (standard cubic feet per hour – SCFH)

F_n = Numeric conversion factor (accounts for certain numeric constants, unit-conversion coefficients, and the velocity of approach factor E_v)

F_c = Orifice calculation factor (a polynomial function of the orifice plate’s β ratio and Reynolds number), appropriate for flange taps

F_{sl} = Slope factor (another polynomial function of the orifice plate’s β ratio and Reynolds number), appropriate for flange taps

$F_c + F_{sl} = C_d$ = Discharge coefficient, appropriate for flange taps

Y = Gas expansion factor (a function of β , differential pressure, static pressure, and specific heats)

F_{pb} = Base pressure factor = $\frac{14.73 \text{ PSI}}{P_b}$, with pressure in PSIA (absolute)

F_{tb} = Base temperature factor = $\frac{T_b}{519.67}$, with temperature in degrees Rankine

F_{tf} = Flowing temperature factor = $\sqrt{\frac{519.67}{T_f}}$, with temperature in degrees Rankine

F_{gr} = Real gas relative density factor = $\sqrt{\frac{1}{G_r}}$

F_{pv} = Supercompressibility factor = $\sqrt{\frac{Z_b}{Z_{f1}}}$

h_W = Differential pressure produced by orifice plate (inches water column)

P_{f1} = Flowing pressure of gas at the upstream tap (PSI absolute)

22.2 Laminar flowmeters

A unique form of differential pressure-based flow measurement deserves its own section in this flow measurement chapter, and that is the *laminar* flowmeter.

Laminar flow is a condition of fluid motion where viscous (internal fluid friction) forces greatly overshadow inertial (kinetic) forces. A flowstream in a state of laminar flow exhibits no turbulence, with each fluid molecule traveling in its own path, with limited mixing and collisions with adjacent molecules. The dominant mechanism for resistance to fluid motion in a laminar flow regime is friction with the pipe or tube walls. Laminar flow is qualitatively predicted by low values of Reynolds number.

This pressure drop created by fluid friction in a laminar flowstream is quantifiable, and is expressed in the Hagen-Poiseuille equation:

$$Q = k \left(\frac{\Delta P D^4}{\mu L} \right)$$

Where,

Q = Flow rate

ΔP = Pressure dropped across a length of pipe


D = Pipe diameter

μ = Fluid viscosity

L = Pipe length

k = Coefficient accounting for units of measurement

Laminar flowmeter elements generally consist of one or more tubes whose length greatly exceeds the inside diameter, arranged in such a way as to produce a slow-moving flow velocity. An example is shown here:



The expanded diameter of the flow element ensures a lower fluid velocity than in the pipes entering and exiting the element. This decreases the Reynolds number to the point where the flow

regime exhibits laminar behavior. The large number of small-diameter tubes packed in the wide area of the element provide adequate wall surface area for the fluid's viscosity to act upon, creating an overall pressure drop from inlet to outlet which is measured by the differential pressure transmitter. This pressure drop is permanent (no recovery of pressure downstream) because the mechanism of pressure drop is friction: total dissipation (loss) of energy in the form of heat.

Another common form of laminar flow element is simply a coiled *capillary tube*: a long tube with a very small inside diameter. The small inside diameter of such a tube makes wall-boundary effects dominant, such that the flow regime will remain laminar over a wide range of flow rates. The extremely restrictive nature of a capillary tube, of course, limits the use of such flow elements to very low flow rates such as those encountered in the sampling networks of certain analytical instruments.

A unique advantage of the laminar flowmeter is its linear relationship between flow rate and developed pressure drop. It is the only pressure-based flow measurement device for filled pipes that exhibits a linear pressure/flow relationship. This means no “square-root” characterization is necessary to obtain linear flow measurements with a laminar flowmeter. The big disadvantage of this meter type is its dependence on fluid viscosity, which in turn is strongly influenced by fluid temperature. Thus, all laminar flowmeters require temperature compensation in order to derive accurate measurements, and some even use temperature *control* systems to force the fluid's temperature to be constant as it moves through the element³¹.

Laminar flow elements find their widest application inside pneumatic instruments, where a linear pressure/flow relationship is highly advantageous (behaving like a “resistor” for instrument air flow) and the viscosity of the fluid (instrument air) is relatively constant. Pneumatic controllers, for instance, use laminar restrictors as part of the derivative and integral calculation modules, the combination of “resistance” from the restrictor and “capacitance” from volume chambers forming a sort of pneumatic time-constant (τ) network.


22.3 Variable-area flowmeters

An *variable-area* flowmeter is one where the fluid must pass through a restriction whose area increases with flow rate. This stands in contrast to flowmeters such as orifice plates and venturi tubes where the cross-sectional area of the flow element remains fixed.

³¹This includes elaborate oil-bath systems where the laminar flow element is submerged in a temperature-controlled oil bath, the purpose of which is to hold temperature inside the laminar element constant despite sudden changes in the measured fluid's temperature.


22.3.1 Rotameters

The simplest example of a variable-area flowmeter is the *rotameter*, which uses a solid object (called a *plummet* or *float*) as a flow indicator, suspended in the midst of a tapered tube:



As fluid flows upward through the tube, a pressure differential develops across the plummet. This pressure differential, acting on the effective area of the plummet body, develops an upward force ($F = PA$). If this force exceeds the weight of the plummet, the plummet moves up. As the plummet moves farther up in the tapered tube, the area between the plummet and the tube walls (through which the fluid must travel) grows larger. This increased flowing area allows the fluid to make it past the plummet without having to accelerate as much, thereby developing less pressure drop across the plummet's body. At some point, the flowing area reaches a point where the pressure-induced force on the plummet body exactly matches the weight of the plummet. This is the point in the tube where the plummet stops moving, indicating flow rate by its position relative to a scale mounted (or etched) on the outside of the tube.

The following rotameter uses a spherical plummet, suspended in a flow tube machined from a solid block of clear plastic. An adjustable valve at the bottom of the rotameter provides a means for adjusting gas flow:



The same basic flow equation used for pressure-based flow elements holds true for rotameters as well:

$$Q = k \sqrt{\frac{P_1 - P_2}{\rho}}$$

However, the difference in this application is that the value inside the radicand is constant, since the pressure difference will remain constant³² and the fluid density will likely remain constant as well. Thus, k will change in proportion to Q . The only variable within k relevant to plummet position is the flowing area between the plummet and the tube walls.

³²If we know that the plummet's weight will remain constant, its drag area will remain constant, and that the force generated by the pressure drop will always be in equilibrium with the plummet's weight for any steady flow rate, then the relationship $F = PA$ dictates a constant pressure. Thus, we may classify the rotameter as a *constant-pressure, variable-area* flowmeter. This stands in contrast to devices such as orifice plates, which are *variable-pressure, constant-area*.


Most rotameters are indicating devices only. They may be equipped to transmit flow information electronically by adding sensors to detect the plummet's position in the tube, but this is not common practice.

Rotameters are very commonly used as purge flow indicators for pressure and level measurement systems requiring a constant flow of purge fluid (see sections 19.6.6 and 20.3.2 for practical examples). Such rotameters are usually equipped with hand-adjustable needle valves for manual regulation of purge fluid flow rate.

22.3.2 Weirs and flumes

A very different style of variable-area flowmeter is used extensively to measure flow rate through open channels, such as irrigation ditches. If an obstruction is placed within a channel, any liquid flowing through the channel must rise on the upstream side of the obstruction. By measuring this liquid level rise, it is possible to infer the rate of liquid flow past the obstruction.

The first form of open-channel flowmeter is the *weir*, which is nothing more than a dam obstructing passage of liquid through the channel. Three styles of weir are shown in the following illustration; the *rectangular*, *Cippoletti*, and *V-notch*:




A rectangular weir has a notch of simple rectangular shape, as the name implies. A Cippoletti weir is much like a rectangular weir, except that the vertical sides of the notch have a 4:1 slope (rise of 4, run of 1; approximately a 14 degree angle from vertical). A V-notch weir has a triangular notch, customarily measuring either 60 or 90 degrees.

The following photograph shows water flowing through a Cippoletti weir made of 1/4 inch steel plate:



At a condition of zero flow through the channel, the liquid level will be at or below the crest (lowest point on the opening) of the weir. As liquid begins to flow through the channel, it must spill over the crest of the weir in order to get past the weir and continue downstream in the channel. In order for this to happen, the level of the liquid upstream of the weir must rise above the weir's crest height. This height of liquid upstream of the weir represents a hydrostatic pressure, much the same as liquid heights in piezometer tubes represent pressures in a liquid flowstream through an enclosed pipe (see section 2.11.15 for examples of this). The height of liquid above the crest of a weir is

analogous to the pressure differential generated by an orifice plate. As liquid flow is increased even more, a greater pressure (head) will be generated upstream of the weir, forcing the liquid level to rise. This effectively increases the cross-sectional area of the weir's "throat" as a taller stream of liquid exits the notch of the weir³³.



³³Orifice plates are *variable-pressure, constant-area* flowmeters. Rotameters are *constant-pressure, variable-area* flowmeters. Weirs are *variable-pressure, variable-area* flowmeters. As one might expect, the mathematical functions describing each of these flowmeter types is unique!

This dependence of notch area on flow rate creates a very different relationship between flow rate and liquid height (measured above the crest) than the relationship between flow rate and differential pressure in an orifice plate:

$$Q = 3.33(L - 0.2H)H^{1.5} \quad \text{Rectangular weir}$$

$$Q = 3.367LH^{1.5} \quad \text{Cippoletti weir}$$

$$Q = 2.48 \left(\tan \frac{\theta}{2} \right) H^{2.5} \quad \text{V-notch weir}$$

Where,


Q = Volumetric flow rate (cubic feet per second – CFS)

L = Width of crest (feet)

θ = V-notch angle (degrees)

H = Head (feet)

As you can see from a comparison of characteristic flow equations between these three types of weirs, the shape of the weir's notch has a dramatic effect on the mathematical relationship between flow rate and head (liquid level upstream of the weir, measured above the crest height). This implies that it is possible to create almost any characteristic equation we might like just by carefully shaping the weir's notch in some custom form. A good example of this is the so-called *proportional* or *Sutro* weir, which is designed to have a linear relationship between head and flow rate:




Sutro weirs are not used very often, due to their inherently weak structure and tendency to clog with debris.

A rare example of a Sutro weir appears in the following photograph, discharging flow from a lake into a stream:



The metal plates forming the weir's shape are quite thick (about 1/2 inch) to give the weir sufficient strength. A good construction practice seen on this Sutro weir, but recommended on *all* weir designs, is to bevel the downstream edge of the weir plate much like a standard orifice plate profile. The beveled edge provides a minimum-friction passageway for the liquid as it spills through the weir's opening.

A variation on the theme of a weir is another open-channel device called a *flume*. If weirs may be thought of as open-channel orifice plates, then flumes may be thought of as open-channel venturi tubes:



Like weirs, flumes generate upstream liquid level height changes indicative of flow rate. One of the most common flume design is the *Parshall flume*, named after its inventor R.L. Parshall when it was developed in the year 1920.

The following formulae relate head (upstream liquid height) to flow rate for free-flowing Parshall flumes³⁴:

$$Q = 0.992H^{1.547} \quad \text{3-inch wide throat Parshall flume}$$

$$Q = 2.06H^{1.58} \quad \text{6-inch wide throat Parshall flume}$$

$$Q = 3.07H^{1.53} \quad \text{9-inch wide throat Parshall flume}$$

$$Q = 4LH^{1.53} \quad \text{1-foot to 8-foot wide throat Parshall flume}$$

$$Q = (3.6875L + 2.5)H^{1.53} \quad \text{10-foot to 50-foot wide throat Parshall flume}$$

Where,

Q = Volumetric flow rate (cubic feet per second – CFS)

L = Width of flume throat (feet)

H = Head (feet)

Flumes are generally less accurate than weirs, but they do enjoy the advantage of being inherently self-cleaning. If the liquid stream being measured is drainage- or waste-water, a substantial amount of solid debris may be present in the flow that could cause repeated clogging problems for weirs. In such applications, flumes are often the more practical flow element for the task (and more accurate

³⁴It is also possible to operate a Parshall flume in fully *submerged* mode, where liquid level must be measured at both the upstream and throat sections of the flume. Correction factors must be applied to these equations if the flume is submerged.

over the long term as well, since even the finest weir will not register accurately once fouled by debris).

Once a weir or flume has been installed in an open channel to measure the flow of liquid, some method must be employed to sense upstream liquid level and translate this level measurement into a flow measurement. Perhaps the most common technology for weir/flume level sensing is *ultrasonic* (see section 20.5.1 beginning on page 1468 for more information on how this technology works). Ultrasonic level sensors are completely non-contact, which means they cannot become fouled by the process liquid (or debris in the process liquid). However, they may be “fooled” by foam or debris floating on top of the liquid, as well as waves on the liquid surface.


The following photograph shows a Parshall flume measuring effluent flow from a municipal sewage treatment plant, with an ultrasonic transducer mounted above the middle of the flume to detect water level flowing through:




Once the liquid level is successfully measured, a computing device is used to translate that level measurement into a suitable flow measurement (and in some cases even integrate that flow measurement with respect to time to arrive at a value for total liquid volume passed through the element, in accordance with the calculus relationship $V = \int Q dt + C$).

A technique for providing a clean and “quiet” (still) liquid surface to measure the level of is called a *stilling well*. This is an open-top chamber connected to the weir/flume channel by a pipe, so the liquid level in the stilling well matches the liquid level in the channel. The following illustration shows a stilling well connected to a weir/flume channel, with the direction of liquid flow in the

channel being perpendicular to the page (i.e. either coming toward your eyes or going away from your eyes):




To discourage plugging of the passageway connecting the stilling well to the channel, a small flow rate of clean water may be introduced into the well. This forms a constant *purge flow* into the channel, flushing out debris that might otherwise find its way into the connecting passageway to plug it up. Note how the purge water enters the stilling well through a submerged tube, so it does not cause splashing on the water's surface inside the well which could cause measurement problems for the ultrasonic sensor:




A significant advantage that weirs and flumes have over other forms of flow measurement is exceptionally high *rangeability*: the ability to measure very wide ranges of flow with a modest pressure (height) span. Another way to state this is to say that the accuracy of a weir or flume is quite high even at low flow rates.

Earlier in this section you saw a three-image representation of liquid flow through a rectangular weir. As fluid flow rate increased, so did the height (head) of the liquid upstream of the weir:



The height of liquid upstream of the weir depends on the flow rate (volumetric Q or mass W) as well as the effective area of the notch through which the fluid must pass. Unlike an orifice plate, this area changes with flow rate in both weirs and flumes. One way to envision this by comparison is to imagine a weir as acting like an elastic orifice plate, whose bore area increases with flow rate. This flow-dependent notch area exhibited by both weirs and flumes means that these devices become *more sensitive* to changes in flow as the flow rate becomes smaller.

A comparison of transfer function graphs for closed-pipe head elements such as orifice plates and venturi tubes versus weirs and flumes shows this striking difference in characteristics:




Looking at the orifice plate / venturi tube graph near the lower-left corner, you can see how small changes in flow result in extremely small changes in head (differential pressure), because the

function has a very low slope (small $\frac{dH}{dQ}$) at that end. By comparison, a weir or flume produces relatively large changes in head (liquid elevation) for small changes in flow near the bottom end of the range, because the function has a very steep slope (large $\frac{dH}{dQ}$) at that end.

The practical advantage this gives weirs and flumes is the ability to maintain high accuracy of flow measurement at very low flow rates – something a fixed-orifice element simply cannot do. It is commonly understood in industry that traditional orifice plate flowmeters do not maintain good measurement accuracy much below a third of their full-range flow (a rangeability or *turndown* of 3:1), whereas weirs (especially the V-notch design) can achieve far greater turndown (up to 500:1 according to some sources³⁵).

22.4 Velocity-based flowmeters

The Law of Continuity for fluids states that the product of mass density (ρ), cross-sectional pipe area (A) and average velocity (\bar{v}) must remain constant through any continuous length of pipe:



If the density of the fluid is not subject to change as it travels through the pipe (a very good assumption for liquids), we may simplify the Law of Continuity by eliminating the density terms from the equation:


$$A_1 \bar{v}_1 = A_2 \bar{v}_2$$

The product of cross-sectional pipe area and average fluid velocity is the volumetric flow rate of the fluid through the pipe ($Q = A\bar{v}$). This tells us that fluid velocity will be directly proportional to volumetric flow rate given a known cross-sectional area and a constant density for the fluid flowstream. Any device able to directly measure fluid velocity is therefore capable of inferring volumetric flow rate of fluid in a pipe. This is the basis for *velocity-based* flowmeter designs.


³⁵These figures are reported in Béla Lipták's excellent reference book *Instrument Engineers' Handbook – Process Measurement and Analysis Volume I* (Fourth Edition). To be fair to closed-pipe elements such as orifice plates and venturi tubes, much improvement in the classic 3:1 rangeability limitation has been achieved through the use of microprocessor-based differential pressure sensors. Lipták reports rangeabilities for orifice plates as great as 10:1 through the use of such modern differential pressure instruments. However, even this pales in comparison to the rangeability of a typical weir or flume, which Lipták reports to be 75:1 for “most devices” in this category.

22.4.1 Turbine flowmeters

Turbine flowmeters use a free-spinning turbine wheel to measure fluid velocity, much like a miniature windmill installed in the flow stream. The fundamental design goal of a turbine flowmeter is to make the turbine element as free-spinning as possible, so no torque will be required to sustain the turbine's rotation. If this goal is achieved, the turbine blades will achieve a rotating (tip) velocity directly proportional to the linear velocity of the fluid, whether that fluid is a gas or a liquid:




The mathematical relationship between fluid velocity and turbine tip velocity – assuming frictionless conditions – is a ratio defined by the *tangent* of the turbine blade angle:



For a 45° blade angle, the relationship is 1:1, with tip velocity equaling fluid velocity. Smaller blade angles (each blade closer to parallel with the fluid velocity vector) result in the tip velocity being a fractional proportion of fluid velocity.

Turbine tip velocity is quite easy to sense using a magnetic sensor, generating a voltage pulse each time one of the ferromagnetic turbine blades passes by. Traditionally, this sensor is nothing more than a coil of wire in proximity to a stationary magnet, called a *pickup coil* or *pickoff coil* because it “picks” (senses) the passing of the turbine blades. Magnetic flux through the coil’s center increases and decreases as the passing of the steel turbine blades presents a varying reluctance (“resistance” to magnetic flux), causing voltage pulses equal in frequency to the number of blades passing by each second. It is the *frequency* of this signal that represents fluid velocity, and therefore volumetric flow rate.

A cut-away demonstration model of a turbine flowmeter is shown in the following photograph. The blade sensor may be seen protruding from the top of the flowtube, just above the turbine wheel:



Note the sets of “flow conditioner” vanes immediately before and after the turbine wheel in the photograph. As one might expect, turbine flowmeters are very sensitive to *swirl* in the process fluid flowstream. In order to achieve high accuracy, the flow profile must not be swirling in the vicinity of the turbine, lest the turbine wheel spin faster or slower than it should to represent the velocity of a straight-flowing fluid. A minimum straight-pipe length of 20 pipe diameters upstream and 5 pipe diameters downstream is typical for turbine flowmeters in order to dissipate swirl from piping disturbances.

Mechanical gears and rotating cables have also been historically used to link a turbine flowmeter’s turbine wheel to indicators. These designs suffer from greater friction than electronic (“pickup coil”) designs, potentially resulting in more measurement error (less flow indicated than there actually is, because the turbine wheel is slowed by friction). One advantage of mechanical turbine flowmeters, though, is the ability to maintain a running total of gas usage by turning a simple odometer-style

totalizer. This design is often used when the purpose of the flowmeter is to track total fuel gas consumption (e.g. natural gas used by a commercial or industrial facility) for billing.

In an electronic turbine flowmeter, volumetric flow is directly and linearly proportional to pickup coil output frequency. We may express this relationship in the form of an equation:

$$f = kQ$$

Where,

f = Frequency of output signal (Hz, equivalent to pulses per second)

Q = Volumetric flow rate (e.g. gallons per second)

k = "K" factor of the turbine element (e.g. pulses per gallon)

Dimensional analysis confirms the validity of this equation. Using units of GPS (gallons per second) and pulses per gallon, we see that the product of these two quantities is indeed pulses per second (equivalent to cycles per second, or Hz):

$$\left[\frac{\text{Pulses}}{\text{s}} \right] = \left[\frac{\text{Pulses}}{\text{gal}} \right] \left[\frac{\text{gal}}{\text{s}} \right]$$

Using algebra to solve for flow (Q), we see that it is the quotient of frequency and k factor that yields a volumetric flow rate for a turbine flowmeter:

$$Q = \frac{f}{k}$$

The inherent linearity of a turbine flowmeter is a tremendous advantage over nonlinear flow elements such as venturi tubes and orifice plates because this linearity results in a much greater turndown ratio for accurate flow measurement. Contrasted against common orifice-type meters which are usually limited to turndown ratios of 4:1 at best, turbine meters commonly exceed turndown ratios of 10:1.

If pickup signal frequency directly represents volumetric flow rate, then the total number of pulses accumulated in any given time span will represent the amount of fluid volume (V) passed through the turbine meter over that same time span. We may express this algebraically as the product of average flow rate (\bar{Q}), average frequency (\bar{f}), k factor, and time:

$$V = \bar{Q}t = \frac{\bar{f}t}{k}$$

A more sophisticated way of calculating total volume passed through a turbine meter requires calculus, representing change in volume as the time-integral of instantaneous signal frequency and k factor over a period of time from $t = 0$ to $t = T$:

$$\Delta V = \int_0^T Q dt \quad \text{or} \quad \Delta V = \int_0^T \frac{f}{k} dt$$

We may achieve approximately the same result simply by using a digital counter circuit to totalize pulses output by the pickup coil and a microprocessor to calculate volume in whatever unit of measurement we deem appropriate.

As with the orifice plate flow element, standards have been drafted for the use of turbine flowmeters as precision measuring instruments in gas flow applications, particularly the custody transfer³⁶ of natural gas. The American Gas Association has published a standard called the Report #7 specifying the installation of turbine flowmeters for high-accuracy gas flow measurement, along with the associated mathematics for precisely calculating flow rate based on turbine speed, gas pressure, and gas temperature.

Pressure and temperature compensation is relevant to turbine flowmeters in gas flow applications because the density of the gas is a strong function of both pressure and temperature. The turbine wheel itself only senses gas *velocity*, and so these other factors must be taken into consideration to accurately calculate mass flow (or *standard* volumetric flow; e.g. SCFM).

In high-accuracy applications, it is important to individually determine the k factor for a turbine flowmeter's calibration. Manufacturing variations from flowmeter to flowmeter make precise duplication of k factor challenging, and so a flowmeter destined for high-accuracy measurement should be tested against a "flow prover" in a calibration laboratory to empirically determine its k factor. If possible, the best way to test the flowmeter's k factor is to connect the prover to the meter on site where it will be used. This way, the any effects due to the piping before and after the flowmeter will be incorporated in the measured k factor.

³⁶ "Custody transfer" refers to measurement applications where a product is exchanging ownership. In other words, someone is selling, and someone else is buying, quantities of fluid as part of a business transaction. It is not difficult to understand why accuracy is important in such applications, as both parties have a vested interest in a fair exchange. Government institutions also have a stake in accurate metering, as taxes are typically levied on the sale of commodity fluids such as natural gas.

The following photograph shows three AGA7-compliant installations of turbine flowmeters for measuring the flow rate of natural gas:




Note the pressure-sensing and temperature-sensing instrumentation installed in the pipe, reporting gas pressure and gas temperature to a flow-calculating computer (along with turbine pulse frequency) for the calculation of natural gas flow rate.

Less-critical gas flow measurement applications may use a “compensated” turbine flowmeter that mechanically performs the same pressure- and temperature-compensation functions on turbine speed to achieve true gas flow measurement, as shown in the following photograph:




The particular flowmeter shown in the above photograph uses a filled-bulb temperature sensor (note the coiled, armored capillary tube connecting the flowmeter to the bulb) and shows total gas flow by a series of pointers, rather than gas flow *rate*.

A variation on the theme of turbine flow measurement is the *paddlewheel* flowmeter, a very inexpensive technology usually implemented in the form of an insertion-type sensor. In this instrument, a small wheel equipped with “paddles” parallel to the shaft is inserted in the flowstream, with half the wheel shrouded from the flow. A photograph of a plastic paddlewheel flowmeter appears here:



A surprisingly sophisticated method of “pickup” for the plastic paddlewheel shown in the photograph uses *fiber-optic cables* to send and receive light. One cable sends a beam of light to the edge of the paddlewheel, and the other cable receives light on the other side of the paddlewheel. As the paddlewheel turns, the paddles alternately block and pass the light beam, resulting in a pulsed light beam at the receiving cable. The frequency of this pulsing is, of course, directly proportional to volumetric flow rate.

The external ends of the two fiber optic cables appear in this next photograph, ready to connect to a light source and light pulse sensor to convert the paddlewheel's motion into an electronic signal:



A problem common to all turbine flowmeters is that of the turbine “coasting” when the fluid flow suddenly stops. This is more often a problem in batch processes than continuous processes, where the fluid flow is regularly turned on and shut off. This problem may be minimized by configuring the measurement system to ignore turbine flowmeter signals any time the automatic shutoff valve reaches the “shut” position. This way, when the shutoff valve closes and fluid flow immediately halts, any coasting of the turbine wheel will be irrelevant. In processes where the fluid flow happens to pulse for reasons other than the control system opening and shutting automatic valves, this problem is more severe.


Another problem common to all turbine flowmeters is lubrication of the turbine bearings. Frictionless motion of the turbine wheel is essential for accurate flow measurement, which is a daunting design goal for the flowmeter manufacturing engineers. The problem is not as severe in applications where the process fluid is naturally lubricating (e.g. diesel fuel), but in applications such as natural gas flow where the fluid provides no lubrication to the turbine bearings, external lubrication must be supplied. This is often a regular maintenance task for instrument technicians: using a hand pump to inject light-weight “turbine oil” into the bearing assemblies of turbine flowmeters used in gas service.

Process fluid viscosity is another source of friction for the turbine wheel. Fluids with high viscosity (e.g. heavy oils) will tend to slow down the turbine's rotation even if the turbine rotates on frictionless bearings. This effect is especially pronounced at low flow rates, which leads to a *minimum linear flow* rating for the flowmeter: a flowrate below which it refuses to register proportionately to fluid flow rate.


22.4.2 Vortex flowmeters

When a fluid moves with high Reynolds number past a stationary object (a “bluff body”), there is a tendency for the fluid to form *vortices* on either side of the object. Each vortex will form, then detach from the object and continue to move with the flowing gas or liquid, one side at a time in alternating fashion. This phenomenon is known as *vortex shedding*, and the pattern of moving vortices carried downstream of the stationary object is known as a *vortex street*.

It is commonplace to see the effects of vortex shedding on a windy day by observing the motion of flagpoles, light poles, and tall smokestacks. Each of these objects has a tendency to oscillate perpendicular to the direction of the wind, owing to the pressure variations caused by the vortices as they alternately form and break away from the object:



This alternating series of vortices was studied by Vincenc Strouhal in the late nineteenth century and later by Theodore von Kármán in the early twentieth century. It was determined that the distance between successive vortices downstream of the stationary object is relatively constant, and directly proportional to the width of the object, for a wide range of Reynolds number values³⁷. If we view these vortices as crests of a continuous wave, the distance between vortices may be represented by the symbol customarily reserved for wavelength: the Greek letter “lambda” (λ).





The proportionality between object width (d) and vortex street wavelength (λ) is called the *Strouhal number* (S), approximately equal to 0.17:

$$\lambda S = d \quad \lambda \approx \frac{d}{0.17}$$

³⁷It is important to note that the vortex-shedding phenomenon ceases altogether if the Reynolds number is too low. Laminar flow produces no vortices, but rather stream-line flow around any object placed in its way.

If a differential pressure sensor is installed immediately downstream of the stationary object in such an orientation that it detects the passing vortices as pressure variations, an alternating signal will be detected:



The *frequency* of this alternating pressure signal is directly proportional to fluid velocity past the object, since the wavelength is constant. This follows the classic frequency-velocity-wavelength formula common to all traveling waves ($\lambda f = v$). Since we know the wavelength will be equal to the bluff body's width divided by the Strouhal number (approximately 0.17), we may substitute this into the frequency-velocity-wavelength formula to solve for fluid velocity (v) in terms of signal frequency (f) and bluff body width (d).

$$v = \lambda f$$

$$v = \frac{d}{0.17} f$$

$$v = \frac{df}{0.17}$$

Thus, a stationary object and pressure sensor installed in the middle of a pipe section constitute a form of flowmeter called a *vortex flowmeter*. Like a turbine flowmeter with an electronic “pickup” sensor to detect the passage of rotating turbine blades, the output frequency of a vortex flowmeter is linearly proportional to volumetric flow rate.

The pressure sensors used in vortex flowmeters are not standard differential pressure transmitters, since the vortex frequency is too high to be successfully detected by such bulky instruments. Instead, the sensors are typically piezoelectric crystals. These pressure sensors need not be calibrated, since the amplitude of the pressure waves detected is irrelevant. Only the frequency of the waves matter

for measuring flow rate, and so nearly any pressure sensor with a fast enough response time will suffice.

Like turbine meters, the relationship between sensor frequency (f) and volumetric flow rate (Q) may be expressed as a proportionality, with the letter k used to represent the constant of proportionality for any particular flowmeter:

$$f = kQ$$

Where,

f = Frequency of output signal (Hz)

Q = Volumetric flow rate (e.g. gallons per second)³⁸

k = “K” factor of the vortex shedding flowtube (e.g. pulses per gallon)

This means vortex flowmeters, like electronic turbine meters, each have a particular “ k factor” relating the number of pulses generated per unit volume passed through the meter³⁹. Counting the total number of pulses over a certain time span yields total fluid volume passed through the meter over that same time span, making the vortex flowmeter readily adaptable for “totalizing” fluid volume just like turbine meters. The direct proportion between vortex frequency and volumetric flow rate also means vortex flowmeters are *linear-responding* instruments just like turbine flowmeters. Unlike orifice plates which exhibit a quadratic response, turbine and vortex flowmeters alike enjoy a wider range (turndown) of flow measurement and do not require special signal characterization to function properly.

Since vortex flowmeters have no moving parts, they do not suffer the problems of wear and lubrication facing turbine meters. There is no moving element to “coast” as in a turbine flowmeter if fluid flow suddenly stops, which means vortex flowmeters are better suited to measuring erratic flows.


A significant disadvantage of vortex meters is a behavior known as *low flow cutoff*, where the flowmeter simply stops working below a certain flow rate. The reason for this is *laminar* flow: at low flow rates (i.e. low Reynolds number values) the effects of fluid viscosity overwhelm fluid momentum, preventing vortices from forming. This cessation of vortices causes the vortex flowmeter to register absolutely no flow at all even when there is still some (laminar) flow through the pipe. At high flow rates (i.e. high Reynolds number values), fluid momentum is enough to overcome viscosity and produce vortices, and the vortex flowmeter works just fine.

The phenomenon of low-flow cutoff for a vortex flowmeter at first seems analogous to the *minimum linear flow* limitation of a turbine flowmeter. However, vortex flowmeter low-flow cutoff is actually a far more severe problem. If the volumetric flow rate through a turbine flowmeter falls below the minimum linear value, the turbine continues to spin, albeit slower than it should. If the volumetric flow rate through a vortex flowmeter falls below the low-flow cutoff value, however, the flowmeter’s signal *goes completely to zero*, indicating no flow at all. This idiosyncrasy makes vortex flowmeters entirely unsuitable in applications where the desired flow measurement range extends all the way down to zero.



³⁸Note that if flow rate is to be expressed in units of gallons per *minute* as is customary, the equation must contain a factor for minutes-to-seconds conversion: $f = \frac{kQ}{60}$

³⁹This k factor is empirically determined for each flowmeter by the manufacturer using water as the test fluid (a factory “wet-calibration”), to ensure optimum accuracy.

The following photograph shows a Rosemount model 8800C vortex flow transmitter:



The next two photographs show close-up views of the flowtube assembly, front (left) and rear (right):




Vortex flowmeters, like other velocity-based meters, are affected by large-scale turbulence in the fluid stream and therefore require some length of straight pipe both upstream and downstream of the flowmeter to properly characterize the flow. It is typical to install vortex flowmeters with 10 pipe diameters of straight-length pipe upstream and 5 pipe diameters downstream.

22.4.3 Magnetic flowmeters

When an electrical conductor moves perpendicular to a magnetic field, a voltage is induced in that conductor perpendicular to both the magnetic flux lines and the direction of motion. This phenomenon is known as *electromagnetic induction*, and it is the basic principle upon which all electro-mechanical generators operate.

In a generator mechanism, the conductor in question is typically a coil (or set of coils) made of copper wire. However, there is no reason the conductor must be made of copper wire. *Any* electrically conductive substance in motion is sufficient to electromagnetically induce a voltage, even if that substance is a liquid⁴⁰. Therefore, electromagnetic induction is a technique applicable to the measurement of liquid flow rates.

Consider water flowing through a pipe, with a magnetic field passing perpendicularly through the pipe:



The direction of liquid flow cuts perpendicularly through the lines of magnetic flux, generating a voltage along an axis perpendicular to both. Metal electrodes opposite each other in the pipe wall intercept this voltage, making it readable to an electronic circuit.

⁴⁰In a practical sense, only liquid flows are measurable using this technique. Gases must be super-heated into a *plasma* state before they are able to conduct electricity, and so electromagnetic flowmeters cannot be used with most industrial gas flowstreams.

A voltage induced by the linear motion of a conductor through a magnetic field is called *motional EMF*, the magnitude of which is predicted by the following formula (assuming perfect perpendicularity between the direction of velocity, the orientation of the magnetic flux lines, and the axis of voltage measurement):

$$\mathcal{E} = Blv$$

Where,

\mathcal{E} = Motional EMF (volts)

B = Magnetic flux density (Tesla)

l = Length of conductor passing through the magnetic field (meters)

v = Velocity of conductor (meters per second)

Assuming a fixed magnetic field strength (constant B) and an electrode spacing equal to the fixed diameter of the pipe (constant $l = d$), the only variable capable of influencing the magnitude of induced voltage is velocity (v). In our example, v is not the velocity of a wire segment, but rather the average velocity of the liquid flowstream (\bar{v}). Since we see that this voltage will be proportional to average fluid velocity, it must also be proportional to volumetric flow rate, since volumetric flow rate is also proportional to average fluid velocity⁴¹. Thus, what we have here is a type of flowmeter based on electromagnetic induction. These flowmeters are commonly known as *magnetic flowmeters* or simply *magflow meters*.

We may state the relationship between volumetric flow rate (Q) and motional EMF (\mathcal{E}) more precisely by algebraic substitution. First, we will write the formula relating volumetric flow to average velocity, and then manipulate it to solve for average velocity:

$$Q = A\bar{v}$$

$$\frac{Q}{A} = \bar{v}$$

Next, we re-state the motional EMF equation, and then substitute $\frac{Q}{A}$ for \bar{v} to arrive at an equation relating motional EMF to volumetric flow rate (Q), magnetic flux density (B), pipe diameter (d), and pipe area (A):

$$\mathcal{E} = Bd\bar{v}$$

$$\mathcal{E} = Bd\frac{Q}{A}$$

$$\mathcal{E} = \frac{BdQ}{A}$$

⁴¹This is an application of the transitive property in mathematics: if two quantities are both equal to a common third quantity, they must also be equal to each other. This property applies to proportionalities as well as equalities: if two quantities are proportional to a common third quantity, they must also be proportional to each other.

Since we know this is a circular pipe, we know that area and diameter are directly related to each other by the formula $A = \frac{\pi d^2}{4}$. Thus, we may substitute this definition for area into the last equation, to arrive at a formula with one less variable (only d , instead of both d and A):

$$\begin{aligned}\mathcal{E} &= \frac{BdQ}{\frac{\pi d^2}{4}} \\ \mathcal{E} &= \frac{BdQ}{1} \frac{4}{\pi d^2} \\ \mathcal{E} &= \frac{4BQ}{\pi d}\end{aligned}$$

If we wish to have a formula defining flow rate Q in terms of motional EMF (\mathcal{E}), we may simply manipulate the last equation to solve for Q :

$$Q = \frac{\pi d \mathcal{E}}{4B}$$

This formula will successfully predict flow rate only for absolutely perfect circumstances. In order to compensate for inevitable imperfections, a “proportionality constant” (k) is usually included in the formula⁴²:

$$Q = k \frac{\pi d \mathcal{E}}{4B}$$

Where,

- Q = Volumetric flow rate (cubic meters per second)
- \mathcal{E} = Motional EMF (volts)
- B = Magnetic flux density (Tesla)
- d = Diameter of flowtube (meters)
- k = Constant of proportionality

Note the linearity of this equation. Nowhere do we encounter a power, root, or other non-linear mathematical function in the equation for a magnetic flowmeter. This means no special characterization is required to calculate volumetric flow rate.

A few conditions must be met for this formula to successfully infer volumetric flow rate from induced voltage:

- The liquid must be a reasonably good conductor of electricity (*note: it is okay if the conducting fluid contains some non-conducting solids; the conductive fluid surrounding the non-conducting solid matter still provides electrical continuity between the electrodes necessary for induction*)
- The pipe must be completely filled with liquid to ensure contact with both probes as well as to ensure flow across the entire cross-section of the pipe
- The flowtube must be properly grounded to avoid errors caused by stray electric currents in the liquid

⁴²The colloquial term in the United States for this sort of thing is *fudge factor*.

The first condition is met by careful consideration of the process liquid prior to installation. Magnetic flowmeter manufacturers will specify the minimum conductivity value of the liquid to be measured. The second and third conditions are met by correct installation of the magnetic flowtube in the pipe. The installation must be done in such a way as to guarantee full flooding of the flowtube (no gas pockets). The flowtube is usually installed with electrodes across from each other horizontally (never vertically!) so even a momentary gas bubble will not break electrical contact between an electrode tip and the liquid flowstream. The following photograph shows how *not* to install magnetic flowmeters:



Note in this example how the electrodes are vertically oriented instead of horizontal, because the pipes for these two magnetic flowmeters were placed too close⁴³ to allow proper clearance for the protruding electrodes to lie horizontally. Sadly, poor flowmeter installation is all too common in new projects, as many piping designers and pipefitters are ignorant of flowmeter operating principles. This is one way instrument engineers and technicians may deter operational problems: by involving themselves in the design phase of a piping system, and helping to educate piping designers.

Magnetic flowmeters exhibit several advantages over other types of flowmeter. They are fairly tolerant of swirl and other large-scale turbulent fluid behavior, because the induced voltage is proportional only to the *perpendicular* velocity of the conductor, in this case the velocity of the fluid along the centerline of the flowtube. As such, magnetic flowmeters do not require the long straight-runs of pipe upstream and downstream that orifice plates do, which is a great advantage in

⁴³The obvious solution to this problem – relocating the pipes to give more clearance between flowmeters – would be quite expensive given the large pipe sizes involved. A “compromise” solution is to tilt the magnetic flowtubes as far as possible without the electrodes touching the adjacent flowtube. Horizontal electrode installation is ideal for horizontal pipes, but an angled installation will be better than a vertical installation.

many piping systems. Upstream straight-pipe requirements of 5 diameters and downstream straight-pipe requirements of 3 diameters is typical⁴⁴.

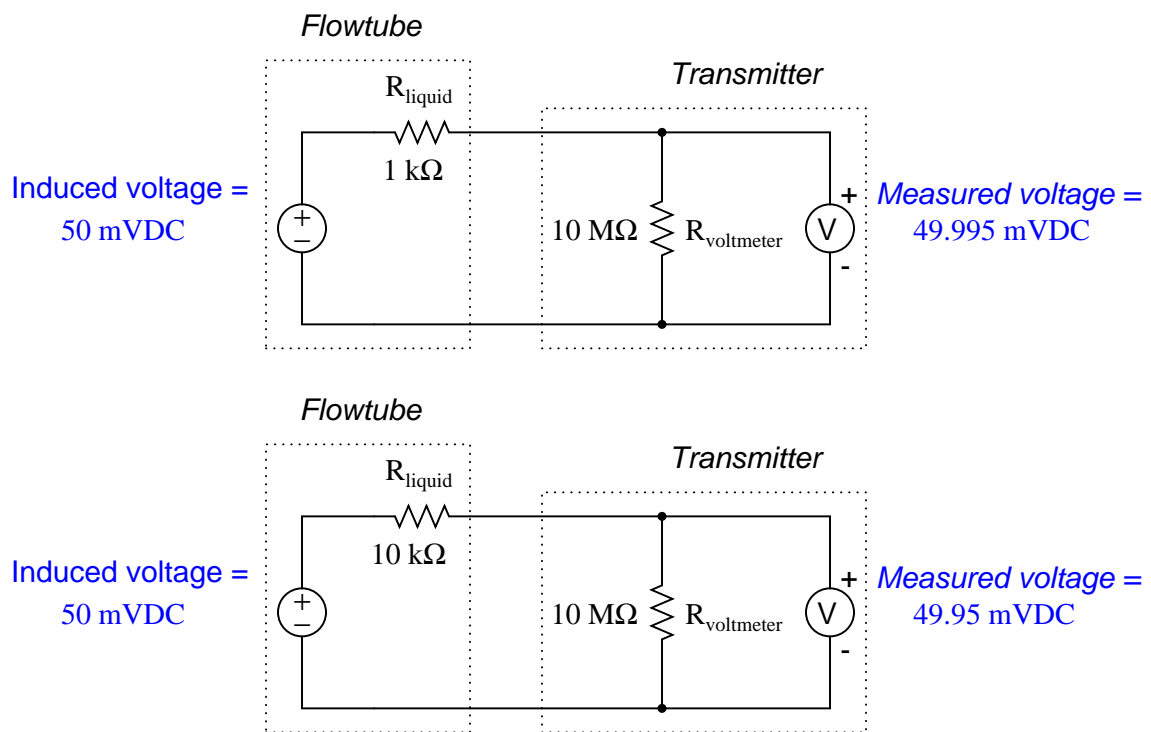
Additionally, the wide-open bore of a magnetic flowmeter's tube means there is absolutely nothing to restrict the flow, resulting in extremely low permanent pressure loss. The lack of any obstruction within the path of fluid flow means magnetic flowmeters are quite tolerant of solids⁴⁵ within the liquid flowstream, making them well-suited for measuring such process liquids as wastewater, slurries, wood pulp, and food products which might clog other types of flowmeters. In fact, magnetic flowmeters are the dominant flowmeter technology used in wastewater, wood pulping, and food processing industries for this very reason.

⁴⁴As always, check the manufacturer's literature for specific requirements, as variations do exist for different models and sizes of magtube.

⁴⁵Even electrically *non-conducting* solid matter is tolerated well by magnetic flowmeters, since the conducting liquid surrounding the solids still provides continuity from one electrode to the other.

Electrical conductivity of the process liquid must meet a certain minimum value, but that is all. It is surprising to some technicians that changes in liquid conductivity have little to no effect on flow measurement accuracy. It is not as though a doubling of liquid conductivity will result in a doubling of induced voltage! Motional EMF is strictly a function of physical dimensions, magnetic field strength, and fluid velocity.

Liquids with poor conductivity present a greater electrical resistance in the voltage-measuring circuit than liquids with good conductivity, but this is of little consequence because the input impedance of the detection circuitry is phenomenally high. The effect of liquid conductivity on flowmeter operation may be modeled by the following DC circuits:



Here, a ten-fold (one order of magnitude) change in liquid resistance barely affects the measured voltage (49.995 mV versus 49.95 mV) because the flow transmitter's voltage-sensing electronic circuit has such a high input impedance. The liquid's equivalent resistance value must increase dramatically beyond the values shown in this example before it will have any significant effect on flow measurement accuracy.

In fact, the only time fluid conductivity is a problem with magnetic flowmeters is when the fluid in question has negligible conductivity. Such fluids include deionized water (e.g. steam boiler feedwater, ultrapure water for pharmaceutical and semiconductor manufacturing) and oils. Most aqueous (water-based) fluids work fine with magnetic flowmeters.

Proper grounding of the flowtube is very important for magnetic flowmeters. The motional EMF generated by most liquid flowstreams is very weak (1 millivolt or less!), and therefore may be easily overshadowed by noise voltage present as a result of stray electric currents in the piping and/or liquid. To combat this problem, magnetic flowmeters are usually equipped with grounding conductors placed to shunt (bypass) stray electric currents around the flowtube so the only voltage intercepted by the electrodes will be the motional EMF produced by liquid flow, and not voltage drops created by stray currents through the resistance of the liquid. The following photograph shows a Rosemount model 8700 magnetic flowtube, with two braided⁴⁶-wire grounding straps clearly visible:




Note how both grounding straps attach to a common junction point on the flowtube housing. This common junction point should also be bonded to a functional earth ground when the flowtube is installed in the process line. On this particular flowtube you can see a stainless steel *grounding ring* on the face of the near flange, connected to one of the braided grounding straps. An identical grounding ring lays on the other flange, but it is not clearly visible in this photograph. These rings

⁴⁶Braided conductors do a better job of shunting radio-frequency currents, because at very high frequencies the *skin effect* makes the surface area of a conductor a greater factor in its conductivity than its cross-sectional area.

provide points of electrical contact with the liquid in installations where the pipe is made of plastic, or where the pipe is metal but lined with a plastic material for corrosion resistance.


If the pipe connecting to a magnetic flowmeter's flowtube is conductive (e.g. metal), grounding may be accomplished by joining the metal pipes' flanges together with grounding straps to a common grounding point on the flowtube body as such:

Stray current(s) shunted around flowtube




If the pipe connecting to a magnetic flowmeter's flowtube is non-conductive (e.g. plastic) or conductive with an insulating lining (e.g. metal pipe with plastic lining), grounding to the pipe flanges will be pointless. In order for flowtube grounding to be effective, the grounding conductors must have electrical continuity to the fluid itself. Special *grounding rings* may be sandwiched between the flanges of non-conducting pipes to provide points of electrical contact with the fluid. These grounding rings are then joined together with grounding straps to a common grounding point on the flowtube body as such:


Stray current(s) shunted around flowtube through grounding rings between flanges



Some magnetic flowmeters have their signal conditioning electronics located integral to the flowtube assembly. A couple of examples are shown here (a pair of small Endress+Hauser flowmeters on the left and a large Toshiba flowmeter on the right):



Other magnetic flowmeters have separate electronics and flowtube assemblies, connected together by shielded cable. In these installations, the electronics assembly is referred to as the flow transmitter (FT) and the flowtube as the flow element (FE):



This next photograph shows an enormous (36 inch diameter!) magnetic flow element (black) and flow transmitter (blue, behind the person's hand shown for scale) used to measure wastewater flow at a municipal sewage treatment plant:




Note the vertical pipe orientation, ensuring constant contact between the electrodes and the water during flowing conditions.

While in theory a permanent magnet should be able to provide the necessary magnetic flux for a magnetic flowmeter to function, this is never done in industrial practice. The reason for this has to do with a phenomenon called *polarization* which occurs when a DC voltage is impressed across a liquid containing ions (electrically charged molecules). Electrically-charged molecules (ions) tend to collect near poles of opposite charge, which in this case would be the flowmeter electrodes. This “polarization” would soon interfere with detection of the motional EMF if a magnetic flowmeter were to use a constant magnetic flux such as that produced by a permanent magnet. A simple solution to this problem is to alternate the polarity of the magnetic field, so the motional EMF polarity also alternates and never gives the fluid ions enough time to polarize.

This is why magnetic flowmeter tubes always employ electromagnet *coils* to generate the magnetic flux instead of permanent magnets. The electronics package of the flowmeter energizes these coils with currents of alternating polarity, so as to alternate the polarity of the induced voltage across the moving fluid. Permanent magnets, with their unchanging magnetic polarities, would only be able to create an induced voltage with constant polarity, leading to ionic polarization and subsequent flow measurement errors.

A photograph of a Foxboro magnetic flowtube with one of the protective covers removed shows these wire coils clearly (in blue):



Perhaps the simplest form of coil excitation is when the coil is energized by 60 Hz AC power taken from the line power source, such as the case with this Foxboro flowtube. Since motional EMF is proportional to fluid velocity and to the flux density of the magnetic field, the induced voltage for such a coil will also be a 60 Hz sine wave whose amplitude varies⁴⁷ with volumetric flow rate.

Unfortunately, if there is any stray electric current traveling through the liquid to produce erroneous voltage drops between the electrodes, chances are it will be 60 Hz AC as well⁴⁸. With the coil energized by 60 Hz AC, any such noise voltage may be falsely interpreted as fluid flow because the sensor electronics has no way to distinguish between 60 Hz noise in the fluid and a 60 Hz motional EMF caused by fluid flow.

A more sophisticated solution to this problem uses a *pulsed* excitation power source for the flowtube coils. This is called *DC* excitation by magnetic flowmeter manufacturers, which is a bit misleading because these “DC” excitation signals often reverse polarity, appearing more like an AC square wave on an oscilloscope display. The motional EMF for one of these flowmeters will exhibit the same waveshape, with amplitude once again being the indicator of volumetric flow rate. The sensor electronics can more easily reject any AC noise voltage because the frequency and waveshape of the noise (60 Hz, sinusoidal) will not match that of the flow-induced motional EMF signal.

The most significant disadvantage of pulsed-DC magnetic flowmeters is slower response time to changing flow rates. In an effort to achieve a “best-of-both-worlds” result, some magnetic flowmeter manufacturers produce *dual-frequency* flowmeters which energize their flowtube coils with two mixed frequencies: one below 60 Hz and one above 60 Hz. The resulting voltage signal intercepted by the electrodes is demodulated and interpreted as a flow rate.

⁴⁷For example, in a condition of no liquid flow through the tube, the electrodes will intercept no voltage at all when the magnetic excitation is 60 Hz AC. When liquid moves slowly in the forward direction through the tube, a low-amplitude 60 Hz millivoltage signal will be detected at the electrodes. When liquid moves rapidly in the forward direction through the tube, the induced 60 Hz AC millivoltage will be greater in amplitude. Any liquid motion in the reverse direction induces a proportional 60 Hz AC voltage signal whose phase is 180° shifted from the excitation signal driving the magnetic coils of the flowtube.

⁴⁸We know this because the largest electrical noise sources in industry are electric motors, transformers, and other power devices operating on the exact same frequency (60 Hz in the United States, 50 Hz in Europe) as the flowtube coils.


22.4.4 Ultrasonic flowmeters

Ultrasonic flowmeters measure fluid velocity by passing high-frequency sound waves along the fluid flow path. Fluid motion influences the propagation of these sound waves, which may then be measured to infer fluid velocity. Two major sub-types of ultrasonic flowmeters exist: *Doppler* and *transit-time*. Both types of ultrasonic flowmeter work by transmitting a high-frequency sound wave into the fluid stream (the *incident* pulse) and analyzing the received pulse.

Doppler flowmeters exploit the *Doppler effect*, which is the shifting of frequency resulting from waves emitted by or reflected by a moving object. A common realization of the Doppler effect is the perceived shift in frequency of a horn's report from a moving vehicle: as the vehicle approaches the listener, the pitch of the horn seems higher than normal; when the vehicle passes the listener and begins to move away, the horn's pitch appears to suddenly "shift down" to a lower frequency. In reality, the horn's frequency never changes, but the velocity of the approaching vehicle relative to the stationary listener acts to "compress" the sonic vibrations in the air. When the vehicle moves away, the sound waves are "stretched" from the perspective of the listener.

The same effect takes place if a sound wave is aimed at a moving object, and the echo's frequency is compared to the transmitted (incident) frequency. If the reflected wave returns from a bubble advancing toward the ultrasonic transducer⁴⁹, the reflected frequency will be greater than the incident frequency. If the flow reverses direction and the reflected wave returns from a bubble traveling away from the transducer, the reflected frequency will be less than the incident frequency. This matches the phenomenon of a vehicle's horn pitch seemingly increasing as the vehicle approaches a listener and seemingly decreasing as the vehicle moves away from a listener.

A Doppler flowmeter bounces sound waves off of bubbles or particulate material in the flow stream, measuring the frequency shift and inferring fluid velocity from the magnitude of that shift.



The requirement for there to be objects in the flow stream large enough to reflect sound waves limits Doppler ultrasonic flowmeters to liquid applications. "Dirty" liquids such as slurries and wastewater, or liquids carrying a substantial number of gas bubbles (e.g. carbonated beverages) are good candidate fluids for this technology. It is unrealistic to expect that any gas stream will be carrying liquid droplets or solid matter large enough to reflect strong echoes, and so Doppler flowmeters cannot be used to measure gas flow.

⁴⁹In the industrial instrumentation world, the word "transducer" usually has a very specific meaning: a device used to process or convert standardized instrumentation signals, such as 4-20 mA converted into 3-15 PSI, etc. In the general scientific world, however, the word "transducer" describes any device converting one form of energy into another. It is this latter definition of the word that I am using when I describe an ultrasonic "transducer" – a device used to convert electrical energy into ultrasonic sound waves, and vice-versa.

The mathematical relationship between fluid velocity (v) and the Doppler frequency shift (Δf) is as follows, for fluid velocities much less than the speed of sound through that fluid ($v \ll c$):

$$\Delta f = \frac{2vf \cos \theta}{c}$$

Where,

Δf = Doppler frequency shift

v = Velocity of fluid (actually, of the particle reflecting the sound wave)

f = Frequency of incident sound wave

θ = Angle between transducer and pipe centerlines

c = Speed of sound in the process fluid

Note how the Doppler effect yields a direct measurement of fluid *velocity* from each echo received by the transducer. This stands in marked contrast to measurements of *distance* based on time-of-flight (time domain reflectometry – where the amount of *time* between the incident pulse and the returned echo is proportional to distance between the transducer and the reflecting surface), such as in the application of ultrasonic liquid level measurement. In a Doppler flowmeter, the time delay between the incident and reflected pulses is irrelevant. Only the *frequency shift* between the incident and reflected signals matters. This frequency shift is also directly proportional to the velocity of flow, making the Doppler ultrasonic flowmeter a linear measurement device.

Re-arranging the Doppler frequency shift equation to solve for velocity (again, assuming $v \ll c$),

$$v = \frac{c\Delta f}{2f \cos \theta}$$

Knowing that volumetric flow rate is equal to the product of pipe area and the average velocity of the fluid ($Q = A\bar{v}$), we may re-write the equation to directly solve for calculated flow rate (Q):

$$Q = \frac{Ac\Delta f}{2f \cos \theta}$$

A very important consideration for Doppler ultrasonic flow measurement is that the calibration of the flowmeter varies with the speed of sound through the fluid (c). This is readily apparent by the presence of c in the above equation: as c increases, Δf must proportionately decrease for any fixed volumetric flow rate Q . Since the flowmeter is designed to directly interpret flow rate in terms of Δf , an increase in c causing a decrease in Δf will thus register as a decrease in Q . This means the speed of sound for a fluid must be precisely known in order for a Doppler ultrasonic flowmeter to accurately measure flow.

The speed of sound through any fluid is a function of that medium's density and bulk modulus (how easily it compresses):

$$c = \sqrt{\frac{B}{\rho}}$$

Where,

c = speed of sound in a material (meters per second)


B = Bulk modulus (pascals, or newtons per square meter)

ρ = Mass density of fluid (kilograms per cubic meter)

Temperature affects liquid density, and composition (the chemical constituency of the liquid) affects bulk modulus. Thus, temperature and composition both are influencing factors for Doppler ultrasonic flowmeter calibration. Pressure is not a concern here, since pressure only affects the density of gases, and we already know Doppler flowmeters only function with liquids.

Following on the theme of requiring bubbles or particles of sufficient size, another limitation of Doppler ultrasonic flowmeters is their inability to measure flow rates of liquids that are too clean and too homogeneous. In such applications, the sound-wave reflections will be too weak to reliably measure. Such is also the case when the solid particles have a speed of sound too close to the that of the liquid, since reflection happens only when a sound wave encounters a material with a markedly different speed of sound. Doppler-type ultrasonic flowmeters are useless in applications where we cannot obtain strong sound-wave reflections.

Transit-time flowmeters, sometimes called *counterpropagation* flowmeters, are an alternative to Doppler ultrasonic flowmeters. A transit-time ultrasonic flowmeter uses a pair of opposed sensors to measure the time difference between a sound pulse traveling with the fluid flow versus a sound pulse traveling against the fluid flow. Since the motion of fluid tends to carry a sound wave along, the sound pulse transmitted downstream will make the journey faster than a sound pulse transmitted upstream⁵⁰:



The rate of volumetric flow through a transit-time flowmeter is a simple function of the upstream and downstream propagation times:

$$Q = k \frac{t_{up} - t_{down}}{(t_{up})(t_{down})}$$

Where,

Q = Calculated volumetric flow rate

k = Constant of proportionality

t_{up} = Time for sound pulse to travel from downstream location to upstream location (upstream, against the flow)

t_{down} = Time for sound pulse to travel from upstream location to downstream location (downstream, with the flow)

An interesting characteristic of transit-time velocity measurement is that the ratio of transit time difference over transit time product remains constant with changes in the speed of sound through the fluid⁵¹. When this equation is cast into terms of path length (L), fluid velocity (v), and sound velocity (c), the equation simplifies to $Q = \frac{2kv}{L}$, proving that the transit-time flowmeter is linear just like the Doppler flowmeter, with the advantage of being immune to changes in the fluid's speed of sound. Changes in bulk modulus resulting from changes in fluid composition, or changes in density resulting from compositional, temperature, or pressure variations therefore have little effect on a transit-time flowmeter's accuracy.

⁵⁰This phenomenon is analogous to paddling a canoe across the width of a river, with the canoe bow angled upstream versus angled downstream. Angled upstream, the canoeist must overcome the velocity of the river and therefore takes longer to reach the other side. Angled downstream, the river's velocity aids the canoeist's efforts and therefore the trip takes less time.

⁵¹If you would like to prove this to yourself, you may do so by substituting path length (L), fluid velocity (v), and sound velocity (c) for the times in the flow formula. Use $t_{up} = \frac{L}{c-v}$ and $t_{down} = \frac{L}{c+v}$ as your substitutions, then algebraically reduce the flow formula until you find that all the c terms cancel. Your final result should be $Q = \frac{2kv}{L}$.

Not only are transit-time ultrasonic flowmeters immune to changes in the speed of sound, but they are also able to measure that sonic velocity independent of the flow rate. The equation for calculating speed of sound based on upstream and downstream propagation times is as follows:

$$c = \frac{L}{2} \left(\frac{t_{up} + t_{down}}{(t_{up})(t_{down})} \right)$$

Where,

c = Calculated speed of sound in fluid

L = Path length

t_{up} = Time for sound pulse to travel from downstream location to upstream location (upstream, against the flow)


t_{down} = Time for sound pulse to travel from upstream location to downstream location (downstream, with the flow)

While not necessary or even particularly relevant for the direct purpose of flow measurement, this inference of the fluid's speed of sound is nevertheless useful as a diagnostic tool. If the true speed of sound for the fluid is known either by direct laboratory measurement of a sample or by chemical analysis⁵² of a sample, this speed may be compared against the flowmeter's reported speed of sound to check the flowmeter's absolute transit time measurement accuracy. Certain problems within the sensors or within the sensor electronics may be detected in this way.

A requirement for reliable operation of a transit-time ultrasonic flowmeter is that the process fluid be free from gas bubbles or solid particles which might scatter or obstruct the sound waves. Note that this is precisely the opposite requirement of Doppler ultrasonic flowmeters, which *require* bubbles or particles to reflect sound waves. These opposing requirements neatly distinguish applications suitable for transit-time flowmeters from applications suitable for Doppler flowmeters, and also raise the possibility of using transit-time ultrasonic flowmeters on gas flowstreams as well as on liquid flowstreams.


⁵²An instrument called a *gas chromatograph* is able to provide live measurement of gas composition, with a computer calculating the average speed of sound for the gas given the known types and percentages of each molecular compound comprising the gas mixture. It just so happens that gas composition analysis by chromatograph is something typically done for custody transfer flow measurement of natural gas anyway, for the primary purpose of calculating the gas's *heating value* as a fuel, and therefore no additional investment of instrumentation is necessary to calculate the gas's speed of sound in this application.

One potential problem with any ultrasonic flowmeter is being able to measure the true average fluid velocity when the flow profile changes with Reynolds number. If just one ultrasonic “beam” is used to probe the fluid velocity, the path this beam takes will likely see a different velocity profile as the flow rate changes (and the Reynolds number changes along with it). Recall the difference in fluid velocity profiles between low Reynolds number flows (left) and high Reynolds number flows (right):



A popular way to mitigate this problem is to use multiple sensor pairs, sending acoustic signals along multiple paths through the fluid (i.e. a *multipath* ultrasonic flowmeter), and to average the resulting velocity measurements. Dual-beam transit-time flowmeters have been in use for well over a decade at the time of this writing (2009), and one manufacturer even has a *five beam* ultrasonic flowmeter model which they claim maintains an accuracy of $\pm 0.15\%$ through the laminar-to-turbulent flow regime transition⁵³. A simplified illustration of a Daniel four-beam (or four “chord”) ultrasonic flowmeter is shown here:

Daniel 4-chord ultrasonic flowmeter



Multipath ultrasonic flowmeters, by virtue of measuring more than one sound wave path through the fluid, also have the ability to detect irregular flow profiles. Each sonic path between sensor pairs in a transit-time ultrasonic flowmeter, called a *chord*, measures flow velocity. The velocities reported for each chord may be compared to the flowmeter’s calculated average flow velocity, and expressed as *velocity ratios*. A particular chord measuring a velocity greater than the flowmeter’s average will report a velocity ratio greater than one (> 1), whereas a chord measuring a velocity less than the meter’s average will report a velocity ratio less than one (< 1). In the Daniel four-chord ultrasonic flowmeter, two chords (B and C) measure velocity near the center of the pipe while the others (A and D) measure velocity closer to the pipe walls. In normal operation, the center chord velocity

⁵³See page 10 of Friedrich Hofmann’s *Fundamentals of Ultrasonic Flow Measurement for industrial applications* paper.

ratios should exceed the outer chord velocity ratios by a small amount, since the flow profiles of laminar and turbulent flow regimes alike exhibit a greater velocity at the center of a pipe than near the walls of a pipe.

Another parameter called *profile factor* expresses the chord velocity factors as a ratio, inner velocity factors over outer velocity factors: $\frac{B+C}{A+D}$, a number which should always exceed one (> 1). The exact value of this profile factor varies with the flowmeter's installation, and may shift over time if the flow profile shifts for some reason (e.g. partial blockage in a flow conditioner, piping change, accumulation of debris on pipe wall). For example, a pipe accumulating dirt or other solid material on its walls over time will tend to slow down the velocity of fluid near the walls compared to the velocity of fluid in the center. This has the effect of decreasing chord A and D velocity ratios while increasing chord B and C velocity ratios, which in turn increases the profile factor value. In this regard, the profile factor value for an operating flowmeter may be an important diagnostic tool, indicating some physical abnormality within the piping.

As previously mentioned, it is possible for a transit-time ultrasonic flowmeter to measure the speed of sound through the fluid from the sum of the upstream and downstream transit times. Multipath transit-time flowmeters may use this feature to cross-check calculated speed of sound values measured by each chord, as a self-diagnostic tool. Since each chord measures the same gas composition, each chord's calculated speed of sound should be exactly equal. Significant differences in calculated speed of sound between chords suggests a failure in the acoustic sensor(s) or electronics for the errant chord.

Some modern ultrasonic flowmeters have the ability to switch back and forth between Doppler and transit-time (counterpropagation) modes, automatically adapting to the fluid being sensed. This capability enhances the suitability of ultrasonic flowmeters to a wider range of process applications.

Ultrasonic flowmeters are adversely affected by swirl and other large-scale fluid disturbances, and as such may require substantial lengths of straight pipe upstream and downstream of the measurement flowtube to stabilize the flow profile.

Like magnetic flowmeters, ultrasonic flowmeters are completely non-obstructive, which means they exhibit extremely low permanent pressure loss and will not accumulate debris.

Advances in ultrasonic flow measurement technology have reached a point where it is now feasible to consider ultrasonic flowmeters for custody transfer measurement of natural gas. The American Gas Association has released a report specifying the use of multipath transit-time ultrasonic flowmeters in this capacity (Report #9). As with the AGA's #3 (orifice plate) and #7 (turbine) high-accuracy gas flow measurement standards, the AGA9 standard requires the addition of pressure and temperature instruments on the gas line to measure gas pressure and temperature in order to calculate flow either in units of mass or in units of standardized volume (e.g. SCFM). The measurement of temperature and pressure for a transit-time ultrasonic flowmeter has nothing to do with correcting errors within the meter itself, since we know transit-time flowmeters are inherently immune to changes in gas density or composition. Temperature and pressure measurements are necessary for custody transfer applications simply because ultrasonic flowmeters, like turbine flowmeters, measure only *volumetric* flow. The fair sale and purchase of a gas requires measurement of *molecular quantity*, not just volume, which is why a flow computer requires measurements of pressure and temperature in order to convert the ultrasonic flowmeter's volumetric output to either mass flow or standardized volumetric flow.


A unique advantage of ultrasonic flow measurement is the ability to measure flow through the use of temporary *clamp-on* sensors rather than a specialized flowtube with built-in ultrasonic transducers. While clamp-on sensors are not without their share of problems⁵⁴, they constitute an excellent solution for certain flow measurement applications.

An important criterion for successful application of a clamp-on flowmeter is that the pipe material be homogeneous in nature, to efficiently conduct sound waves between the process fluid and the clamp-on transducers. Porous pipe materials such as clay and concrete are therefore unsuitable for clamp-on ultrasonic flow measurement.

⁵⁴Most notably, the problem of achieving good acoustic coupling with the pipe wall so signal transmission to the fluid and signal reception back to the sensor may be optimized. Also, there is the potential for sound waves to “ring around the pipe” instead of travel through the fluid with clamp-on ultrasonic flowmeters because the sound waves must travel through the full thickness of the pipe walls in order to enter and exit the fluid stream.

22.4.5 Optical flowmeters

A relatively recent development in industrial flow measurement is the use of light to measure the velocity of a fluid through a pipe. One such technology referred to as *Laser-Two-Focus* (L2F) uses two laser beams to detect the passage of any light-scattering particles carried along by the moving fluid:



Where,

v = Velocity of particle

d = Distance separating laser beams

t = Time difference between sensor pulses


As a particle passes through each laser beam, it redirects the light away from its normal straight-line path in such a way that an optical sensor (one per beam) detects up the scattered light and generates a pulse signal. As that same particle passes through the second beam, the scattered light excites a second optical sensor to generate a corresponding pulse signal. The time delay between two successive pulses is inversely proportional to the velocity of that particle. This technique is analogous to the that used by law-enforcement officers to measure the speed of a vehicle on a highway when

viewed from an aircraft: measure how much time elapses as the vehicle passes between two marks on the road spaced a known distance from each other.

L2F flowmeters of course rely on the continual presence of light-scattering particles within the fluid. These particles could be either liquid droplets or solids within a gas stream, or they could be solid particles or bubbles in a liquid stream.

An alternative to passing laser beams across the entire width of the pipe is to shrink the assembly down to the size of a probe which may be inserted into a pipe. While minimizing installation and maintenance costs, this approach suffers the disadvantage of sensing velocity at only one point within the flow stream, much like a classic Pitot tube. In order to obtain a measurement of bulk (average) fluid velocity, the raw velocity measurement provided by the sensor must be corrected based on the expected Reynolds number for the process fluid, which is why insertion L2F flowmeters are equipped with pressure and temperature transmitters in addition to the optical probe. When the three measurements (pressure, temperature, and velocity at the probe) are combined, the Reynolds number may be calculated which then predicts how “flat” (consistent) the velocity profile is for the fluid stream:


L2F insertion-style flowmeter



For example, if the probe velocity, pressure, temperature, and dimensional parameters indicate a high Reynolds number, it means the velocity profile will be relatively flat, and thus the single-point velocity measurement at the probe will be a fair representation of velocity across the entire width of the pipe. However, if the parameters indicate a low Reynolds number, it means the velocity profile will have a more pronounced “bullet” shape, and therefore the velocities near the pipe walls will be substantially less than the velocity at the center. The flow computer uses this predicted velocity profile to interpret the probe’s single-point velocity measurement and calculate the average velocity of the entire flowstream.

A more sophisticated technique for optical flow measurement relies on the principle of *scintillation*, whereby the fluid itself warps the path of light passing through, rather than entrained particulate matter scattering the light. Scintillation is the same phenomenon responsible for the “twinkling” of stars and city lights viewed from a long distance: as air passes between the light and the observer, pockets of air having different density (due to differences in temperature) and/or sufficient turbulence cause some of the light to be refracted away from an otherwise straight-line path, making it appear as though the light source is randomly vibrating or oscillating. In fact, the phenomenon of scintillation has been used for the measurement of air velocity (anemometry) for many years before its application to industrial flow measurement.


Velocity measurements are inferred by a scintillation flowmeter much the same as they are by an L2F optical flowmeter: measuring the time difference between two sensors’ detection of the same scintillation pattern. Therefore, the scintillation flowmeter applies the same basic formula $v = \frac{d}{t}$ to calculate fluid velocity. A simplified diagram of a scintillation flowmeter is shown here:



Scintillation-style optical flowmeters require a long optical path in order to maximize the angle at which light will be refracted. Thus, they function best when used to measure across the full diameter of a pipe. An interesting feature of this flowmeter technology is that it functions *best* when the flow regime is highly turbulent, since increased fluid turbulence leads to greater scintillation.

Both L2F and scintillation flowmeters have phenomenal turndown ratios. At the low end of the velocity measurement range, an L2F meter is limited by the number of particles entrained in the flow stream and also by the random motions of particles which might be misinterpreted by the flowmeter as bulk motion of the fluid. Scintillation flowmeters are limited at the low end of their measurement range by loss of turbulence, which of course is one of the driving mechanisms of the scintillation effect. However, it should be noted that in both cases the low-velocity measuring limit is quite low in comparison to other types of flowmeters. With no moving parts and using light as the sensing medium, the upper limit for an optical flowmeter can extend into supersonic velocities. This combination of excellent low-velocity and high-velocity sensing yields practical turndown ratios of 1000:1 or more.

Optical flowmeters have been successfully used in one of the more challenging industrial flow measurement applications in existence: *flare gas flow metering*. A “flare” is a continually-ignited burner used as a safety relief point for flammable gases in industries such as petroleum production and refining where a need exists to occasionally vent these gases to the atmosphere:




Ideally, flares operate with no flow of gas to them except during emergency conditions, and then when there is an emergency the flow rate can be very high. In order to accurately measure the low “quiescent” flow of gas to a flare during non-emergency operation (due to pressure relief valve leakage and other inefficiencies) while still being able to register full flow rates during emergencies, the flowmeter must have an extremely large turndown ratio (rangeability).

Gas sent to a flare typically varies widely in composition, temperature, and pressure, especially when multiple processing units share a common flare. The random variability of the gas eliminates many types of flowmeters from consideration. Pressure-based flowmeters such as orifice plates suffer from calibration errors due to density changes, as well as poor turndown ratio. Thermal mass flowmeters suffer from calibration errors due to changes in the specific heat value of the flare gas. Velocity-based gas flowmeters (e.g. turbine, vortex, ultrasonic) are not affected by compositional changes to the extent that pressure-based and thermal flowmeters are, but few exhibit the degree of turndown necessary to accurately measure the full range of flare gas flow. Multi-path transit-time ultrasonic flowmeters show promise in this challenging application, but their relatively high cost (especially at the large pipe diameters typical of flare headers) is a limiting factor.

22.5 Positive displacement flowmeters

A *positive displacement* flowmeter is a cyclic mechanism built to pass a fixed volume of fluid through with every cycle. Every cycle of the meter's mechanism *displaces* a precisely defined (“positive”) quantity of fluid, so that a count of the number of mechanism cycles yields a precise quantity for the total fluid volume passed through the flowmeter. Many positive displacement flowmeters are rotary in nature, meaning each shaft revolution represents a certain volume of fluid has passed through the meter. Some positive displacement flowmeters use pistons, bellows, or expandable bags working on an alternating fill/dump cycle to measure discrete fluid quantities.

Positive displacement flowmeters have been the traditional choice for residential and commercial natural gas flow and water flow measurement in the United States (a simple application of *custody transfer* flow measurement, where the fluid being measured is a commodity bought and sold). The cyclic nature of a positive displacement meter lends itself well to total gas quantity measurement (and not just flow *rate*), as the mechanism may be coupled to a mechanical counter which is read by utility personnel on a monthly basis. A rotary gas flowmeter is shown in the following photograph. Note the odometer-style numerical display on the left-hand end of the meter, totalizing gas usage over time:



Positive displacement flowmeters rely on moving parts to shuttle quantities of fluid through them, and these moving parts must effectively seal against each other to prevent leakage past the mechanism (which will result in the instrument indicating less fluid passing through than there actually is). In fact, the defining characteristic of any positive displacement device is that fluid *cannot* move through without actuating the mechanism, and that the mechanism *cannot* move

without fluid passing through. This stands in contrast to machines such as centrifugal pumps and turbines, where it is possible for the moving part (the impeller or turbine wheel) to jam in place and still have fluid pass through the mechanism. If a positive displacement mechanism jams, fluid flow absolutely halts.

The finely-machined mechanical components of a positive displacement flowmeter will suffer damage from grit or other abrasive materials present in the fluid, which means these flowmeters are applicable only to clean fluid flowstreams. Even with clean fluid flowing through, the sealing surfaces of the mechanisms are subject to wear and accumulating inaccuracies over time. However, there is really nothing more definitive for measuring volumetric flow rate than an instrument built to measure individual volumes of fluid with each mechanical cycle. As one might guess, these instruments are completely immune to swirl and other large-scale fluid turbulence, and may be installed nearly anywhere in a piping system (no need for long sections of straight-length pipe upstream or downstream). Positive displacement flowmeters are also very linear, since mechanism cycles are directly proportional to fluid volume.

A large positive displacement flowmeter used to measure the flow of liquid (registering total accumulated volume in units of gallons) is shown here, having been cut away for use as an instructional display:



The left-hand photograph shows the gear mechanism used to convert rotor motion into a visible total readout. The right-hand photograph shows a close-up of the interlocking rotors (one with three lobes, the other with four slots which those lobes mesh with). Both the lobes and slots are spiral-shaped, such that fluid passing along the spiral pathways must “push” the lobes out of the slots and cause the rotors to rotate. So long as there is no leakage between rotor lobes and slots, rotor turns will have a precise relationship to fluid volume passed through the flowmeter.

22.6 Standardized volumetric flow

The majority of flowmeter technologies operate on the principle of interpreting fluid flow based on the *velocity* of the fluid. Magnetic, ultrasonic, turbine, and vortex flowmeters are prime examples, where the sensing elements (of each meter type) respond directly to fluid velocity. Translating fluid velocity into volumetric flow is quite simple, following this equation:

$$Q = A\bar{v}$$

Where,

Q = Volumetric flow rate (e.g. cubic feet per minute)

A = Cross-sectional area of flowmeter throat (e.g. square feet)

\bar{v} = Average fluid velocity at throat section (e.g. feet per minute)



Positive displacement flowmeters are even more direct than velocity-sensing flowmeters. A positive displacement flowmeter directly measures volumetric flow, counting discrete volumes of fluid as it passes through the meter.

Even pressure-based flowmeters such as orifice plates and venturi tubes are usually calibrated to measure in units of volume over time (e.g. gallons per minute, barrels per hour, cubic feet per second, etc.). For a great many industrial fluid flow applications, measurement in volumetric units makes sense.

This is especially true if the fluid in question is a liquid. Liquids are essentially incompressible: that is, they do not easily yield in volume to applied pressure. This makes volumetric flow measurement relatively simple for liquids: one cubic foot of a liquid at high pressure inside a process vessel will occupy approximately the same volume ($\approx 1 \text{ ft}^3$) when stored in a barrel at atmospheric pressure.

Gases and vapors, however, easily change volume under the influences of both pressure and temperature. In other words, a gas will yield to an increasing pressure by decreasing in volume as the gas molecules are forced closer together, and it will yield to a decreasing temperature by decreasing in volume as the kinetic energy of the individual molecules is reduced. This makes volumetric flow measurement more complex for gases than for liquids. One cubic foot of gas at high pressure and temperature inside a process vessel will *not* occupy one cubic foot under ambient pressure and temperature conditions.

The practical difference between volumetric flow measurement for liquids versus gases is easily seen through an example where we measure the volumetric flow rate before and after a pressure-reduction valve:




The volumetric flow rate of liquid before and after the pressure-reducing valve is the same, since the volume of a liquid does not depend on the pressure applied to it (i.e. liquids are *incompressible*). The volumetric flow rate of the gas, however, is significantly greater following the pressure-reducing valve than before, since the reduction in pressure allows the gas to expand (i.e. less pressure means the gas occupies a greater volume).

What this tells us is that volumetric flow measurement for gas is virtually meaningless without accompanying data on pressure and temperature. A flow rate of “430 ft³/min” reported by a flowmeter measuring gas at 250 PSIG means something completely different than the same volumetric flow rate (430 ft³/min) reported at a different line pressure.

One solution to this problem is to dispense with volumetric flow measurement altogether in favor of *mass flow measurement*, constructing the flowmeter in such a way that the actual *mass* of the gas molecules is measured as they pass through the instrument. This approach is explored in more detail in section 22.7 beginning on page 1708. A more traditional approach to this problem is to specify gas flow in volume units per time, at some agreed-upon (standardized) set of pressure and temperature conditions. This is known as *standardized* volumetric flow measurement.

Referring to our pictorial example previously shown, imagine if we took a sample of the gas flowing at a line pressure of 250 PSIG and let that sample expand to atmospheric pressure (0 PSIG), measuring its new volume under those new conditions. For the sake of keeping this example simple, we will consider only a change in pressure, but not a change in temperature. Obviously, one cubic foot of gas at 250 PSIG would expand to a far greater volume than 1 ft³ at atmospheric pressure. This ratio of “standard volume” to “actual volume” (in the pressurized pipe) could then be used to *scale* the flowmeter’s measurement, so that the flowmeter registers in *standard cubic feet per minute*, or *SCFM*:



Assuming the same temperature inside the pipe as outside, the expansion ratio for this gas will be approximately 18:1, meaning the “actual” flowing rate of 430 ft³/min is equivalent to approximately 7743 ft³/min of gas flow at “standard” atmospheric conditions. In order to unambiguously distinguish “actual” volumetric flow rate from “standardized” volumetric flow rate, we commonly preface each unit with a letter “A” or letter “S” (e.g. ACFM and SCFM).

If a volumetric-registering flowmeter is equipped with pressure and temperature sensors, it may automatically scale its own output signal to measure gas flow rate in standard volumetric units. All we need to determine is the mathematical procedure to scale actual conditions to standard conditions for a gas. To do this, we will refer to the Ideal Gas Law ($PV = nRT$), which is a fair approximation for most real gases at conditions far from their critical phase-change points. First, we shall write two versions of the Ideal Gas Law, one for the gas under “standard” atmospheric conditions, and one for the gas under “actual” flowing conditions (using the subscripts “S” and “A” to distinguish one set of variables from the other):

$$P_S V_S = nRT_S$$

$$P_A V_A = nRT_A$$

What we need to do is determine the *ratio* of standard volume to actual volume ($\frac{V_S}{V_A}$). To do this, we may divide one equation by the other⁵⁵:

$$\frac{P_S V_S}{P_A V_A} = \frac{nRT_S}{nRT_A}$$

Seeing that both the n variables are identical and R is a constant, we may cancel them both from the right-hand side of the equation:

$$\frac{P_S V_S}{P_A V_A} = \frac{T_S}{T_A}$$

Solving for the ratio $\frac{V_S}{V_A}$:

$$\frac{V_S}{V_A} = \frac{P_A T_S}{P_S T_A}$$

Since we know the definition of volumetric flow (Q) is volume over time ($\frac{V}{t}$), we may divide each V variable by t to convert this into a volumetric flow *rate* correction ratio⁵⁶:

$$\frac{\frac{V_S}{t}}{\frac{V_A}{t}} = \frac{P_A T_S}{P_S T_A}$$


This leaves us with a ratio of “standardized” volumetric flow (Q_S) to “actual” volumetric flow (Q_A), for any known pressures and temperatures, standard to actual:

$$\frac{Q_S}{Q_A} = \frac{P_A T_S}{P_S T_A}$$

⁵⁵Recall from algebra that we may perform any arithmetic operation we wish to any equation, so long as we apply that operation equally to both sides of the equation. Dividing one equation by another equation obeys this principle, because both sides of the second equation are equal. In other words, we could divide both sides of the first equation by $P_A V_A$ (although that would not give us the solution we are looking for), but dividing the left side by $P_A V_A$ and the right side by nRT_A is really doing the same thing, since nRT_A is identical in value to $P_A V_A$.

⁵⁶Division by t does not alter the equation at all, since we are essentially multiplying the left-hand side by $\frac{t}{t}$ which is multiplication by 1. This is why we did not have to apply t to the right-hand side of the equation.

We may apply this to a practical example, assuming flowing conditions of 250 PSIG and 90 degrees Fahrenheit, and “standard” conditions of 0 PSIG and 60 degrees Fahrenheit. It is very important to ensure all values for pressure and temperature are expressed in *absolute* units (e.g. PSIA and degrees Rankine), which is what the Ideal Gas Law assumes:




$$\frac{Q_S}{Q_A} = \frac{P_A T_S}{P_S T_A}$$

$$Q_S = Q_A \left(\frac{P_A T_S}{P_S T_A} \right)$$

$$7320 \text{ SCFM} = (430 \text{ ACFM}) \left(\frac{(264.7 \text{ PSIA})(519.67^\circ\text{R})}{(14.7 \text{ PSIA})(549.67^\circ\text{R})} \right)$$

This figure of 7320 SCFM indicates the volumetric flow rate of the gas through the pipe *had it been allowed to expand to atmospheric pressure and cool to ambient temperature*. Although we know these are definitely not the same conditions inside the gas pipe, the correction of actual volumetric flow measurement to these imagined conditions allows us to express gas flow rates in an equitable fashion regardless of the process line pressure or temperature. To phrase this in colloquial terms, standardized volumetric flow figures allow us to compare different process gas flow rates on an “apples to apples” basis, instead of the “apples to oranges” problem we faced earlier where flowmeters would register different volumetric flow values because their line pressures and/or flowing temperatures differed.

With pressure and temperature compensation integrated into volumetric flowmeters, we should be able to measure the exact same *standardized* flow rate at any point in a series gas piping system regardless of pressure or temperature changes:



An unfortunate state of affairs is the existence of multiple⁵⁷ “standard” conditions of pressure and temperature defined by different organizations. In the previous example, 14.7 PSIA was assumed to be the “standard” atmospheric pressure, and 60 degrees Fahrenheit (519.67 degrees Rankine) was assumed to be the “standard” ambient temperature. This conforms to the API (American Petroleum Institute) standards in the United States of America, but it does *not* conform to other standards in America or in Europe. The ASME (American Society of Mechanical Engineers) uses 14.7 PSIA and 68 degrees Fahrenheit (527.67 degrees Rankine) as their “standard” conditions for calculating SCFM. In Europe, the PNEUROP agency has standardized with the American CAGI (Compressed Air and Gas Institute) organization on 14.5 PSIA and 68 degrees Fahrenheit being the “standard” conditions.

For gas flows containing condensable vapors, the *partial pressures* of the vapors must be subtracted from the absolute pressures (P_A and P_S) in order that the correction factor accurately reflects gas behavior alone. A common application of standardized gas flow involving partial pressure correction is found in compressed air systems, where water vapor (relative humidity) is a factor. Here too, standards differ as to the humidity conditions of “standard” cubic feet. The API and PNEUROP/CAGI standards call for 0% humidity (perfectly dry air) as the “standard,” while the ASME defines 35% relative humidity as “standard” for compressed air calculations.

⁵⁷The wonderful thing about standards is that there are so many to choose from!

22.7 True mass flowmeters

Many traditional flowmeter technologies respond to the *volumetric flow rate* of the moving fluid. Velocity-based flowmeters such as magnetic, vortex, turbine, ultrasonic, and optical generate output signals proportional to the speed of fluid molecules and nothing else. This means that if the fluid flowing through one of these flowmeter types were to suddenly become denser (while still flowing by at the same number of volumetric units per minute), the flowmeter's response would not change at all.

The information provided by a volumetric flowmeter may not be what is actually best for the process being measured, however. If the flowmeter in question happens to be measuring the flow rate of feed into a chemical reactor vessel, for example, what we're really concerned with is *how many molecules per unit time* of feed is entering that reactor, not how many cubic meters or how many gallons. We know that changes in temperature will cause gases and liquids alike to change density, which means each volumetric unit will contain a different number of molecules after a temperature change than before. Pressure has a similar influence on gases: increased pressure means more gas molecules occupying each cubic foot (or other volumetric unit), all other factors being equal. If a process requires an accounting of molecular flow rate, a volumetric flowmeter will not provide relevant information.

In steam boiler control systems, the flow rate of water into the boiler and the flow rate of steam coming out of the boiler must be matched in order to maintain a constant quantity of water within the boiler tubes and drums. However, water is a liquid and steam is a vapor, so flow measurements based on volume are meaningless: a cubic foot of steam will *never* contain the same number of molecules as a cubic foot of water. The only reasonable way for the control system to balance both flow rates is to measure them as *mass* flows rather than volumetric flows. No matter what form (phase) the H₂O molecules take, every kilogram going into the boiler must be matched by a kilogram coming out of the boiler in accordance with the Law of Mass Conservation: every H₂O molecule entering the boiler must be matched by one H₂O molecule exiting the boiler in order to maintain an unchanging quantity of H₂O molecules within the boiler. This is why boiler feedwater and steam flowmeters alike are typically calibrated to measure in units of lbm (pounds mass) per unit time.

A similar problem arises in instances where the flowmeter is used for *custody transfer*. This term denotes scenarios where a particular material is being bought and sold, and where accuracy of flow measurement is a matter of monetary importance. Again, in such instances, it is usually the *number of molecules* being bought and sold that really matters, not how many cubic meters or gallons those molecules occupy⁵⁸. Here, as with the chemical reactor feed flow application, a volumetric flowmeter does not provide the most relevant information.

We know from the study of chemistry that all elements have fixed mass values: one *mole*⁵⁹ of any element in monatomic form (single, unbound atoms) will have a mass equal to the atomic mass of that element. For example, one mole of carbon (C) atoms has a mass of 12 grams because the element carbon has an atomic mass of 12. Similarly, one mole of oxygen (O) atoms is guaranteed to have a mass of 16 grams⁶⁰ because 16 is the atomic mass for the element oxygen. Consequently, one

⁵⁸In some applications, such as the custody transfer of natural gas, we are interested in something even more abstract: *heating value*. However, in order to calculate the gross heating value of a fuel gas stream, we must begin with an accurate mass flow measurement – volumetric flow is not really helpful.

⁵⁹A “mole” is equal to a value of 6.022×10^{23} entities. Therefore, one mole of carbon atoms is 602,200,000,000,000,000,000,000 carbon atoms. For a more detailed examination of this subject, refer to section 3.7 beginning on page 262.

⁶⁰I am purposely ignoring the fact that naturally occurring carbon has an average atomic mass of 12.011, and

mole of carbon monoxide (CO) molecules will have a mass of 28 grams (12 + 16), and one mole of carbon dioxide (CO₂) molecules will have a mass of 44 grams (12 + 16×2). These molecule/mass relationships are fixed regardless of how dense or sparse the substances are: one mole of CO₂ will have a mass of 44 grams regardless of pressure or temperature conditions affecting the density of that gas sample. The relationship between molecule count and mass for any given chemical compound is fixed, because mass is an intrinsic property of matter. If our desire is to account for the number of molecules passed through a pipe, and we happen to know the chemical composition of those molecules, measuring the *mass* of the fluid passing through is the most practical way to do it.

The mathematical relationship between volumetric flow (Q) and mass flow (W) is one of proportionality with mass density (ρ):

$$W = \rho Q$$

Dimensional analysis confirms this relationship. Volumetric flow is always measured in volume units (m³, ft³, cc, in³, gallons, etc.) over time, whereas mass flow is always measured in mass units (g, kg, lbm⁶¹, or slugs) over time. To use a specific example, a mass flow rate in pounds (mass) per minute will be obtained by multiplying a mass density in pounds per cubic foot by a volumetric flow rate in cubic feet per minute:

$$\left[\frac{\text{lbm}}{\text{min}} \right] = \left[\frac{\text{lbm}}{\text{ft}^3} \right] \left[\frac{\text{ft}^3}{\text{min}} \right]$$

For example, a volumetric flow rate of 1000 cubic feet per minute of water is equivalent to 62400 pounds (mass) per minute, or 1040 lbm/s, with water having a density of 62.4 lbm/ft³.

With modern sensing and computational technology, it is possible to combine pressure, temperature, and volumetric flow measurements in such a way to *derive* a measurement of mass flow. This is precisely the goal with AGA3 flow measurement (orifice plates), AGA7 flow measurements (turbines), and AGA9 flow measurement (ultrasonic): “compensating” the fundamentally volumetric nature⁶² of these flow-measuring elements with pressure and temperature data to calculate the flow rate in mass units over time.

However, compensated flowmeter systems require much more calibration effort to maintain their long-term accuracy, not to mention a significant capital investment in the multiple transmitters and

naturally occurring oxygen has an atomic mass of 15.9994.

⁶¹The British unit of the “pound” is technically a measure of *force* or *weight* and not *mass*. The proper unit of mass measurement in the British system is the “slug.” However, for better or worse, the “slug” is rarely used, and so engineers have gotten into the habit of using “pound” as a mass measurement. In order to distinguish the use of “pound” to represent mass (an intrinsic property of matter) as opposed to the use of “pound” to represent weight (an incidental property of matter), the former is abbreviated *lbm* (literally, “pounds mass”). In Earth gravity, “lbm” and “lb” are synonymous. However, the standard Newtonian equation relating force, mass, and acceleration ($F = ma$) does not work when “lbm” is the unit used for mass and “lb” is used for force (it does when “slug” is used for mass and “lb” is used for force, though!). A weird unit of force invented to legitimize “pound” as an expression of mass is the *poundal* (“pdl”): one “poundal” of force is the reaction of one “pound” of mass (lbm) accelerated one foot per second squared. By this definition, a one-pound mass (1 lbm) in Earth gravity weighs 32 poundals!

⁶²One could argue that orifice plates and other pressure-based flowmeters respond primarily to mass flow rather than volumetric flow, since their operation is based on the pressure created by *accelerating a mass*. However, fluid density does affect the relationship between mass flow rate and differential pressure (note how the density term ρ appears in the mass flow equation $W = k\sqrt{\rho(P_1 - P_2)}$, where it would not if differential pressure were a strict function of mass flow rate and nothing else), and so the raw output of these instruments must still be “compensated” by pressure and temperature measurements.

flow computer required to gather all the necessary data and perform the mass flow calculations. It would be much simpler if there existed flowmeter technologies naturally responsive to the mass flow rate of a fluid! Fortunately, such flowmeter technologies *do* indeed exist, which is the subject of this section.

For each of the following mass flowmeter technologies, it should be clearly understood that the instrument in question *naturally* responds to mass flow rate. To use our hypothetical example of a fluid stream whose density suddenly increases while the volumetric rate remains constant, a true mass flowmeter will immediately recognize the increase in mass flow (same volume rate, but more mass per unit volume) without the need for additional compensating measurements or computer calculations. True mass flowmeters operate on principles directly related to the mass of the fluid molecules passing through the meter, making them fundamentally different from other flowmeter types.

In the case of the Coriolis flowmeter, the instrument works on the principle of *inertia*: the force generated by an object when it is accelerated or decelerated. This basic property of mass (opposition to change in velocity) forms the basis of the Coriolis flowmeter's function. The inertial force generated inside a Coriolis flowmeter will thus double if the volumetric flowrate of a constant-mass fluid doubles; the inertial force will likewise double if the density of a constant volumetric flow of fluid doubles. Either way, the inertial force becomes a representation of *how fast mass is moving through the flowmeter*, and so the Coriolis flowmeter is a true mass flow instrument.

In the case of the thermal flowmeter, the instrument works on the principle of *convective heat transfer*: heat energy extracted from a hot object as cooler molecules pass by. The ability for fluid molecules to transport heat is a function of the *specific heat* of each molecule and the number of molecules moving past the warmer object. So long as the chemical composition of the fluid remains unchanged, the convective transfer of heat is a function of how many fluid molecules pass by in a given time. The heat transfer rate inside a thermal flowmeter will thus double if the volumetric flowrate of a given fluid doubles and all else remains constant; the heat transfer rate will likewise double if the density of a given fluid doubles and all else remains constant (i.e. twice the number of molecules passing by with each time interval). Either way, the convective heat transfer rate becomes a representation of *how many molecules of fluid are moving through the flowmeter*, which for any given fluid type is proportional to the fluid's mass flow rate. This makes the thermal flowmeter a true mass flow instrument for any (calibrated) fluid composition.

Some older, mechanical technologies⁶³ exist for measuring true mass flow, but these are being supplanted by Coriolis and thermal mass flowmeter technologies. Coriolis and thermal mass flowmeters are also fast becoming the technology of choice for applications formerly the domain of compensated orifice plate (e.g. AGA3) and turbine (e.g. AGA7) flowmeters.

⁶³The impeller-turbine and twin-turbine mass flowmeter types are examples of mechanical true-mass flow technologies. Both work on the principle of fluid inertia. In the case of the impeller-turbine flowmeter, an impeller driven by a constant-speed electric motor imparts a "spin" to a moving fluid, which then impinges on a stationary turbine wheel to generate a measurable torque. The greater the mass flow rate, the greater the impulse force imparted to the turbine wheel. In the twin-turbine mass flowmeter, two rotating turbine wheels with different blade pitches are coupled together by a flexible coupling. As each turbine wheel attempts to spin at its own speed, the inertia of the fluid causes a differential torque to develop between the two wheels. The more mass flow rate, the greater the angular displacement (offset) between the two wheels.

22.7.1 Coriolis flowmeters

Simplified explanation

Coriolis flowmeters represent the state-of-the-art in mass flow measurement at the time of this writing (2010). While incredibly versatile and accurate, their internal operation can be difficult to understand. Put into very simple terms, a Coriolis flowmeter works by shaking one or more tubes carrying the flowing fluid, then precisely measuring the frequency and phase of that shaking. The back-and-forth shaking is driven by an electromagnetic coil, powered by an electronic amplifier circuit to shake the tube(s) at their mechanical resonant frequency. Since this frequency depends on the mass of each tube, and the mass of the tubes depends on the density of the fluid filling the fixed volume of the tubes, the resonant frequency becomes an inverse indication of fluid *density*⁶⁴ regardless of fluid flow through the tubes. As fluid begins to move through the tubes, the inertia of the moving fluid adds another dimension to the tubes' motion: the tubes begin to *undulate*⁶⁵, twisting slightly instead of just shaking back and forth. This twisting motion is directly proportional to the mass flow rate, and is internally measured by comparing the phase shift (θ) between motion at one point on the tube versus another point on the tube: the greater the undulation, the greater the phase shift between these two points' vibrations.

Expressed as proportionalities:

$$\text{Tube frequency} \propto \frac{1}{\text{Density}} \qquad f \propto \frac{1}{\rho}$$

$$\text{Tube twisting} \propto \text{Mass flowrate} \qquad \theta \propto W$$

⁶⁴In fact, this density-measuring function of Coriolis flowmeters is so precise that they often find use *primarily* as density meters, and only secondarily as flowmeters!

⁶⁵An interesting experiment to perform consists of holding a water hose in a U-shape and gently swinging the hose back and forth like a pendulum, then flowing water through that same hose while you continue to swing it. The hose will begin to undulate, its twisting motion becoming visually apparent.


The Coriolis force

In physics, certain types of forces are classified as *fictitious* or *pseudoforces* because they only appear to exist when viewed from an accelerating perspective (called a *non-inertial reference frame*). The feeling you get in your stomach when you accelerate either up or down in an elevator, or when riding a roller-coaster at an amusement park, feels like a force acting against your body when it is really nothing more than the reaction of your body's inertia to being accelerated by the vehicle you are in. The real force is the force of the vehicle against your body, causing it to accelerate. What you perceive is merely a reaction to that force, and not the primary cause of your discomfort as it might appear to be.

Centrifugal force is another example of a "pseudoforce" because although it may appear to be a real force acting on any rotating object, it is in fact nothing more than an inertial reaction. Centrifugal force is a common experience to any child who has ever played on a "merry-go-round:" that perception of a force drawing you away from the center of rotation, toward the rim. The real force acting on any rotating object is toward the center of rotation (a *centripetal* force) which is necessary to make the object radially accelerate toward a center point rather than travel in a straight line as it normally would without any forces acting upon it. When viewed from the perspective of the spinning object, however, it would seem there is a force drawing the object away from the center (a *centrifugal* force).

Yet another example of a “pseudoforce” is the *Coriolis force*, more complicated than centrifugal force, arising from motion perpendicular to the axis of rotation in a non-inertial reference frame. The example of a merry-go-round works to illustrate Coriolis force as well: imagine sitting at the center of a spinning merry-go-round, holding a ball. If you gently toss the ball away from you and watch the trajectory of the ball, you will notice it curve rather than travel away in a straight line. In reality, the ball *is* traveling in a straight line (as viewed from an observer standing on the ground), but from your perspective on the merry-go-round, it appears to be deflected by an invisible force which we call the Coriolis force.

In order to generate a Coriolis force, we must have a mass moving at a velocity perpendicular to an axis of rotation:



The magnitude of this force is predicted by the following vector equation⁶⁶:

$$\vec{F}_c = -2\vec{\omega} \times \vec{v}'m$$

Where,

\vec{F}_c = Coriolis force vector


$\vec{\omega}$ = Angular velocity (rotation) vector

\vec{v}' = Velocity vector as viewed from the rotating reference frame

m = Mass of the object


⁶⁶This is an example of a vector *cross-product* where all three vectors are perpendicular to each other, and the directions follow the right-hand rule.

If we replace the ball with a fluid moving through a tube, and we introduce a rotation vector by tilting that tube around a stationary axis (a fulcrum), a Coriolis force develops on the tube in such a way as to oppose the direction of rotation just like the Coriolis force opposed the direction of rotation of the rotating platform in the previous illustration:



To phrase this in anthropomorphic terms, the fluid “fights” against this rotation because it “wants” to keep traveling in a straight line. For any given rotational velocity, the amount of “fight” will be directly proportional to the product of fluid velocity and fluid mass. In other words, the magnitude of the Coriolis force will be in direct proportion to the fluid’s mass flow rate. This is the basis of a *Coriolis mass flowmeter*.

A demonstration of this Coriolis force may be made by modifying the nozzles on a rotary lawn sprinkler so they point straight out from the center rather than angle in one direction. As water squirts through the now-straight nozzles, they no longer generate a rotational reaction force to spin the nozzle assembly, and so the nozzles remain in place (this much should be obvious). However, if someone were to try rotating the nozzle assembly by hand, they would discover the Coriolis force *opposes* the rotation, acting to keep the nozzle assembly from rotating. The greater the mass flow rate of water through the nozzles, the stronger the inhibiting Coriolis force. Instead of a rotating lawn sprinkler, you are now the proud owner of an *anti-rotating* lawn sprinkler that actually *fights* any attempt to rotate it:




This is a very non-intuitive concept, so it deserves further explanation. The “anti-rotating” sprinkler doesn’t just fail to rotate on its own – it actually *opposes* any attempt to rotate from an external force (e.g. a person trying to push the tubes by hand).

This opposition would not occur if the tubes were merely capped off at the ends and filled with stagnant water. If this were the case, the tubes would simply be heavy with the water's weight, and they would rotate freely about the axis just like any pair of heavy metal tubes would (whether hollow and filled with water, or solid metal). The tubes would have inertia, but they would not *actively* oppose any external effort to rotate.

Having liquid water *move* through the tubes is what makes the difference, and the reason becomes clear once we imagine what each water molecule experiences as it flows from the center (axis of rotation) to the nozzle at the tube tip. Each water molecule originating from the center begins with no lateral velocity, but must *accelerate* as it travels farther along the tube toward the circumference of the tips' rotation where the lateral velocity is at a maximum. The fact that new water molecules are continually making this journey from center to tip means there will always be a *new* set of water molecules requiring acceleration from center velocity (zero) to tip velocity (maximum). In capped tubes filled with stagnant water, the acceleration would only occur in getting the tubes' rotation up to speed – once there, the lateral velocity of each water molecule sitting stagnant inside the tubes would remain the same. However, with water *flowing* from center to tip, this process of acceleration from zero velocity to tip velocity must occur over and over again (continually) for each new water molecule flowing through. This continual acceleration of *new mass* is what generates the Coriolis force, and what actively opposes any external force trying to rotate the “anti-rotating” sprinkler.

As you might guess, it can be difficult to engineer a tubing system capable of spinning in circles while carrying a flowstream of pressurized fluid. To bypass the practical difficulties of building a spinning tube system, Coriolis flowmeters are instead built on the principle of a flexible tube that *oscillates* back and forth, producing the same effect in a cyclic rather than continuous fashion. The effect is not unlike shaking a hose⁶⁷ side to side as it carries a stream of water:



(This illustration is from a vertical view, looking down. The Coriolis force acts laterally, bending the hose to the side.)


The Coriolis force opposes the direction of rotation. The greater the mass flow rate of water through the hose, the stronger the Coriolis force. If we had a way to precisely measure the Coriolis

⁶⁷The Coriolis force generated by a flowing fire hose as firefighters work to point it in a different direction can be quite significant, owing to the high mass flow rate of the water as it flows through the hose and out the nozzle!

force imparted to the hose by the water stream, and to precisely wave the hose so its rotational velocity held constant for every wave, we could directly infer the water's mass flow rate.


Practical Coriolis flowmeter construction

We cannot build a Coriolis flowmeter exactly like the water hose or lawn sprinkler unless we are willing to let the process fluid exit the tubing, so a common Coriolis flowmeter design uses a U-shaped tube that redirects the fluid flow back to the center of rotation. The curved end of the flexible U-tube is forced to shake back and forth by an electromagnetic force coil (like the force coil on an audio speaker) while the tube ends anchor to a stationary manifold:




If fluid inside the tube is stagnant (no flow), the tube will simply vibrate back and forth with the applied force. However, if fluid *flows* through the tube, the moving fluid molecules will experience acceleration as they travel from the anchored base to the tube's rounded end, then experience *deceleration* as they travel back to the anchored base. This continual acceleration and subsequent deceleration of new mass generates a Coriolis force altering the tube's motion.


This Coriolis force causes the U-tube assembly to *twist*. The tube portion carrying fluid from the anchored base to the end tends to *lag* in motion because the fluid molecules in that section of the tube are being accelerated to a greater lateral velocity. The tube portion carrying fluid from the end back to the anchored base tends to *lead* in motion because those molecules are being decelerated back to zero lateral velocity. As mass flow rate through the tube increases, so does the degree of twisting. By monitoring the severity of this twisting motion, we may infer the mass flow rate of the fluid passing through the tube:



In order to reduce the amount of vibration generated by a Coriolis flowmeter, and more importantly to reduce the effect any external vibrations may have on the flowmeter, two identical U-tubes are built next to each other and shaken in complementary fashion (always moving in opposite directions)⁶⁸. Tube twist is measured as *relative* motion from one tube to the next, not as motion between the tube and the stationary housing of the flowmeter. This (ideally) eliminates the effect of any common-mode vibrations on the inferred flow measurement:



Viewed from the end, the complimentary shaking and twisting of the tubes looks like this:




Great care is taken by the manufacturer to ensure the two tubes are as close to identical as possible: not only are their physical characteristics precisely matched, but the fluid flow is split very evenly between the tubes⁶⁹ so their respective Coriolis forces should be identical in magnitude.


⁶⁸For those readers with an automotive bent, this is the same principle applied in opposed-cylinder engines (e.g. Porsche “boxer” air-cooled 6-cylinder engine, Volkswagen air-cooled 4-cylinder engine, BMW air-cooled motorcycle twin engine, Citroen 2CV 2-cylinder engine, Subaru 4- and 6-cylinder opposed engines, etc.). Opposite piston pairs are *always* 180° out of phase for the purpose of maintaining mechanical balance: both moving away from the crankshaft or both moving toward the crankshaft, at any given time.

⁶⁹An alternative to splitting the flow is to plumb the tubes in series so they *must* share the exact same flow rate, like series-connected resistors sharing the exact same amount of electrical current.

A photograph of a Rosemount (Micro-Motion) U-tube Coriolis flowmeter demonstration unit shows the U-shaped tubes (one tube is directly above the other in this picture, so you cannot tell there are actually two U-tubes):



A closer inspection of this flowmeter shows that there are actually two U-tubes, one positioned directly above the other, shaken in complementary directions by a common electromagnetic force coil:



The force coil works on the same principle as an audio speaker: AC electric current passed through a wire coil generates an oscillating magnetic field, which acts against a permanent magnet's field to produce an oscillating force. In the case of an audio speaker, this force causes a lightweight cone to move, which then creates sound waves through the air. In the case of the Coriolis meter assembly, the force pushes and pulls between the two metal tubes, causing them to alternately separate and come together (shake in opposite directions).

Two magnetic displacement sensors monitor the relative motions of the tubes and transmit signals to an electronics module for digital processing. One of those sensor coils may be seen in the previous photograph. Both the force coil and the sensor coil are nothing more than permanent magnets surrounded by movable copper wire coils. The main difference between the force coil and the sensor coil is that the force coil is powered by an AC electric current to impart a vibratory force to the tubes, whereas the sensor coils are both unpowered so they can detect tube motion by generating AC voltage signals to be sensed by the electronics module. The force coil is shown in the left-hand photograph, while one of the two sensor coils appears in the right-hand photograph:




The two AC signals generated by the sensor coils provide data from which the electronics package may interpret fluid density and mass flow rate. The *frequency* of the two coils' signals is inversely related to fluid density, because a denser fluid will cause the tubes to have greater mass and therefore vibrate at a lower frequency⁷⁰. The *phase shift* of the two coils' signals is directly related to mass flow rate, because a greater mass flow rate will cause the tubes to twist to a greater degree, causing the sensors' signals to shift further out of phase with each other.

Advances in sensor technology and signal processing have allowed the construction of Coriolis flowmeters employing straighter tubes than the U-tube unit previously illustrated and photographed. Straighter tubes are advantageous for reasons of reduced plugging potential and the ability to easily drain all liquids out of the flowmeter when needed. In straight-tube Coriolis flowmeters, we still find the same general design of a force coil flanked by matching sensor coils measuring vibration frequency (for density measurement) and phase shift (for mass flow measurement).

⁷⁰The force coil is powered by an electronic amplifier circuit, which receives feedback from the sensor coils. Like any amplifier circuit given positive (regenerative) feedback, it will begin to oscillate at a frequency determined by the feedback network. In this case, the feedback "network" consists of the force coil, tubes, and sensor coils. The tubes, having both resilience and mass, naturally possess their own resonant frequency. This mechanical resonance dominates the feedback characteristic of the amplifier loop, causing the amplifier circuit to oscillate at that same frequency.

Matched tubes and electronics

The tubes inside a Coriolis flowmeter are not just conduits for fluid flow, they are also precision spring elements and volume chambers. As such, it is important to precisely know the spring characteristics and precise dimensions of these tubes so both mass flow and density may be inferred from tube motion. Every Coriolis flow element is factory-tested to determine the flow tubes' mechanical properties, then the electronic transmitter is programmed with these parameters describing the tubes' mechanical properties. The following photograph shows a close-up view of the nameplate on a Rosemount (Micro-Motion) Coriolis mass flowmeter, showing the physical constant values determined for that specific flowtube assembly at the time of manufacture:



This means every Coriolis flowmeter element (the tube and sensor assembly) is unique, with no two identical in behavior. Consequently, the transmitter (the electronics package outputting the process variable signals) must be programmed with values describing the element's behavior, and the complete flowmeter is shipped from the manufacturer as a *matched set*. You cannot interchange elements and transmitters without re-programming the transmitters with the new elements' physical constant values.

Density and temperature measurement

The tubes within a Coriolis flowmeter are shaken at their mechanical resonant frequency to maximize their shaking motion while minimizing electrical power applied to the force coil. The electronics module uses a feedback loop⁷¹ between the sensor coils and the shaker coil to maintain the tubes in a continuous state of resonant oscillation. This resonant frequency changes with process fluid density, since the effective mass of the fluid-filled tubes changes with process fluid density⁷², and mass is one of the variables influencing the mechanical resonant frequency of any elastic structure. Note the “mass” term in the following formula, describing the resonant frequency of a tensed string:

$$f = \frac{1}{2L} \sqrt{\frac{F_T}{\mu}}$$

Where,

f = Fundamental resonant frequency of string (Hertz)

L = String length (meters)

F_T = String tension (newtons)

μ = Unit mass of string (kilograms per meter)

A fluid-filled tube is a close analogue to a tensed string, with tube stiffness analogous to string tension and liquid density analogous to unit mass. So long as the spring constant (tube stiffness) is known, the resonant frequency of the tubes' vibration serves to indicate the unit mass of the tubes, which in turn represents fluid density given the known internal volume of the tubes.

Temperature changes have the potential to interfere with this density measurement, because temperature affects the elasticity of metal (Young's modulus) as well as the tubes' physical dimensions. This is why all Coriolis flowmeters are equipped with RTD temperature sensors to continuously monitor the temperature of the vibrating tubes. The flowmeter's microprocessor takes this tube temperature measurement and uses it to compensate for the resulting elasticity and dimensional changes based on a prior modeling of the tube metal characteristics. In other words, the flowmeter's microprocessor continuously updates the force variable (F_T) representing tube stiffness in the resonant frequency equation so that the frequency will always be a reliable indicator of unit mass (fluid density). This temperature measurement happens to be accessible as an auxiliary output signal, which means a Coriolis flowmeter may double as a (very expensive!) temperature⁷³ transmitter in addition to measuring mass flow rate and fluid density.

The ability to simultaneously measure three process variables (mass flow rate, temperature, and density) makes the Coriolis flowmeter a very versatile instrument indeed. This is especially true


⁷¹This usually takes the form of a simple analog oscillator circuit, using the tubes and sensors as feedback elements. It is not unlike a *crystal* oscillator circuit where the mechanical resonance of a quartz crystal stabilizes the circuit's frequency at one value. The feedback system naturally finds and maintains resonance, just as a crystal oscillator circuit naturally finds and maintains the resonant frequency of the quartz crystal when provided with ample regenerative (positive) feedback. As fluid density inside the tubes changes, the tubes' mass changes accordingly, thus altering the resonant frequency of the system. The analog nature of this mechanical oscillator explains why some early versions of Coriolis flowmeters sometimes required a minor shake or tap to the flowtubes to start their oscillation!

⁷²If you consider each tube as a container with a fixed volume capacity, a change in fluid density (e.g. pounds per cubic foot) must result in a change in mass for each tube.

⁷³An important caveat is that the RTD sensing tube temperature in a Coriolis flowmeter really senses the tubes' *outside skin temperature*, which may not be exactly the same as the temperature of the fluid inside the tube. If the ambient temperature near the flowmeter differs substantially from the fluid's temperature, the tube skin temperature reading may not be accurate enough for the flowmeter to double as a fluid temperature transmitter.

when the flowmeter in question communicates digitally using a “fieldbus” standard rather than an analog 4-20 mA signal. Fieldbus communication allows multiple variables to be transmitted by the device to the host system (and/or to other devices on the same fieldbus network), allowing the Coriolis flowmeter to do the job of three instruments!


An example of a Coriolis mass flowmeter being used as a multi-variable transmitter appears in the following photographs. Note the instrument tag labels in the close-up photograph (FT, TT, and DT), documenting its use as a flow transmitter, temperature transmitter, and density transmitter, respectively:



Proper installation

Although Coriolis flowmeters are immune to fluid turbulence and therefore have no upstream or downstream straight-pipe length requirements, they are still susceptible to other installation-related problems. One of these is vibration: attaching a Coriolis flowmeter a machine that vibrates, or to piping that vibrates from attachment to such a machine, can be problematic because sufficient external vibration may interfere with the resonant vibration of the flowmeter tubes, causing errors in density and/or flow measurement.

A problem unique to bent-tube Coriolis flowmeters is the entrapment of gas bubbles (in a liquid process) or liquid droplets (in a gas process). Either condition will create an uneven distribution of mass inside the flowmeter’s tubes, potentially leading to measurement errors in density and/or flow. The bent tubes of a Coriolis flowmeter should be oriented such that bubbles or droplets cannot collect within them, similar to how a differential pressure sensor should be oriented in relation to an orifice plate: for liquid processes, the bent tubes should be located below the pipe’s centerline; for gas processes, the bent tubes should be located above the pipe’s centerline:



Coriolis flowmeter capabilities and limitations

Even though a Coriolis flowmeter inherently measures *mass* flow rate, the continuous measurement of fluid density allows the meter to calculate *volumetric flow rate* if this is the preferred means of expressing fluid flow. The relationship between mass flow (W), volumetric flow (Q), and mass density (ρ) is quite simple:

$$W = \rho Q \qquad Q = \frac{W}{\rho}$$

All the flowmeter's computer must do to output a volumetric flow measurement is take the mass flow measurement value and divide that by the fluid's measured density. A simple exercise in dimensional analysis (performed with metric units of measurement) validates this concept for both forms of the equation shown above:

$$\left[\frac{\text{kg}}{\text{s}} \right] = \left[\frac{\text{kg}}{\text{m}^3} \right] \left[\frac{\text{m}^3}{\text{s}} \right] \qquad \left[\frac{\text{m}^3}{\text{s}} \right] = \frac{\left[\frac{\text{kg}}{\text{s}} \right]}{\left[\frac{\text{kg}}{\text{m}^3} \right]}$$

Coriolis mass flowmeters are very accurate and dependable. They are also completely immune to swirl and other fluid disturbances, which means they may be located nearly anywhere in a piping system with no need at all for straight-run pipe lengths upstream or downstream of the meter. Their natural ability to measure true mass flow, along with their characteristic linearity and accuracy, makes them ideally suited for custody transfer applications (where the flow of fluid represents product being bought and sold).

The American Gas Association (AGA) formalized the use of Coriolis mass flowmeters for the measurement of natural gas with their Report #11. This standard is equivalent to AGA #3 for orifice meters, AGA #7 for turbine meters, and AGA #9 for ultrasonic meters.

Perhaps the greatest disadvantage of Coriolis flowmeters is their high initial cost, especially for large pipe sizes. Coriolis flowmeters are also more limited in operating temperature than other types of flowmeters and may have difficulty measuring low-density fluids (gases) and mixed-phase⁷⁴ (liquid/vapor) flows. The bent tubes used to sense process flow may also trap process fluid inside to the point where it becomes unacceptable for hygienic (e.g. food processing, pharmaceuticals) applications. Straight-tube Coriolis flowmeter designs, and designs where the angle of the tubes is slight, fare better in this regard than the traditional U-tube Coriolis flowmeter design. However, a disadvantage of straight tubes is that they are stiffer than U-shaped tubes, and so straight-tube Coriolis flowmeters tend to be less sensitive to low flow rates than their U-tube counterparts.

⁷⁴Significant technological progress has been made on mixed-phase Coriolis flow measurement, to the point where this may no longer be a serious consideration in the future.

22.7.2 Thermal flowmeters

Wind chill is a phenomenon common to anyone who has ever lived in a cold environment. When the ambient air temperature is substantially colder than the temperature of your body, heat will transfer from your body to the surrounding air. If there is no breeze to move air past your body, the air molecules immediately surrounding your body will begin to warm up as they absorb heat from your body, which will then decrease the rate of heat loss. However, if there is even a slight breeze of air moving past your body, your body will come into contact with more cool (unheated) air molecules than it would otherwise, causing a greater rate of heat loss. Thus, your perception of the surrounding temperature will be cooler than if there were no breeze.



We may exploit this principle to measure mass flow rate, by placing a heated object in the midst of a fluid flowstream, and measuring how much heat the flowing fluid convects away from the heated object. The “wind chill” experienced by that heated object is a function of true mass flow rate (and not just volumetric flow rate) because the mechanism of heat loss is the rate at which fluid molecules contact the heated object, with each of those molecules having a definite mass.

The simplest form of thermal mass flowmeter is the *hot-wire anemometer*, used to measure air speed. This flowmeter consists of a metal wire through which an electric current passes to heat it up. An electric circuit monitors the resistance of this wire (which is directly proportional to wire temperature because most metals have a definite temperature coefficient of resistance). If air speed past the wire increases, more heat will be drawn away from the wire and cause its temperature to drop. The circuit senses this temperature change and compensates by increasing current through the wire to bring its temperature back up to setpoint. The amount of electrical power required to maintain the hot wire at a constant elevated temperature is a direct function of mass air flow rate past the wire.

Most mass air flow sensors used in automotive engine control applications employ this principle. It is important for engine control computers to measure *mass* air flow and not just volumetric air flow because it is important to maintain proper air/fuel ratio even if the air density changes due to changes in altitude. In other words, the computer needs to know how many air molecules are entering the engine per second in order to properly meter the correct amount of fuel into the engine for complete and efficient combustion. The “hot wire” mass air flow sensor is simple and inexpensive to produce in quantity, which is why it finds common use in automotive applications.

Industrial thermal mass flowmeters usually consist of a specially designed “flowtube” with two temperature sensors inside: one that is heated and one that is unheated. The heated sensor acts as the mass flow sensor (cooling down as flow rate increases) while the unheated sensor serves to compensate for the “ambient” temperature of the process fluid.

A typical thermal mass flowtube appears in the following photographs (note the swirl vanes in the close-up photograph, designed to introduce large-scale turbulence into the flowstream to maximize the convective cooling effect of the fluid against the heated sensor element):



Thermal mass flowmeters lend themselves well to “insertion” style probes, sensing the passage of fluid molecules at one point within the flowstream. An example is shown in the next two photographs, where a thermal mass flowmeter (manufactured by Sage) senses the amount of gas sent to a flare. The insertion probe appears in the left-hand photo (mounted in the vertical flare pipe) while the transmitter head appears in the right-hand photo (located inside of a weather-sheltered building):



The simple construction of thermal mass flowmeters allows them to be manufactured in very small sizes. The following photograph shows a small device that is not only a mass flow meter, but also a mass flow *controller* with its own built-in throttling valve mechanism and control electronics. To give you a sense of scale, the tube fittings seen on the left- and right-hand sides of this device are 1/4 inch, making this photograph nearly full-size:



An important factor in the calibration of a thermal mass flowmeter is the *specific heat* of the process fluid. “Specific heat” is a measure of the amount of heat energy needed to change the temperature of a standard quantity of substance by some specified amount⁷⁵. Some substances have much greater specific heat values than others, meaning those substances have the ability to absorb (or release) a lot of heat energy without experiencing a great temperature change. Fluids with high specific heat values make good *coolants*, because they are able to remove much heat energy from hot objects without experiencing great increases in temperature themselves. Since thermal mass

⁷⁵For example, the specific heat of water is 1.00 kcal / kg · °C, meaning that the addition of 1000 calories of heat energy is required to raise the temperature of 1 kilogram of water by 1 degree Celsius, or that we must remove 1000 calories of heat energy to cool that same quantity of water by 1 degree Celsius. Ethyl alcohol, by contrast, has a specific heat value of only 0.58 kcal / kg · °C, meaning it is almost twice as easy to warm up or cool down as water (little more than half the energy required to heat or cool water needs to be transferred to heat or cool the same mass quantity of ethyl alcohol by the same amount of temperature).

flowmeters work on the principle of convective cooling, this means a fluid having a high specific heat value will elicit a greater response from a thermal mass flowmeter than the exact same mass flow rate of a fluid having a lesser specific heat value (i.e. a fluid that is not as good of a coolant).

This means we must know the specific heat value of whatever fluid we plan to measure with a thermal mass flowmeter, and we must be assured its specific heat value will remain constant. For this reason, thermal mass flowmeters are not suitable for measuring the flow rates of fluid streams whose chemical composition is likely to change over time. This limitation is analogous to that of a pressure sensor used to hydrostatically measure the level of liquid in a vessel: in order for this level-measurement technique to be accurate, we must know the density of the liquid and also be assured that density will be constant over time.

Thermal mass flowmeters are simple and reliable instruments. While not as accurate or tolerant of piping disturbances as Coriolis mass flowmeters, they are far less expensive.


Perhaps the greatest disadvantage of thermal mass flowmeters is their sensitivity to changes in the specific heat of the process fluid. This makes the calibration of any thermal mass flowmeter specific for one composition of fluid only. In some applications such as automotive engine intake air flow, where the fluid composition is constant, this limitation is not a factor. In many industrial applications, however, this limitation is severe enough to prohibit the use of thermal mass flowmeters. Industrial applications for thermal mass flowmeters include natural gas flow measurement (non-custody transfer), and the measurement of purified gas flows (oxygen, hydrogen, nitrogen) where the composition is known to be very stable.

Another (potential) limitation of thermal mass flowmeters is the sensitivity of some designs to changes in flow regime. Since the measurement principle is based on heat transfer by fluid convection, any factor influencing the convective heat-transfer efficiency will translate into a perceived difference in mass flow rate. It is a well-known fact in fluid mechanics that turbulent flows are more efficient at convecting heat than laminar flows, because the “stratified” nature of a laminar flowstream impedes heat transfer across the fluid width⁷⁶. In some thermal flowmeter designs, the walls of a heated metal *tube* serve as the “hot” element cooled by the fluid, and the difference between the rate of heat transferred by a laminar flowstream from the walls of a heated tube versus a turbulent flowstream can be great. Therefore, a change in flow regime (from turbulent to laminar, and vice-versa) will cause a calibration shift for this design of thermal mass flowmeter.

⁷⁶In a laminar flowstream, individual molecules do not cross paths, but rather travel in parallel lines. This means only those molecules traveling near the wall of a tube will be exposed to the temperature of the wall. The lack of “mixing” in a laminar flowstream means molecules traveling in the inner portions of the stream never contact the tube wall, and therefore never serve to transfer heat directly to or from the wall. At best, those inner-path molecules transfer heat by *conduction* with adjacent molecules which is a less efficient transfer mechanism than convection.

22.8 Weighfeeders

A special type flowmeter suited for powdered or granular solids is the *weighfeeder*. One of the most common weighfeeder designs consists of a conveyor belt with a section supported by rollers coupled to one or more load cells, such that a fixed length of the belt is continuously weighed:



The load cell measures the weight of a fixed-length belt section, yielding a figure of material weight per linear distance on the belt. A tachometer (speed sensor) measures the speed of the belt. The product of these two variables is the mass flow rate of solid material “through” the weighfeeder:

$$W = \frac{Fv}{d}$$

Where,

W = Mass flow rate (e.g. pounds per second)

F = Force of gravity acting on the weighed belt section (e.g. pounds)

v = Belt speed (e.g. feet per second)

d = Length of weighed belt section (e.g. feet)

A small weighfeeder (about two feet in length) is shown in the following photograph, the weighfeeder being used to feed powdered soda ash into water at a municipal filtration plant to neutralize pH:



In the middle of the belt's span (hidden from view) is a set of rollers supporting the weight of the belt and of the soda ash piled on the belt. This load cell array provides a measurement of pounds material per foot of belt length (lb/ft).

As you can see in this next picture, the soda ash powder simply falls off the far end of the conveyor belt, into the water below:



The speed sensor measures belt speed in feet per minute (ft/min). This measurement, when multiplied by the pounds-per-foot measurement sensed by the load cells, translates into a mass flow rate (W) in units of pounds per minute (lb/min). A simple unit conversion ($\times 60$) expresses the mass flow rate in units of pounds per hour (lb/h). A photograph of this weighfeeder's display screen shows these variables:



Note that the belt loading of 1.209 lb/ft and the belt speed of 0.62 feet per minute do *not* exactly equate⁷⁷ to the displayed mass flow rate of 43.7 lb/h. The reason for this discrepancy is that the camera's snapshot of the weighfeeder display screen happened to capture an image where the values were not simultaneous. Weighfeeders often exhibit fluctuations in belt loading during normal operation, leading to fluctuations in calculated mass flow rate. Sometimes these fluctuations in measured and calculated variables do not coincide on the display screen, given the latency inherent to the mass flow calculation (delaying the flow rate value until after the belt loading has been measured and displayed).

⁷⁷The proper mass flow rate value corresponding to these two measurements would be 45.0 lb/h.

22.9 Change-of-quantity flow measurement

Flow, by definition, is the passage of material from one location to another over time. So far this chapter has explored technologies for measuring flow rate en route from source to destination. However, a completely different method exists for measuring flow rates: measuring how much material has either departed or arrived at the terminal locations over time.

Mathematically, we may express flow as a ratio of quantity to time. Whether it is volumetric flow or mass flow we are referring to, the concept is the same: quantity of material moved per quantity of time. We may express average flow rates as ratios of changes:

$$\bar{W} = \frac{\Delta m}{\Delta t} \qquad \bar{Q} = \frac{\Delta V}{\Delta t}$$

Where,

\bar{W} = Average mass flow rate


\bar{Q} = Average volumetric flow rate

Δm = Change in mass

ΔV = Change in volume

Δt = Change in time

Suppose a water storage vessel is equipped with load cells to precisely measure weight (which is directly proportional to mass with constant gravity). Assuming only one pipe entering or exiting the vessel, any flow of water through that pipe will result in the vessel's total weight changing over time:



If the measured mass of this vessel decreased from 74688 kilograms to 70100 kilograms between 4:05 AM and 4:07 AM, we could say that the average mass flow rate of water leaving the vessel is 2294 kilograms per minute over that time span.

$$\bar{W} = \frac{\Delta m}{\Delta t} = \frac{70100 \text{ kg} - 74688 \text{ kg}}{4:07 - 4:05} = \frac{-4588 \text{ kg}}{2 \text{ min}} = -2294 \frac{\text{kg}}{\text{min}}$$

Note that this average flow measurement may be determined without any flowmeter of any kind installed in the pipe to intercept the water flow. All the concerns of flowmeters studied thus far (turbulence, Reynolds number, fluid properties, etc.) are completely irrelevant. We may measure practically any flow rate we desire simply by measuring stored weight (or volume) over time. A computer may do this calculation automatically for us if we wish, on practically any time scale desired.

Now suppose the practice of determining average flow rates every two minutes was considered too infrequent. Imagine that operations personnel require flow data calculated and displayed more often than just 30 times an hour. All we must do to achieve better time resolution is take weight (mass) measurements more often. Of course, each mass-change interval will be expected to be less with more frequent measurements, but the amount of time we divide by in each calculation will be proportionally smaller as well. If the flow rate happens to be absolutely steady, we may sample mass as frequently as we might like and we will still arrive at the same flow rate value as before (sampling mass just once every two minutes). If, however, the flow rate is not steady, sampling more often will allow us to better see the immediate “ups” and “downs” of flow behavior.

Imagine now that we had our hypothetical “flow computer” take weight (mass) measurements at an infinitely fast pace: an infinite number of samples per second. Now, we are no longer *averaging* flow rates over finite periods of time; instead we would be calculating *instantaneous* flow rate at any given *point* in time.

Calculus has a special form of symbology to represent such hypothetical scenarios: we replace the Greek letter “delta” (Δ , meaning “change”) with the roman letter “d” (meaning *differential*). A simple way of picturing the meaning of “d” is to think of it as meaning an *infinitesimal* change in whatever variable follows the “d” in the equation⁷⁸. When we set up two differentials in a quotient, we call the $\frac{d}{d}$ fraction a *derivative*. Re-writing our average flow rate equations in derivative (calculus) form:

$$W = \frac{dm}{dt} \qquad Q = \frac{dV}{dt}$$

Where,

W = Instantaneous mass flow rate

Q = Instantaneous volumetric flow rate


dm = Infinitesimal (infinitely small) change in mass

dV = Infinitesimal (infinitely small) change in volume

dt = Infinitesimal (infinitely small) change in time

⁷⁸While this may seem like a very informal definition of differential, it is actually rooted in a field of mathematics called *nonstandard analysis*, and closely compares with the conceptual notions envisioned by calculus’ founders.

We need not dream of hypothetical computers capable of infinite calculations per second in order to derive a flow measurement from a mass (or volume) measurement. Analog electronic circuitry exploits the natural properties of resistors and capacitors to essentially do this very thing in real time⁷⁹:



In the vast majority of applications you will see digital computers used to calculate average flow rates rather than analog electronic circuits calculating instantaneous flow rates. The broad capabilities of digital computers virtually ensures they will be used somewhere in the measurement/control system, so the rationale is to use the existing digital computer to calculate flow rates (albeit imperfectly) rather than complicate the system design with additional (analog) circuitry. As fast as modern digital computers are able to process simple calculations such as these anyway, there is little practical reason to prefer analog signal differentiation except in specialized applications where high speed performance is paramount.

Perhaps the single greatest disadvantage to inferring flow rate by differentiating mass or volume measurements over time is the requirement that the storage vessel have only one flow path in and out. If the vessel has multiple paths for liquid to move in and out (simultaneously), any flow rate calculated on change-in-quantity will be a *net* flow rate only. It is impossible to use this flow measurement technique to measure one flow out of multiple flows common to one liquid storage vessel.


A simple “thought experiment” confirms this fact. Imagine a water storage vessel receiving a flow rate in at 200 gallons per minute. Next, imagine that same vessel emptying water out of a second pipe at the exact same flow rate: 200 gallons per minute. With the exact same flow rate both entering and exiting the vessel, the water level in the vessel will remain constant. Any change-of-quantity flow measurement system would register zero change in mass or volume over time, consequently calculating a flow rate of absolutely zero. Truly, the *net* flow rate for this vessel is zero, but this tells us nothing about the flow in each pipe, except that those flow rates are equal in magnitude and opposite in direction.

⁷⁹To be precise, the equation describing the function of this analog differentiator circuit is: $V_{out} = -RC \frac{dV_{in}}{dt}$. The negative sign is an artifact of the circuit design – being essentially an inverting amplifier with negative gain – and not an essential element of the math.


22.10 Insertion flowmeters

This section does not describe a particular type of flowmeter, but rather a design that may be implemented for several different kinds of flow measurement technologies. When the pipe carrying process fluid is large in size, it may be impractical or cost-prohibitive to install a full-diameter flowmeter to measure fluid flow rate. A practical alternative for many applications is the installation of an *insertion* flowmeter: a probe that may be inserted into or extracted from a pipe, to measure fluid velocity in one region of the pipe's cross-sectional area (usually the center).

A classic example of an insertion flowmeter element is the *Annubar*, a form of averaging pitot tube pioneered by the Dieterich Standard corporation. The Annubar flow element is inserted into a pipe carrying fluid where it generates a differential pressure for a pressure sensor to measure:



The Annubar element may be extracted from the pipe by loosening a “gland nut” and pulling the assembly out until the end passes through a hand ball valve. Once the element has been extracted this far, the ball valve may be shut and the Annubar completely removed from the pipe:




For safety reasons, a “stop” is usually built into the assembly to prevent someone from accidentally pulling the element all the way out with the valve still open.

Other flowmeter technologies manufactured in insertion form include vortex, turbine, and thermal mass. An insertion-type turbine flowmeter appears in the following photographs:




If the flow-detection element is compact rather than distributed (as is certainly the case with the turbine flowmeter shown above), care must be taken to ensure correct positioning within the pipe. Since flow profiles are never completely flat, any insertion meter element will register a greater flow rate at the center of the pipe than near the walls. Wherever the insertion element is placed in the pipe diameter, that placement must remain consistent through repeated extractions and re-insertions or else the effective calibration of the insertion flowmeter will change every time it is removed and re-inserted into the pipe. Care must also be taken to insert the flowmeter so the flow element points directly upstream, and not at an angle.


A unique advantage of insertion instruments is that they may be installed in an operating pipe by using specialized *hot-tapping* equipment. A “hot tap” is a procedure whereby a safe penetration is made into a pipe while the pipe is carrying fluid under pressure. The first step in a hot-tapping operation is to weld a “saddle tee” fitting on the side of the pipe:




Next, a ball valve is bolted onto the saddle tee flange. This ball valve will be used to isolate the insertion instrument from the fluid pressure inside the pipe:



A special hot-tapping drill is then bolted to the open end of the ball valve. This drill uses a high-pressure seal to contain fluid pressure inside the drill chamber as a motor spins the drill bit. The ball valve is opened, then the drill bit is advanced toward the pipe wall where it cuts a hole into the pipe. Fluid pressure rushes into the empty chamber of the ball valve and hot-tapping drill as soon as the pipe wall is breached:



Once the hole has been completely drilled, the bit is extracted and the ball valve shut to allow removal of the hot-tapping drill:



Now there is a flanged and isolated connection into the “hot” pipe, through which an insertion flowmeter (or other instrument/device) may be installed.

Hot-tapping is a technical skill, with many safety concerns specific to different process fluids, pipe types, and process applications. This brief introduction to the technique is not intended to be instructional, but merely informational.

22.11 Process/instrument suitability

Every flow-measuring instrument exploits a physical principle to measure the flow rate of fluid stream. Understanding each of these principles as they apply to different flow-measurement technologies is the first and most important step in properly applying a suitable technology to the measurement of a particular process stream flow rate. The following table lists the specific operating principles exploited by different flow measurement technologies:

Flow measurement technology	Operating principle	Linearity	2-way flow
Differential pressure	Fluid mass self-acceleration, potential-kinetic energy exchange	$\sqrt{\Delta P}$	(some)
Laminar	Viscous fluid friction	linear	yes
Weirs & flumes	Fluid mass self-acceleration, potential-kinetic energy exchange	H^n	no
Turbine (velocity)	Fluid velocity spinning a vaned wheel	linear	yes
Vortex	von Kármán effect	linear	no
Magnetic	Electromagnetic induction	linear	yes
Ultrasonic	Sound wave time-of-flight	linear	yes
Coriolis	Fluid inertia, Coriolis effect	linear	yes
Turbine (mass)	Fluid inertia	linear	(some)
Thermal	Convective cooling, specific heat of fluid	linear	no
Positive displacement	Movement of fixed volumes	linear	(some)

A potentially important factor in choosing an appropriate flowmeter technology is energy loss caused by pressure drop. Some flowmeter designs, such as the common orifice plate, are inexpensive to install but carry a high price in terms of the energy lost in *permanent pressure drop* (the total, non-recoverable loss in pressure from the inlet of the device to the outlet, not the temporary pressure difference between inlet and vena contracta). Energy costs money, and so industrial facilities would be wise to consider the long-term cost of a flowmeter before settling on the one that is cheapest to install. It could very well be, for example, that an expensive venturi tube will cost less after years of operation than a cheap orifice plate⁸⁰.

In this regard, certain flowmeters stand above the rest: those with obstructionless flowtubes. Magnetic and ultrasonic flowmeters have no obstructions whatsoever in the path of the flow. This translates to (nearly) zero permanent pressure loss along the length of the tube, and therefore. Thermal mass and straight-tube Coriolis flowmeters are nearly obstructionless, while vortex and turbine meters are only slightly worse.

⁸⁰This is not always the case, as primary elements are often found on throttled process lines. In such cases where a control valve normally throttles the flow rate, any energy dissipated by the orifice plate is simply less energy that the valve would otherwise be required to dissipate. Therefore, the presence or absence of an orifice plate has no net impact on energy dissipation when used on a process flow throttled by a control valve, and therefore does not affect cost over time due to energy loss.

22.12 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Basic geometrical quantities:** *distance* and *velocity* are both one-dimensional measurements. *Area* is a two-dimensional measurement. *Volume* is a three-dimensional measurement. Relevant to unit conversions for flow measurements. For example, fluid velocity in ft/s, fluid pressure in pounds per ft², volumetric flow in ft³/s.
- **Density:** the ratio of mass to volume for a particular substance. Relevant to the kinetic energy of a moving fluid, as well as true-mass flow measurement.
- **Viscosity:** the resistance of a fluid to *shear*, which may be thought of as the “internal friction” of that fluid. Relevant to whether a fluid moves in *laminar* or *turbulent* fashion.
- **Laminar flow:** a condition where a fluid's molecules move in parallel paths, never crossing. Relevant to flowmeter selection, because many flowmeters require the flow regime to be turbulent rather than laminar (e.g. orifice plates, thermal mass). Also relevant to certain industrial processes such as mixing and heat transfer, because laminar flow impedes both these processes.
- **Turbulent flow:** a condition where a fluid's molecules move chaotically, randomly crossing pathways. Relevant to flowmeter selection, because many flowmeters require the flow regime to be turbulent rather than laminar (e.g. orifice plates, thermal mass). Also required for certain industrial processes such as mixing and heat transfer to efficiently occur.
- **Conservation of energy:** energy cannot be created or destroyed, only converted between different forms. Relevant to fluid velocities and pressures inside of flow elements such as venturi tubes, orifice plates, Pitot tubes, etc.
- **Bernoulli's equation:** $z_1\rho g + \frac{v_1^2\rho}{2} + P_1 = z_2\rho g + \frac{v_2^2\rho}{2} + P_2$, which is an application of the Law of Energy Conservation, stating that the sum of all forms of energy in a moving fluid stream (height, kinetic, and pressure) must remain the same. Relevant to calculations of pressure drop and pressure recovery across restrictions such as venturi tubes, orifice plates, etc.
- **Conservation of mass:** mass is an intrinsic property of matter, and as such cannot be created or destroyed. Relevant to the Continuity Principle for moving fluids, where the mass flow rate of a fluid entering a pipe must equal the mass flow rate exiting the pipe, assuming no accumulation or depletion (storage) of mass occurs within the pipe.
- **Reynolds number:** a unitless value representing the ratio of kinetic to viscous forces in a fluid. The greater the Reynolds number, the more turbulent the flow. The smaller the Reynolds number, the more likely the fluid will move in a *laminar* fashion. Relevant to many types of flowmeters, which operate accurately only within certain ranges of Reynolds number.

- **Flow profile or velocity profile:** the relative velocities of a fluid as it moves through a pipe, the velocity at the center being greater than the velocity at the pipe wall. Laminar flow is characterized by large differences in velocity along the profile, while turbulent flow exhibits a “flatter” profile with more consistent velocity across the pipe diameter. Relevant to insertion-type flowmeters such as Pitot tubes where the flowing velocity is sampled at only one point in the flowstream.
- **Ideal Gas Law:** $PV = nRT$, describing the relationship between gas pressure, chamber volume, gas quantity (in moles), and gas temperature. Relevant to measurements of gas flow rate at different pressures and temperatures (e.g. converting between “actual” or “flowing” units and “standard” units of gas measurement).
- **Newton’s Second Law of motion:** $F = ma$, describing how the acceleration of an object is directly proportional to the amount of applied (resultant) force and inversely proportional to its mass. Relevant to the development of a pressure difference across a flow element where the fluid molecules must either accelerate (positive a) or decelerate (negative a).
- **Inverse mathematical functions:** an inverse function, when applied to the result of its counterpart function, “un-does” the operation and leaves you with the original quantity. Relevant to the application of “square-root” in DP-based flow measurements. The natural characteristic of an accelerating or decelerating flow element (e.g. orifice plate) is to generate a pressure drop proportional to the square of the flow rate. Therefore, we must “square-root” that pressure signal in order to infer flow rate.
- **Wavelength vs. frequency:** $v = \lambda f$, describing the relationship between wavelength (λ) and frequency (f) for a wave. Relevant to vortex-shedding flowmeters, where the frequency is directly proportional to the velocity of the fluid. Also relevant to ultrasonic flowmeters.
- **Time, velocity, and distance:** $x = vt$, describing the relationship between velocity (v), time of travel (t), and distance traveled (x). Relevant to ultrasonic flowmeters, where travel time of an ultrasonic sound wave is used to calculate fluid velocity.
- **Speed of sound through a substance:** varies directly with the bulk modulus of the substance and inversely with the mass density of the substance as described by the formula $c = \sqrt{\frac{E}{\rho}}$. Relevant to Doppler-style ultrasonic flowmeters, where the Doppler frequency shift depends on the speed of sound through the fluid.
- **Resonance:** when something oscillates at its natural frequency. Relevant to Coriolis flowmeters, whose tubes are made to vibrate at their resonant frequency in order to measure fluid density. For a vibrating string (which closely approximates a vibrating tube), resonant frequency is directly proportional to string tension and inversely proportional to both length and mass as described by the formula $f = \frac{1}{2L} \sqrt{\frac{F_T}{\mu}}$.
- **Specific heat:** the amount of heat necessary to change the temperature of a some substance per unit mass and per unit of temperature. Relevant to thermal flowmeters, which work on the principle of heat transfer from a heated object to the moving fluid. The greater the mass flow rate of the fluid, the greater the heat transfer rate for any given specific heat value of the fluid. Specific heat is a function of the fluid’s chemical composition.

- **Scintillation:** the random and time-varying warping of light rays due to pockets of fluid having different refractive indices. Variations in temperature will cause this, as will turbulent motion of the fluid. Relevant to certain types of optical flowmeter, which pass light through a flowstream and look for patterns of scintillation after the light has traveled through the moving fluid.
- **Differentiation (calculus):** where one variable is proportional to the rate-of-change of two others. Differentiation always results in a division (quotient) of units. Relevant to calculations of flow rate based on mass or volume. Volumetric flow rate (Q) is equal to the rate of change in fluid volume over time ($Q = \frac{dV}{dt}$). Mass flow rate (W) is equal to the rate of change in fluid mass over time ($W = \frac{dm}{dt}$)
- **Integration (calculus):** where one variable is proportional to the accumulation of the product of two others. Integration always results in a multiplication of units. Relevant to calculations of mass or volume based on flow rate. Total volume of fluid passed by a point in a pipe (V) equal to the integral of volumetric flow rate times time: $V = \int Q dt$. Total mass of fluid passed by a point in a pipe (m) equal to the integral of mass flow rate times time: $m = \int W dt$.

References

AGA Report No. 3 – Orifice metering of natural gas and other related hydrocarbon fluids, Part 1 (General Equations and Uncertainty Guidelines), Catalog number XQ9017, American Gas Association and American Petroleum Institute, Washington D.C., Third Edition October 1990, Second Printing June 2003.

AGA Report No. 3 – Orifice metering of natural gas and other related hydrocarbon fluids, Part 2 (Specification and Installation Requirements), Catalog number XQ0002, American Gas Association and American Petroleum Institute, Washington D.C., Fourth Edition April 2000, Second Printing June 2003.

AGA Report No. 3 – Orifice metering of natural gas and other related hydrocarbon fluids, Part 3 (Natural Gas Applications), Catalog number XQ9210, American Gas Association and American Petroleum Institute, Washington D.C., Third Edition August 1992, Second Printing June 2003.

AGA Report No. 3 – Orifice metering of natural gas and other related hydrocarbon fluids, Part 4 (Background, Development, Implementation Procedure, and Subroutine Documentation for Empirical Flange-Tapped Discharge Coefficient Equation), Catalog number XQ9211, American Gas Association and American Petroleum Institute, Washington D.C., Third Edition October 1992, Second Printing August 1995, Third Printing June 2003.

Chow, Ven Te., *Open-Channel Hydraulics*, McGraw-Hill Book Company, Inc., New York, NY, 1959.

“Daniel Gas Ultrasonic Flow Meter Brochure”, document DAN-USM-B-FAMILY-0405, Daniel Measurement and Control, Inc., Emerson Process Management, 2005.

“Daniel Ultrasonic Gas Flowmeter Reference, Installation and Operations Manual”, part number 3-9000-740, revision H, Daniel Measurement and Control, Inc., Emerson Process Management, 2007.

“Flow Measurement User Manual”, Form Number A6043, Part Number D301224X012, Emerson Process Management, 2005.

Freund, William; Zanker, Klaus; Goodson, Dale; Hall, James E.; Jamieson, Andrew W.; “Operation of Ultrasonic Flow Meters at Conditions Different from their Calibration”, Paper 2.2, North Sea Flow Measurement Workshop, October 22-25, 2002.

Fribance, Austin E., *Industrial Instrumentation Fundamentals*, McGraw-Hill Book Company, New York, NY, 1962.

General Specifications: “EJX910A Multivariable Transmitter”, Document GS 01C25R01-01E, 5th edition, Yokogawa Electric Corporation, Tokyo, Japan, 2005.

Giancoli, Douglas C., *Physics for Scientists & Engineers*, Third Edition, Prentice Hall, Upper Saddle River, NJ, 2000.

Hanlon, Paul C., *Compressor Handbook*, The McGraw-Hill Companies, New York, NY, 2001.

Hofmann, Friedrich, *Fundamentals of Ultrasonic Flow Measurement for industrial applications*, Krohne Messtechnik GmbH & Co. KG, Duisburg, Germany, 2000.

Hofmann, Friedrich, *Fundamental Principles of Electromagnetic Flow Measurement*, 3rd Edition, Krohne Messtechnik GmbH & Co. KG, Duisburg, Germany, 2003.

“How Today’s Ultrasonic Meter Diagnostics Solve Metering Problems”, Daniel Measurement and Control, Inc., Emerson Process Management, 2010.

Improving Compressed Air System Performance – a sourcebook for industry, U.S. Department of Energy, Washington, DC, 2003.

Kallen, Howard P., *Handbook of Instrumentation and Controls*, McGraw-Hill Book Company, Inc., New York, NY, 1961.

Keisler, H. Jerome, *Elementary Calculus – An Infinitesimal Approach*, Second Edition, University of Wisconsin, 2000.

Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.

Miller, Richard W., *Flow Measurement Engineering Handbook*, Second Edition, McGraw-Hill Publishing Company, New York, NY, 1989.

Parker, Jody; Stobie, Gordon; Melnyk, Ivan; Letton, Chip; “Flare Metering With Optics – From Blue-Sky Technology to the Real World”, 25th International North Sea Flow Measurement workshop, Oslo, Norway, October 16-19, 2007.

Parshall, R. L., “Measuring Water in Irrigation Channels”, Farmers’ Bulletin number 1683, pages 1-29, Washington D.C., 1941.

Price, James F., *A Coriolis Tutorial*, version 3.3, Woods Hole Oceanographic Institution, Woods Hole, MA, 2006.

“Proving Coriolis Flowmeters”, document 1004732, Revision A, Micro Motion, Inc., Boulder, CO, October 1998.

Spink, L. K., *Principles and Practice of Flow Meter Engineering*, Ninth Edition, The Foxboro Company, Foxboro, MA, 1967.

Tech-Spec: “SCFM (Standard CFM) vs. ACFM (Actual CFM)”, Reference 15-010504.006, Sullair Corporation, 2004.

“Top Mount Installation for DP Flowmeters in Steam Service”, whitepaper number 00870-0200-4809, Emerson Process Management, August 2009.

Vennard, John K., *Elementary Fluid Mechanics*, 3rd Edition, John Wiley & Sons, Inc., New York, NY, 1954.

Wang, Ting-I; Buhr, Eric; “Optical Flow Sensing: A New Approach to an Old Problem”, May 2001.

Zanker, Klaus J., “Diagnostic Ability of the Daniel Four Path Ultrasonic Flow Meter”, Daniel Measurement and Control, Inc., Emerson Process Management, 2010.

Chapter 23

Continuous analytical measurement

In the field of industrial instrumentation and process control, the word *analyzer* generally refers to an instrument tasked with measuring the concentration of some substance, usually mixed with other substances. Unlike the other “bulk” measurement devices for sensing such general variables as pressure, level, temperature, or flow, an analytical device must *discriminately detect* one specific substance while ignoring all other substances present in the sample. This single problem accounts for much of the complexity of analytical instrumentation: *how do we achieve a high degree of selectivity in our measurement?*

Analytical instruments generally achieve selectivity by measuring some property of the substance of interest unique to that substance alone, or at least unique to it among the possible substances likely to be found in the process sample. For example, an optically-based analyzer might achieve selectivity by measuring the intensities of only those particular wavelengths of light absorbed by the compound of interest, and absorbed by none of the other wavelengths. A “paramagnetic” oxygen gas analyzer achieves selectivity by exploiting the paramagnetic properties of oxygen gas, since no other industrial gas is nearly as paramagnetic as oxygen. A pH analyzer achieves hydrogen ion selectivity by using a specially-prepared glass membrane constructed to pass only hydrogen ions.

Problems are sure to arise if the measured property of the substance of interest is not as unique as originally thought. This may occur due to oversight on the part of the person originally choosing the analyzer technology, or it may occur as a result of changes made to the process chemistry, whether by intentional modification of the process equipment or by abnormal operating conditions. For example, a gas that happens to absorb some (or all!) of the same light wavelengths as the gas of interest will cause false measurements in an optical absorption analyzer. Nitric oxide (NO) gas is considered an interferent for paramagnetic oxygen analyzers, since this gas is one of the few gases besides oxygen also exhibiting significant paramagnetism. A pH analyzer immersed in a liquid solution containing an abundance of sodium ions may fall victim to measurement errors, because sodium ions also happen to interact with the glass membrane of a pH electrode to generate a voltage. These are but a few practical examples of analyzer non-selectivity.

For this reason, the student of analytical instrumentation must always pay close attention to the *underlying principle of measurement* for any analyzer technology, looking out for any ways that

analyzer may be “fooled” by the presence of some *other* substance than the one the analyzer was designed to measure.

Some chemical analyzers are known for their unreliability, requiring a variety of conditions to be just right in order for to operate as designed. Proper conditioning of the sample to be analyzed (e.g. filtering, heating or cooling, drying) is one of the many points of failure potentially plaguing analytical instruments, as a typical sample conditioning system is a complex arrangement of tubes, valves, and instruments in its own right. Analyzers also tend to be expensive, both in their initial and consumable costs. These and other reasons are why analytical instrument maintenance is considered an advanced skill. Keeping analyzers in good working order is a challenging technical task for any technician, and usually requires special training above and beyond the knowledge and skill base required for general instrument maintenance.

An interesting historical reference on this point comes from the book *Instrumentation and Control in the German Chemical Industry*, written after the end of World War II as British investigators toured a variety of chemical manufacturing plants in Germany to learn about their process instrumentation. The instrument head at the I.G. Chemische Werke in the city of Hüls communicated the following sentiment to the authors of this text regarding the maintenance of analytical instrumentation by a special department called the Physical Laboratory:

[He] considers that the ordinary instrument man is not capable of giving the skilled maintenance necessary with this class of instrument. From a maintenance point of view he considers them in a different class to flowmeters, pressure gauges, etc. The section instrument manager is responsible for the routine day to day maintenance such as changing of charts and filling of pens, but special maintenance and calibration is carried out in the laboratory or in the plant by skilled men from the physical laboratory. In [his] opinion this is the only satisfactory way of maintaining complicated analysis instruments as only men who are having continuous experience can carry out the special maintenance, repair and calibration. (page 115)

Mind you, this was during a time when the state of the art for optical absorption analyzers required 2 to 3 days of labor to complete a full calibration, but in some very fundamental ways the challenge of analyzer maintenance remains unaltered. Chemically-selective measurements are by their very nature more complex than bulk measurements such as pressure, level, temperature, or flow, and as such they are prone to a greater number and more complex set of faults. *If your career goal is to become as knowledgeable and skillful as possible in the field of industrial process measurement, analytical instrumentation is your specialty of choice!*

23.1 Conductivity measurement

The electrical conductivity of liquids is an important analytical measurement in many industrial processes. This measurement is one of the more non-specific types of analytical technologies, because it does not discriminate between different conductive substances dissolved in the solution. For this reason, conductivity measurement is found in process applications where the type of conductive substance is irrelevant (e.g. ultra-pure water treatment for semiconductor “chip” manufacturing, where *any* conductive substance dissolved in the water is undesirable), or where the substance of interest is known to be the only conductive substance present in significant quantity (e.g. controlling the salinity of a brine solution, where large quantities of salt are added to water).

Electrical conductivity in metals is the result of free electrons drifting within a “lattice” of atomic nuclei comprising the metal object. When a voltage is applied across two points of a metal object, these free electrons immediately drift toward the positive pole (anode) and away from the negative pole (cathode).

Electrical conductivity in liquids is another matter entirely. Here, the charge carriers are *ions*: electrically imbalanced atoms or molecules that are free to drift because they are not “locked” into a lattice structure as is the case with solid substances. The degree of electrical conductivity of any liquid is therefore dependent on the ion density of the solution (how many ions freely exist per unit volume of liquid). When a voltage is applied across two points of a liquid solution, negative ions will drift toward the positive pole (anode) and positive ions will drift toward the negative pole (cathode). In honor of this directional drifting, negative ions are sometimes called *anions* (attracted to the *anode*), while positive ions are sometimes called *cations* (attracted to the *cathode*).

Electrical conductivity in gases is much the same: ions are the charge carriers. However, with gases at room temperature, ionic activity is virtually nonexistent. A gas must be superheated into a *plasma* state before substantial ions exist which can support an electric current.

23.1.1 Dissociation and ionization in aqueous solutions

Pure water is a very poor conductor of electricity. Some water molecules will “ionize” into unbalanced halves (instead of H_2O , you will find some negatively charged hydroxyl ions (OH^-) and some positively charged hydrogen ions¹ (H^+), but the percentage is extremely small at room temperature.

Any substance that enhances electrical conductivity when dissolved in water is called an *electrolyte*. This enhancement of conductivity occurs due to the molecules of the electrolyte separating into positive and negative ions, which are then free to serve as electrical charge carriers. If the electrolyte in question is an *ionically-bonded* compound² (table salt is a common example), the ions forming that compound naturally separate in solution, and this separation is called *dissociation*. If the electrolyte in question is a *covalently-bonded* compound³ (hydrogen chloride is an example), the separation of those molecules into positive and negative ions is called *ionization*.

Both *dissociation* and *ionization* refer to the separation of formerly joined atoms upon entering a solution. The difference between these terms is the type of substance that splits: “dissociation” refers to the division of ionic compounds (such as table salt), while “ionization” refers to covalent-bonded (molecular) compounds such as HCl which are not ionic in their pure state.

Ionic impurities added to water (such as salts and metals) immediately dissociate and become available to act as charge carriers. Thus, the measure of a water sample’s electrical conductivity is a function⁴ of its ionic impurity concentration. Conductivity is therefore an important analytical measurement for certain water purity applications, such as the treatment of boiler feedwater, and the preparation of high-purity water used for semiconductor manufacturing.

It should be noted that conductivity measurement is a very *non-specific* form of analytical measurement. The conductivity of a liquid solution is a gross indication of its ionic content, but it tells us nothing specific about the *type* or *types* of ions present in the solution. Therefore, conductivity measurement is meaningful only when we have prior knowledge of the particular ionic species present in the solution (or when the purpose is to eliminate all ions in the solution such as in the case of ultra-pure water treatment, in which case we do not care about types of ions because our ideal goal is zero conductivity).

¹Truth be told, free hydrogen ions are extremely rare in an aqueous solution. You are far more likely to find them bound to normal water molecules to form positive hydronium ions (H_3O^+). For simplicity’s sake, though, professional literature often refers to these positive ions as “hydrogen” ions and even represent them symbolically as H^+ .


²Ionic compounds are formed when oppositely charged atomic ions bind together by mutual attraction. The distinguishing characteristic of an ionic compound is that it is a conductor of electricity in its pure, liquid state. That is, it readily separates into anions and cations all by itself. Even in its solid form, an ionic compound is already ionized, with its constituent atoms held together by an imbalance of electric charge. Being in a liquid state simply gives those atoms the physical mobility needed to dissociate.

³Covalent compounds are formed when electrically neutral atoms bind together by the mutual sharing of valence electrons. Such compounds are not good conductors of electricity in their pure, liquid states.

⁴It should be noted that the relationship between conductivity and electrolyte concentration in a solution is typically non-linear. Not only does the electrical conductivity of a solution *not* follow a set proportion to concentration, but even the slope of the relationship may change from positive to negative over a wide range of concentrations. This fact makes conductivity measurement in liquid solutions useful for concentration analysis only over limited ranges.


23.1.2 Two-electrode conductivity probes

We may measure the electrical conductivity of a liquid solution by passing an electric current through it. The most primitive form of conductivity sensor (sometimes referred to as a conductivity *cell*) consists of two metal electrodes inserted in the solution, connected to a circuit designed to measure conductance (G), the reciprocal of resistance ($\frac{1}{R}$):



A general problem faced with electrical measurements of liquid conductance is that the derived conductance value (G) does not tell us much about the liquid itself, because that measurement depends just as much on the geometry of the plates (their area A and separation distance d) as it does on the ionic activity of the liquid solution. If we are trying to analyze the liquid all by itself, what we really need is a measurement of specific conductivity (k , or *conductance*) independent of plate geometry.

We face the same essential problem when trying to quantify the resistivity of metal conductors. If we measure the resistance of a piece of wire in the same manner shown in the previous illustration measuring liquid conductance, we arrive at a result that is every bit as much dependent on the length and area of the wire specimen as it is on the resistivity of the metal itself:



In other words, the calculated value in ohms (from direct voltage and current measurements) for the resistance of this metal specimen doesn't tell us much about that type of metal in general, but rather it tells us the resistance of that *particular specimen* of wire. In order to calculate the specific resistance (ρ , or *resistivity*) of the metal, we must also account for the specimen's length (d) and cross-sectional area (A).

The mathematical relationship between conductance (G), plate area (A), plate distance (d), and the actual conductivity of the liquid (k) is expressed in the following formula:

$$G = k \frac{A}{d}$$

Where,

G = Conductance, in Siemens (S)

k = Specific conductance (conductivity) of liquid, in Siemens per centimeter (S/cm)

A = Electrode area (each), in square centimeters (cm^2)

d = Electrode separation distance, in centimeters (cm)

Manipulating this formula to solve for conductivity (k) of the liquid:

$$k = \frac{Gd}{A}$$

The unit of Siemens per centimeter for liquid conductivity may seem odd at first, but it is necessary to account for all the units present in the variables of the equation. A simple dimensional analysis proves this:

$$k = \frac{Gd}{A} \quad \left[\frac{\text{S}}{\text{cm}} \right] = \frac{[\text{S}][\text{cm}]}{[\text{cm}^2]}$$

In order to quantify the plate geometry for any particular cell, manufacturers typically express the fraction $\frac{d}{A}$ as a single value called the *cell constant*, symbolized by the Greek letter “theta” (θ) and expressed in the unit of inverse centimeters (cm^{-1}):

$$\theta = \frac{d}{A} \quad \left[\frac{1}{\text{cm}} \right] = [\text{cm}^{-1}] = \frac{[\text{cm}]}{[\text{cm}^2]}$$

Substituting θ for the quotient $\frac{d}{A}$ in the conductivity formula reveals conductivity to be the simple product of measured conductance (G) and the cell constant:

$$k = G\theta$$


Where,

k = Specific conductivity of liquid, in Siemens per centimeter (S/cm)

G = Conductance, in Siemens (S)

θ = Cell constant, in inverse centimeters (cm^{-1})

The following photograph shows an example of such a direct-contact style of conductivity probe, consisting of stainless steel electrodes contacting the fluid flowing through a glass tube:




Two-electrode conductivity cells are not very practical in real applications, because mineral and metal ions attracted to the electrodes tend to “foul” the electrodes over time forming solid, insulating barriers on the electrodes. While this “electroplating” action may be substantially reduced by using AC instead of DC⁵ to excite the sensing circuit, it is usually not enough. Over time, the conductive barriers formed by ions bonded to the electrode surfaces will create calibration errors by making the instrument “think” the liquid is less conductive than it actually is.

⁵The use of alternating current forces the ions to switch directions of travel many times per second, thus reducing the chance they have of bonding to the metal electrodes.

23.1.3 Four-electrode conductivity probes

A very old electrical technique known as the *Kelvin* or *four-wire* resistance-measuring method is a practical solution to the problem of electrode fouling faced by two-electrode conductivity probes. Commonly employed to make precise resistance measurements for scientific experiments in laboratory conditions, as well as measuring the electrical resistance of strain gauges and other resistive sensors such as RTDs, the four-wire technique uses four conductors to connect the resistance under test to the measuring instrument:




$$R_{\text{specimen}} = \frac{\text{Voltmeter indication}}{\text{Current source}}$$

Only the outer two conductors carry substantial current. The inner two conductors connecting the voltmeter to the test specimen carry negligible current (due to the voltmeter's extremely high input impedance) and therefore drop negligible voltage along their lengths. Voltage dropped across the current-carrying (outer) wires is irrelevant, since that voltage drop is never detected by the voltmeter.

Since the voltmeter only measures voltage dropped across the specimen (the resistor under test), and not the test resistance plus wiring resistance, the resulting resistance measurement is much more accurate than if only two wires were used to connect the test meters to the specimen.

In the case of conductivity measurement, it is not wire resistance that we care to ignore, but rather the added resistance caused by fouling of the electrodes. By using four electrodes instead of two, we are able to measure voltage dropped across a length of liquid solution *only*, and completely ignore the resistive effects of electrode fouling:



In the 4-wire conductivity cell, any electrode fouling will merely burden the current source by causing it to output a greater voltage, but it will *not* affect the amount of voltage detected by the two inner electrodes as that electric current passes through the liquid. Any fouling that happens to occur⁶ on the two inner electrodes is of no effect to our conductivity measurement because these inner electrodes carry negligible current. With little or no current through the inner electrodes, there will be negligible voltage dropped across any resistive coating that happens to form on them, and thus the voltmeter will still register the true voltage dropped by the liquid solution.

If the solution's conductivity is defined as the product of the measured conductance and the cell constant ($k = G\theta$), and conductance is defined as the ratio of current to voltage ($G = \frac{I}{V}$), then we may determine conductivity from voltage and current measurements by combining these two equations:

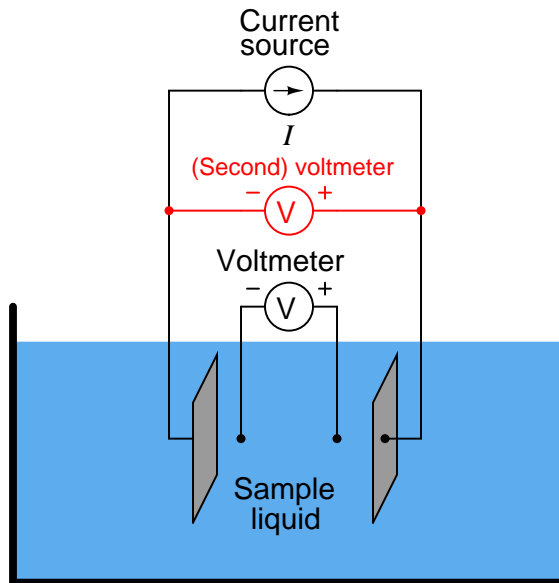
$$k = G\theta \quad G = \frac{I}{V}$$

. . . substituting $\frac{I}{V}$ for G . . .

$$k = \frac{I}{V}\theta \quad \text{or} \quad k = \frac{I\theta}{V}$$

⁶There will be very little if any fouling on these electrodes anyway because they carry no current, and thus provide no reason for ions to migrate toward them.


Some conductivity instruments employ a second voltmeter to measure the voltage dropped between the “excitation” electrodes, to indicate electrode fouling:



Any form of electrode fouling will cause this secondary voltage measurement to disproportionately exceed the first, thus providing an indicator that instrument technicians may use for predictive maintenance (telling them when the probes need cleaning or replacement). Meanwhile, the primary voltmeter will do its job of accurately measuring liquid conductivity so long as the current source is still able to output its normal amount of current.

23.1.4 Electrodeless conductivity probes

An entirely different design of conductivity cell called *electrodeless* uses electromagnetic induction rather than direct electrical contact to detect the conductivity of the liquid solution. This cell design enjoys the distinct advantage of virtual immunity to fouling⁷, since there is no direct electrical contact between the measurement circuit and the liquid solution. Instead of using two or four electrodes inserted into the solution for conductivity measurement, this cell uses two *toroidal* inductors (one to induce an AC voltage in the liquid solution, and the other to measure the strength of the resulting current through the solution):




The basic idea of this instrument is that a primary coil energized by AC power induces an electric current that passes through the sample liquid. This current, in turn, induces a measurable voltage in a secondary coil. Since ferromagnetic toroids do an excellent job of containing their own magnetic fields, there will be negligible mutual inductance between the two wire coils. The *only* way a voltage will be induced in the secondary coil is if there is an AC current passing through the center of that coil, through the liquid itself. If the liquid is non-conductive, the secondary coil will see no induced voltage at all despite being situated near the energized primary coil. The more conductive the liquid solution, the more current will pass through the center of both coils (through the liquid), thus producing a greater induced voltage at the secondary coil. Secondary coil voltage therefore is directly proportional to liquid conductivity⁸.

⁷Toroidal conductivity sensors may suffer calibration errors if the fouling is so bad that the hole becomes choked off with sludge, but this is an extreme condition. These sensors are far more tolerant to fouling than any form of contact-type (electrode) conductivity cell.

⁸Note that this is opposite the behavior of a direct-contact conductivity cell, which produces *less* voltage as the liquid becomes more conductive.

The equivalent electrical circuit for a toroidal conductivity probe looks like a pair of transformers, with the liquid acting as a resistive path for current to connect the two transformers together:



Toroidal conductivity cells are preferred over direct-contact conductivity cells whenever possible, due to their ruggedness and virtual immunity to fouling. However, toroidal cells are not sensitive enough for conductivity measurement in high-purity applications such as boiler feedwater treatment and ultra-pure water treatment necessary for pharmaceutical and semiconductor manufacturing. As always, the manufacturer's specifications are the best source of information for conductivity cell applicability in any particular process.

The following photograph shows a toroidal conductivity probe mounted on a display board at a trade-show, along with a conductivity transmitter (to both display the conductivity measurement in millisiemens per centimeter and also transmit the measurement as a 4-20 mA analog signal):



23.2 pH measurement

pH is the measurement of the hydrogen ion activity in a liquid solution. It is one of the most common forms of analytical measurement in industry, because pH has a great effect on the outcome of many chemical processes. Food processing, water treatment, pharmaceutical production, steam generation (thermal power plants), and alcohol manufacturing are just some of the industries making extensive use of pH measurement (and control). pH is also a significant factor in the corrosion of metal pipes and vessels carrying aqueous (water-based) solutions, so pH measurement and control is important in the life-extension of these capital investments.

In order to understand pH measurement, you must first understand the chemistry of pH. Please refer to section 3.12 beginning on page 284 for a theoretical introduction to pH.

23.2.1 Colorimetric pH measurement

One of the simplest ways to measure the pH of a solution is by color. Some chemical compounds dissolved in an aqueous solution will change color if the pH value of that solution falls within a certain range. *Litmus paper* is a common laboratory application of this principle, where a color-changing chemical substance infused on a paper strip changes color when dipped in the solution. Comparing the final color of the litmus paper to a reference chart yields an approximate pH value for the solution.

A natural example of this phenomenon is well-known to flower gardeners, who recognize that hydrangea blossoms change color with the pH value of the soil. In essence, these plants act as organic litmus indicators⁹. This hydrangea plant indicates acidic soil by the violet color of its blossoms:




⁹Truth be told, the color of a hydrangea blossom is only indirectly determined by soil pH. Soil pH affects the plant's uptake of aluminum, which is the direct cause of color change. Interestingly, the pH-color relationship of a hydrangea plant is exactly opposite that of common laboratory litmus paper: red litmus paper indicates an acidic solution while blue litmus paper indicates an alkaline solution; whereas red hydrangea blossoms indicate alkaline soil while blue (or violet) hydrangea blossoms indicate acidic soil.

Another example of a natural colorimetric pH indicator is *red cabbage*. If some red cabbage is chopped and cooked, the juices released by the cabbage¹⁰ will be sensitive to pH. This makes a very easy demonstration for the home kitchen. In these two photographs, you see how liquid may be collected from the cabbage in a steaming pot, and then transferred to three glasses for testing:



Adding vinegar (acid) to one glass, baking soda (caustic/base/alkaline) to another glass, and leaving the third glass unaltered (as an experimental “control”), we see striking differences in the color of each solution. Vinegar turns the cabbage juice red, while baking soda turns it dark green, compared to its original purple color:




In fact, you may make your own crude form of litmus paper by soaking paper strips with red cabbage juice!

¹⁰ *Flavin*, classified as an *anthocyanin*, is the pigment in red cabbage responsible for the pH-indicating behavior. This same pigment also changes color according to soil pH while the cabbage plant is growing, much like a hydrangea. Unlike hydrangeas, the coloring of a red cabbage is more akin to litmus paper, with red indicating acidic soil.

23.2.2 Potentiometric pH measurement

Color-change is a common pH test method used for manual laboratory analyses, but it is not well-suited to continuous process measurement. By far the most common pH measurement method in use is *electrochemical*: special pH-sensitive electrodes inserted into an aqueous solution will generate a voltage dependent upon the pH value of that solution.

Like all other potentiometric (voltage-based) analytical measurements, electrochemical pH measurement is based on the *Nernst equation*, which describes the electrical potential created by ions migrating through a permeable membrane. The “textbook example” of this is a device called a *concentration cell*, where two halves of an electrochemical cell are filled with solutions having different concentrations of ions (i.e. different molarities):



Where,

V = Voltage produced across membrane due to ion exchange (volts)

R = Universal gas constant (8.315 J/mol·K)

T = Absolute temperature (Kelvin)

n = Number of electrons transferred per ion exchanged (unitless)

F = Faraday constant, in coulombs per mole (96485 C/mol e^-)

C_1 = Concentration of ion in measured solution (moles per liter of solution, M)

C_2 = Concentration of ion in reference solution (moles per liter of solution, M)

As ions naturally migrate through this membrane in an attempt¹¹ to equalize the two

¹¹Of course, ions possess no agency and therefore cannot literally “attempt” anything. What is happening here is the normal process of *diffusion* whereby the random motions of individual molecules tends to evenly distribute those molecules throughout a space. If a membrane divides two solutions of differing ionic concentration, ions from the more concentrated region will, over time, migrate to the region of lower concentration until the two concentrations are

concentrations, a voltage corresponding to the *difference* in ion concentrations between the two cell halves will develop between the two electrodes. The greater the difference in concentrations between the two sides, the greater the voltage produced by the cell. The Nernst voltage may be used to infer the concentration of a specific type of ion if the membrane is *selectively permeable* to that one type of ion.

We may write the Nernst equation using either natural logarithms (ln) or common logarithms (log). Either form of the Nernst equation works to predict the voltage generated by a concentration cell. The typical form applied to pH measurement calculations is the common log version, which makes more intuitive sense since pH is defined as the common logarithm of hydrogen ion activity:


$$V = \frac{RT}{nF} \ln \left(\frac{C_1}{C_2} \right) \qquad V = \frac{2.303RT}{nF} \log \left(\frac{C_1}{C_2} \right)$$

Both forms of the Nernst equation predict a greater voltage developed across the thickness of a membrane as the concentrations on either side of the membrane differ to a greater degree. If the ionic concentration on both sides of the membrane are equal, no Nernst potential will develop¹².

equal to each other. Truth be told, ions are continually migrating in *both* directions through the porous membrane at all times, but the rate of migration from the high concentration to the low concentration solutions is greater than the other direction simply because there are more ions present to migrate that way. After the two solutions have become equal in ionic concentration, the random migration still proceeds in both directions, but now the rates in either direction are equal and therefore there is zero *net* migration.

¹²This is apparent from a mathematical perspective by examination of the Nernst equation: if the concentrations are equal (i.e. $C_1 = C_2$), then the ratio of $\frac{C_1}{C_2}$ will be equal to 1. Since the logarithm of 1 is zero, this predicts zero voltage generated across the membrane. From a chemical perspective, this corresponds to the condition where random ion migration through the porous membrane is equal in both directions. In this condition, the Nernst potentials generated by the randomly-migrating ions are equal in magnitude and opposite in direction (polarity), and therefore the membrane generates zero overall voltage.

In the case of pH measurement, the Nernst equation describes the amount of electrical voltage developed across a special *glass* membrane due to hydrogen ion exchange between the process liquid solution and a *buffer solution* inside the bulb formulated to maintain a constant pH value of 7.0 pH. Special pH-measurement electrodes are manufactured with a closed end made of this glass, a small quantity of buffer solution contained within the glass bulb:



Any concentration of hydrogen ions in the process solution differing from the hydrogen ion concentration in the buffer solution ($[H^+] = 1 \times 10^{-7} M$) will cause a voltage to develop across the

thickness of the glass. Thus, a standard pH measurement electrode produces no potential when the process solution's pH value is exactly 7.0 pH (i.e. when the process solution has the same hydrogen ion activity as the buffer solution within the bulb).

Given the knowledge that the measurement bulb is filled with a buffer solution having a pH value of 7, we may conclude that one of the concentrations for the glass membrane will always have a value of $1 \times 10^{-7} M$. We may manipulate the Nernst equation to reflect this knowledge, and to express the potential developed in terms of the pH of both solutions, since we know pH is defined the negative logarithm of hydrogen ion molarity:

$$V = \frac{2.303RT}{nF} \log \left(\frac{C_1}{C_2} \right)$$

$$V = \frac{2.303RT}{nF} (\log C_1 - \log C_2)$$

If we know that $\text{pH} = -\log[\text{H}^+]$

$$V = \frac{2.303RT}{nF} (-\text{pH}_1 - (-\text{pH}_2))$$

$$V = \frac{2.303RT}{nF} (\text{pH}_2 - \text{pH}_1)$$

$$V = \frac{2.303RT}{nF} (7 - \text{pH}_1)$$


Thus, the Nernst voltage produced by a glass pH electrode is directly proportional to the difference in pH value between the measured solution (pH_1) and the probe's internal 7.0 pH buffer.

The glass used to manufacture this electrode is no ordinary glass. Rather, it is specially manufactured to be *selectively permeable* to hydrogen ions¹³. If it were not for this fact, the electrode might generate voltage as it contacted any number of different ions in the solution. This would make the electrode non-specific, and therefore useless for pH measurement.

Manufacturing processes for pH-sensitive glass are highly guarded trade secrets. There seems to be something of an art to the manufacture of an accurate, reliable, and long-lived pH electrode. A variety of different measurement electrode designs exist for different process applications, including high pressure and high temperature services.


¹³It is a proven fact that sodium ions in relatively high concentration (compared to hydrogen ions) will also cause a Nernst potential across the glass of a pH electrode, as will certain other ion species such as potassium, lithium, and silver. This effect is commonly referred to as *sodium error*, and it is usually only seen at high pH values where the hydrogen ion concentration is extremely low. Like any other analytical technology, pH measurement is subject to "interference" from species unrelated to the substance of interest.

Actually measuring the voltage developed across the thickness of the glass electrode wall, however, presents a bit of a problem: while we have a convenient electrical connection to the solution *inside* the glass bulb, we do not have any place to connect the other terminal of a sensitive voltmeter to the solution *outside* the bulb¹⁴. In order to establish a complete circuit from the glass membrane to the voltmeter, we must create a zero-potential electrical junction with the process solution. To do this, we use another special electrode called a *reference electrode* immersed in the same liquid solution as the measurement electrode:



¹⁴Remember that voltage is always measured *between two points!*

Together, the measurement and reference electrodes provide a voltage-generating element sensitive to the pH value of whatever solution they are submerged in:



The most common configuration for modern pH probe sets is what is called a *combination electrode*, which combines both the glass measurement electrode and the porous reference electrode in a single unit. This photograph shows a typical industrial combination pH electrode:



The red-colored plastic cap on the right-hand end of this combination electrode covers and

protects a gold-plated coaxial electrical connector, to which the voltage-sensitive pH indicator (or transmitter) attaches.

Another model of pH probe appears in the next photograph. Here, there is no protective plastic cap covering the probe connector, allowing a view of the gold-plated connector bars:



A close-up photograph of the probe tip reveals the glass measurement bulb, a weep hole for process liquid to enter the reference electrode assembly (internal to the white plastic probe body), and a metal *solution ground* electrode:



It is extremely important to always keep the glass electrode wet. Its proper operation depends on complete *hydration* of the glass, which allows hydrogen ions to penetrate the glass and develop the Nernst potential. The probes shown in these photographs are shown in a dry state only because they have already exhausted their useful lives and cannot be damaged any further by dehydration.

The process of hydration – so essential to the working of the glass electrode – is also a mechanism of wear for pH probes. Layers of glass “slough” off over time when continuously hydrated, which means that glass pH electrodes have a limited life whether they are being used to measure the pH of a process solution (continuously wet) or if they are being stored on a shelf (maintained in a wet state by a small quantity of potassium hydroxide held close to the glass probe by a liquid-tight cap). It is therefore impossible to extend the shelf life of a glass pH electrode indefinitely.

A common installation for industrial pH probe assemblies is to simply dip them into an open vessel containing the solution of interest. This arrangement is very common in water treatment applications, where the water mostly flows in open vessels by gravity at the treatment facility. A photograph showing a pH measurement system for the “outfall” flow of water from an industrial facility appears here:



Water flowing from the discharge pipe of the facility enters an open-top stainless steel tank where the pH probe hangs from a bracket. An overflow pipe maintains a maximum water level in the tank as water continuously enters it from the discharge pipe. The probe assembly may be easily removed for maintenance:



An alternative design for industrial pH probes is the *insertion* style, designed to install in a pressurized pipe. Insertion probes are designed to be removed while the process line remains pressurized, to facilitate maintenance without interrupting continuous operation:



The probe assembly inserts into the process line through the open bore of a 90° turn ball valve. The left-hand photograph (above) shows the retaining nut loosened, allowing the probe to slide up and out of the pipe. The right-hand photograph shows the ball valve shut to block process liquid pressure from escaping, while the technician unlatches the clamps securing the probe to the pipe fitting.

Once the clamp is unlatched, the probe assembly may be completely detached from the pipe, allowing cleaning, inspection, calibration, repair, and/or replacement:



The voltage produced by the measurement electrode (glass membrane) is quite modest. A calculation for voltage produced by a measurement electrode immersed in a 6.0 pH solution shows this. First, we must calculate hydrogen ion concentration (activity) for a 6.0 pH solution, based on the definition of pH being the negative logarithm of hydrogen ion molarity:

$$\text{pH} = -\log[\text{H}^+]$$

$$6.0 = -\log[\text{H}^+]$$

$$-6.0 = \log[\text{H}^+]$$

$$10^{-6.0} = 10^{\log[\text{H}^+]}$$

$$10^{-6.0} = [\text{H}^+]$$

$$[\text{H}^+] = 1 \times 10^{-6} M$$

This tells us the concentration of hydrogen ions in the 6.0 pH solution¹⁵. We know that the buffer solution inside the glass measurement bulb has a stable value of 7.0 pH (hydrogen ion concentration of $1 \times 10^{-7} M$, or 0.0000001 moles per liter), so all we need to do now is insert these values into the Nernst equation to see how much voltage the glass electrode should generate. Assuming a solution temperature of 25 °C (298.15 K), and knowing that n in the Nernst equation will be equal to 1 (since each hydrogen ion has a single-value electrical charge):

$$V = \frac{2.303RT}{nF} \log \left(\frac{C_1}{C_2} \right)$$

$$V = \frac{(2.303)(8.315)(298.15)}{(1)(96485)} \log \left(\frac{1 \times 10^{-6} M}{1 \times 10^{-7} M} \right)$$

$$V = (59.17 \text{ mV})(\log 10) = 59.17 \text{ mV}$$

We get the same result using our modified version of the Nernst equation:

$$V = \frac{2.303RT}{nF} (7 - \text{pH}_1)$$

$$V = \frac{(2.303)(8.315)(298.15)}{(1)(96485)} (7 - 6)$$

$$V = (59.17 \text{ mV})(1) = 59.17 \text{ mV}$$

¹⁵Hydrogen ion *concentration* being practically the same as hydrogen ion *activity* for dilute solutions. In highly concentrated solutions, hydrogen ion concentration may exceed hydrogen ion activity because the ions may begin to interact with each other and with other ion species rather than act as independent entities. The ratio of activity to concentration is called the *activity coefficient* of the ion in that solution.

If the measured solution had a value of 7.0 pH instead of 6.0 pH, there would be no voltage generated across the glass membrane since the two solutions' hydrogen ion activities would be equal. Having a solution with one decade (ten times more: exactly one "order of magnitude") greater hydrogen ions activity than the internal buffer solution produces 59.17 millivolts at 25 degrees Celsius. If the pH were to drop to 5.0 (two units away from 7.0 instead of one unit), the output voltage would be double: 118.3 millivolts. If the solution's pH value were more alkaline than the internal buffer (for example, 8.0 pH), the voltage generated at the glass bulb would be the opposite polarity (e.g. 8.0 pH = -59.17 mV ; 9.0 pH = -118.3 mV, etc.).

The following table shows the relationship between hydrogen ion activity, pH value, and probe voltage¹⁶:

Hydrogen ion activity	pH value	Probe voltage (at 25 °C)
$1 \times 10^{-3} M = 0.001 M$	3.0 pH	236.7 mV
$1 \times 10^{-4} M = 0.0001 M$	4.0 pH	177.5 mV
$1 \times 10^{-5} M = 0.00001 M$	5.0 pH	118.3 mV
$1 \times 10^{-6} M = 0.000001 M$	6.0 pH	59.17 mV
$1 \times 10^{-7} M = 0.0000001 M$	7.0 pH	0 mV
$1 \times 10^{-8} M = 0.00000001 M$	8.0 pH	-59.17 mV
$1 \times 10^{-9} M = 0.000000001 M$	9.0 pH	-118.3 mV
$1 \times 10^{-10} M = 0.0000000001 M$	10.0 pH	-177.5 mV
$1 \times 10^{-11} M = 0.00000000001 M$	11.0 pH	-236.7 mV

This numerical progression is reminiscent of the *Richter scale* used to measure earthquake magnitudes, where each ten-fold (decade) multiplication of power is represented by one more increment on the scale (e.g. a 6.0 Richter earthquake is ten times more powerful than a 5.0 Richter earthquake). The logarithmic nature of the Nernst equation means that pH probes – and in fact all potentiometric sensors based on the same dynamic of voltage produced by ion exchange across a membrane – have astounding rangeability: they are capable of representing a wide range of conditions with a modest signal voltage span.


Of course, the disadvantage of high rangeability is the potential for large pH measurement errors if the voltage detection within the pH instrument is even just a little bit inaccurate. The problem is made even worse by the fact that the voltage measurement circuit has an extremely high impedance due to the presence of the *glass* membrane¹⁷. The pH instrument measuring the voltage produced by a pH probe assembly must have an input impedance that is orders of magnitude greater yet, or else the probe's voltage signal will become "loaded down" by the voltmeter and not register accurately.

¹⁶The mathematical sign of probe voltage is arbitrary. It depends entirely on whether we consider the reference (buffer) solution's hydrogen ion activity to be C_1 or C_2 in the equation. Which ever way we choose to calculate this voltage, though, the polarity will be opposite for acidic pH values as compared to alkaline pH values

¹⁷Glass is a very good insulator of electricity. With a thin layer of glass being an essential part of the sensor circuit, the typical impedance of that circuit will lie in the range of *hundreds* of mega-ohms!


Fortunately, modern operational amplifier circuits with field-effect transistor input stages are sufficient for this task¹⁸:

Equivalent electrical circuit of a pH probe and instrument



The voltage sensed by the pH instrument very nearly equals V_{pH} because $(R_{glass} + R_{ref}) \ll R_{input}$

Even if we use a high-input-impedance pH instrument to sense the voltage output by the pH probe assembly, we may still encounter a problem created by the impedance of the glass electrode: an RC time constant created by the parasitic capacitance of the probe cable connecting the electrodes to the sensing instrument. The longer this cable is, the worse the problem becomes due to increased capacitance:



¹⁸Operational amplifier circuits with field-effect transistor inputs may easily achieve input impedances in the *tera-ohm* range ($1 \times 10^{12} \Omega$).

This time constant value may be significant if the cable is long and/or the probe resistance is abnormally large. Assuming a combined (measurement and reference) electrode resistance of $700\text{ M}\Omega$ and a 30 foot length of RG-58U coaxial cable (at 28.5 pF capacitance per foot), the time constant will be:

$$\tau = RC$$

$$\tau = (700 \times 10^6 \Omega) ((28.5 \times 10^{-12} \text{ F/ft})(30 \text{ ft}))$$

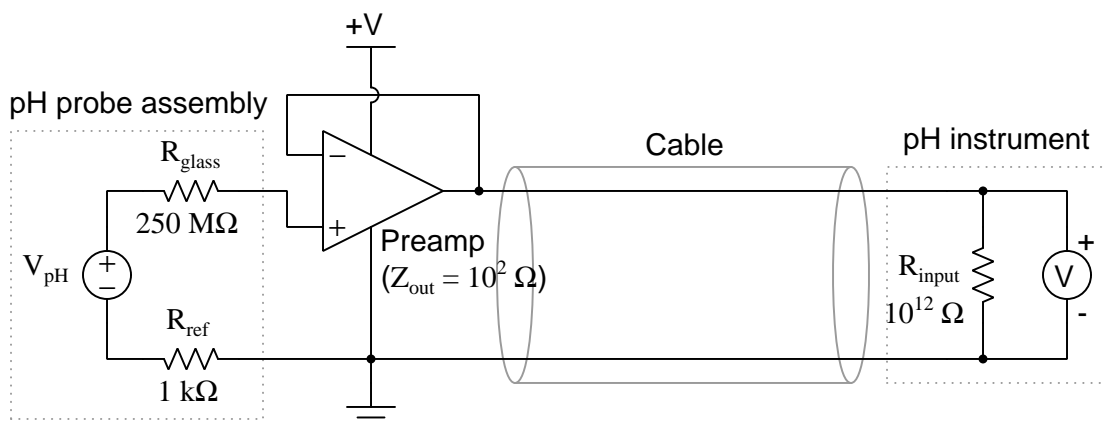
$$\tau = (700 \times 10^6 \Omega)(8.55 \times 10^{-10} \text{ F})$$

$$\tau = 0.599 \text{ seconds}$$

Considering the simple approximation of 5 time constants being the time necessary for a first-order system such as this to achieve within 1% of its final value after a step-change, this means a sudden change in voltage at the pH probe caused by a sudden change in pH will not be fully registered by the pH instrument until almost 3 seconds after the event has passed!

It may seem impossible for a cable with capacitance measured in *picofarads* to generate a time constant easily within the range of human perception, but it is indeed reasonable when you consider the exceptionally large resistance value of a glass pH measurement electrode. For this reason, and also for the purpose of limiting the reception of external electrical “noise,” it is best to keep the cable length between pH probe and instrument as short as possible.

When short cable lengths are simply not practical, a *preamplifier* module may be connected between the pH probe assembly and the pH instrument. Such a device is essentially a unity-gain (gain = 1) amplifier designed to “repeat” the weak voltage output of the pH probe assembly in a much stronger (i.e. lower output impedance) form so the effects of cable capacitance will not be as severe. A unity-gain operational amplifier “voltage buffer” circuit illustrates the concept of a preamplifier, with a typical output impedance in the hundreds (rather than *millions*) of ohms:




A preamplifier module appears in this next photograph:



The preamplifier does not boost the probes' voltage output at all. Rather, it serves to decrease the impedance (the Thévenin equivalent resistance) of the probes by providing a low-resistance (relatively high-current capacity) voltage output to drive the cable and pH instrument. By providing a voltage gain of 1, and a very large current gain, the preamplifier practically eliminates RC time constant problems caused by cable capacitance, and also helps reduce the effect of induced electrical noise. As a consequence, the practical cable length limit is extended by orders of magnitude.

Some pH probe assemblies have built-in preamplifier circuits to boost the current-sourcing ability of the probe, rather than relying on a separate preamplifier module connected between the probe and pH instrument. Preamplified pH probes have multi-conductor cables with extra wires used to conduct DC power from the pH transmitter to the pH probe to power the preamplifier:



A feature seen in the above amplified probe is an *RTD* sensor for detecting the temperature of the liquid process solution. This is important because the Nernst equation contains a term for membrane temperature, which means the Nernst potential depends just as much on temperature as it does on ionic concentration. The calculations we performed earlier predicting the amount of voltage produced by different solution pH values all assumed the same temperature: 25 degrees Celsius (298.15 Kelvin). If the solution is not at room temperature, however, the voltage output by the pH probe will not be 59.17 millivolts per pH unit. For example, if a glass measurement electrode is immersed in a solution having a pH value of 6.0 pH at 70 degrees Celsius (343.15 Kelvin), the voltage generated by that glass membrane will be 68.11 mV rather than 59.17 mV as it would be at 25 degrees Celsius. That is to say, the *slope* of the pH-to-voltage function will be 68.11 millivolts per pH unit rather than 59.17 millivolts per pH unit as it was at room temperature.

The portion of the Nernst equation to the left of the logarithm function defines this slope value:

$$\text{Nernst potential} = \frac{2.303RT}{nF} \log \left(\frac{C_1}{C_2} \right) = \frac{2.303RT}{nF} (7 - \text{pH}_1)$$

$$\text{Slope} = \frac{2.303RT}{nF}$$

Recall that R and F are fundamental constants, and n is fixed at a value of 1 for pH measurement (since there is exactly one electron exchanged for every H^+ ion migrating through the membrane). This leaves temperature (T) as the only variable capable of influencing the theoretical slope of the function.

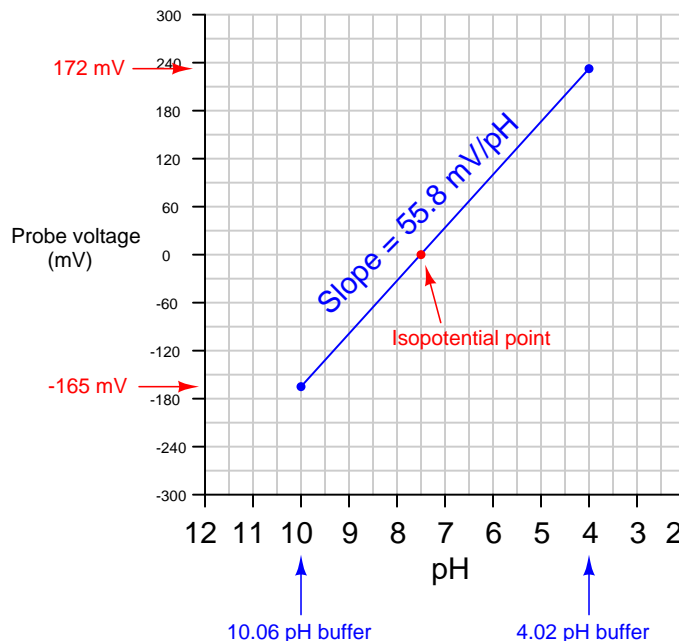
In order for a pH instrument to accurately infer a solution's pH value from the voltage generated by a glass electrode, it must “know” the expected slope of the Nernst equation. Since the only variable in the Nernst equation besides the two ion concentration values (C_1 and C_2) is temperature (T), a simple temperature measurement will provide the pH instrument the information it needs to function accurately. For this reason, many pH instruments are equipped with RTD inputs for solution temperature sensing, and many pH probe assemblies have built-in RTD temperature sensors ready to sense solution temperature.

While the theoretical slope for a pH instrument depends on no variable but temperature, the *actual* slope also depends on the condition of the measurement electrode. For this reason, pH instruments need to be calibrated for the probes they connect to.

A pH instrument is generally calibrated by performing a two-point test using *buffer solutions* as the pH calibration standard. A buffer solution is a specially formulated solution that maintains a stable pH value even under conditions of slight contamination. For more information on pH buffer solutions, see section 18.10.5 beginning on page 1302. The pH probe assembly is inserted into a cup containing a buffer solution of known pH value, then the pH instrument is calibrated to that pH value¹⁹. After establishing the first calibration point, the pH probe is removed from the buffer, rinsed, then placed into another cup containing a second buffer with a different pH value. After another stabilization period, the pH instrument is calibrated to this second pH value.

¹⁹With all modern pH instruments being digital in design, this calibration process usually entails pressing a pushbutton on the faceplate of the instrument to “tell” it when the probe has stabilized in the buffer solution. Clean and healthy pH probes typically stabilize to the buffer solution's pH value within 30 seconds of immersion.

Only two points are necessary to define a line, so these two buffer measurements are all that is required by a pH instrument to define the linear transfer function relating probe voltage to solution pH:



Most modern pH instruments will display the calculated slope value after calibration. This value should (ideally) be 59.17 millivolts per pH unit at 25 degrees Celsius, but it will likely be a bit less than this. The voltage-generating ability of a glass electrode decays with age, so a low slope value may indicate a probe in need of replacement.

Another informative feature of the voltage/pH transfer function graph is the location of the *isopotential* point: that point on the graph corresponding to zero probe voltage. In theory, this point should correspond to a pH value of 7.0 pH. However, if there exist stray potentials in the pH measurement circuit – for example, voltage differences caused by ion mobility problems in the porous junction of the reference electrode, or contamination of the buffer solution inside the glass electrode bulb – this point will shift. Sufficient contamination of the buffer solution inside the measurement electrode (enough to drive its pH value from 7.0) will also cause an isopotential point shift, since the Nernst equation predicts zero voltage when ion concentrations on both sides of the membrane are equal.

A quick way to check the isopotential point of a pH probe assembly is to short the input terminals on the pH measuring instrument together (forcing V_{input} to be equal to 0 millivolts) and note the pH indication on the instrument's display²⁰. This test should be performed *after* calibrating the instrument using accurate pH buffer solutions. The instrument characterized by the previous graph,

²⁰A more obvious test would be to directly measure the pH probe assembly's voltage while immersed in 7.0 pH buffer solution. However, most portable voltmeters lack sufficient input impedance to perform this measurement, and so it is easier to calibrate the pH instrument in 7.0 pH buffer and then check *its* zero-voltage pH value to see where the isopotential point is at.

for example, will register approximately 7.5 pH with its input terminals shorted because that is the pH value at which its probe happens to output zero millivolts.

When calibrating a pH instrument, you should choose buffers that most closely “bracket” the expected range of pH measurement in the process. The most common buffer pH values are 4, 7, and 10. For example, if you expect to measure pH values in the process ranging between 7.5 and 9, for example, you should calibrate that pH instrument using 7 and 10 buffers.

Potentiometric pH probes require regular care and servicing for long life and accurate measurement. For some pH probes, regular service includes refilling a liquid electrolyte reservoir for the reference electrode. Cleaning is a common requirement for pH electrodes in dirty process applications such as wastewater. Both the glass sensing bulb and the reference electrode must have direct contact with the process liquid, with no coating, plugging or other barriers to interfere with ion transfer. If a pH probe is dirty, cleaning should be done with no solid contact made to the glass electrode because the glass electrode is very fragile. Never use a toothbrush, a towel, or any sort of abrasive tool to clean a glass electrode! Liquid probe-cleaning solutions are manufactured to dissolve a number of different types of fouling commonly found on pH probes:

- **Fats, greases, and oils** – use a non-ionic surfactant or methanol solution
- **Proteins** – use an acidic pepsin solution
- **Minerals** – use an acidic solution
- **Sulfides** – use thiourea solution
- **Microbial growth** – use thiourea solution
- **Salts** – use deionized (distilled) water

Pressure-rinsing is a practical technique for cleaning stubborn deposits from pH probes. A hand-pump squirt nozzle using cleaning solution (or in some cases, simply deionized water) is often able to dislodge the matter from a fouled pH probe. Soaking in warm cleaning solution or deionized water is also recommended, especially for dislodging material from the small reference hole (allowing process fluid to reach the reference junction) found in many combination pH probes. Regardless of the cleaning solution used, a thorough rinsing with deionized water is recommended as a final step before returning the pH probe to service.

23.3 Chromatography

Imagine a major marathon race, where hundreds of runners gather in one place to compete. When the starting gun is fired, all the runners begin running the race, starting from the same location (the starting line) at the same time. As the race progresses, the faster runners distance themselves from the slower runners, resulting in a dispersion of runners along the race course over time.

Now imagine a marathon race where certain runners share the exact same running speeds. Suppose a group of runners in this marathon all run at exactly 8 miles per hour (MPH), while another group of runners in the race run a bit slower at exactly 6 miles per hour, and another group of runners plod along at exactly 5 miles per hour. What would happen to these three groups of runners over time, supposing they all begin the race at the same location and at the exact same time?

As you can probably imagine, the runners within each speed group will stay with each other throughout the race, with the three groups becoming further spread apart over time. The first of these three groups to cross the finish line will be the 8 MPH runners, followed by the 6 MPH runners a bit later, and then followed by the 5 MPH runners after that. To an observer at the very start of the race, it would be difficult to tell exactly how many 6 MPH runners there were in the crowd, but to an observer at the finish line with a stop watch, it would be very easy to tell how many 6 MPH runners competed in the race, by counting how many runners crossed the finish line as a distinct group at the exact time corresponding to a speed of 6 MPH.

Now imagine a mixture of chemicals in a fluid state traveling through a very small-diameter “capillary” tube filled with an inert, porous material such as sand. Some of those fluid molecules will progress more easily down the length of the tube than others, with similar molecules sharing similar propagation speeds. Thus, a small sample of that chemical mixture injected into such a capillary tube, and carried along the tube by a continuous flow of solvent (gas or liquid), will tend to separate into its constituent components (called *species*) over time just like the crowd of marathon runners separate over time according to running speed. Slower-moving molecules will experience greater *retention time* inside the capillary tube, while faster-moving molecules experience less. A detector placed at the outlet of the capillary tube, configured to detect any chemical different from the solvent, will indicate the different species exiting the tube at different times. If the retention time of each chemical species is known from prior tests, this device may be used to identify the composition of the original chemical mix (and even how much of each species was present in the injected sample) based solely on the *time delay* of each species exiting the column. More importantly, this one device will be able to identify and quantify a great many chemical compounds present in the original sample, a feat unmatched by most analytical technologies.

This is the essence of *chromatography*: the technique of chemical separation by time-delayed travel down the length of a stationary medium (called a *column*). In chromatography, the chemical solution traveling down the column is called the *mobile phase*, while the solid and/or liquid substance residing within the column is called the *stationary phase*. Chromatography was first applied to chemical analysis by a Russian botanist named Mikhail Tswett, who was interested in separating mixtures of plant pigments. The colorful bands left behind in the stationary phase by the separated pigments gave rise to the name “chromatography,” which literally means “color writing.”


Modern chemists often apply chromatographic techniques in the laboratory to purify chemical samples, and/or to measure the concentrations of different chemical substances within mixtures. Some of these techniques are manual (such as in the case of *thin-layer chromatography*, where liquid

solvents carry liquid chemical mixtures along a flat plate covered with an inert coating such as alumina, and the positions of the chemical drops after time distinguishes one chemical species from another). Other techniques are automated, with machines called *chromatographs* performing the timed analysis of chemical travel through tightly-packed tubular columns. The main focus of this section will be automated chromatography, as is used for continuous process analysis.

23.3.1 Manual chromatography methods

One of the manual chromatography methods taught to beginning chemists is *thin-layer chromatography*, also known as *TLC*. An illustrated sequence showing thin-layer chromatography appears here:

Thin-layer chromatography




As solvent wicks up the surface of the plate, it carries along with it all components of the sample spot. Each component travels at a different speed, separating the components along the plate over time.

The simplest forms of chromatography reveal the chemical composition of the analyzed mixture as residue retained by the stationary phase. In the case of thin-layer chromatography, the different liquid compounds of the mobile phase remain embedded in the stationary phase at distinct locations after sufficient “developing” time. The same is true in *paper-strip chromatography* where a simple strip of filter paper serves as the stationary phase through which the mobile phase (liquid sample and solvent) travels: the different species of the sample remain in the paper as residue, their relative positions along the paper’s length indicating their extent of travel during the test period. If the chemical species happen to have different colors, the result will be a stratified pattern of colors on the paper strip²¹.

²¹This effect is particularly striking when paper-strip chromatography is used to analyze the composition of *ink*. It is really quite amazing to see how many different colors are contained in plain “black” ink!

23.3.2 Automated chromatographs

The most common type of chromatography used in continuous process analysis is the *gas chromatograph* (abbreviated “GC”), so named because the mobile phase is a gas (or a vapor²² rather than a liquid. In a GC, the sample to be analyzed is injected at the head of a very long and very narrow tube packed with solid and/or liquid material²³. This long and narrow tube (called the *column*) is designed to impede the passage of the sample molecules. A continuous flow of “carrier gas” washes the sample compounds down the length of the column, allowing them to separate over time according to how they interact with the stationary phase packed inside the column. A simplified schematic of a process GC shows how it functions:



The *sample valve* periodically injects a very precise quantity of sample into the entrance of the column tube and then shuts off to allow the constant-flow *carrier* gas to wash this sample through the length of the column tube. Each chemical species within the sample travels through the column at different rates, exiting the column at different times. All the *detector* needs to do is be able to tell the difference between pure carrier gas and carrier gas mixed with anything else.

An important detail to note here is that the detector need not be discriminatory in its response to different chemical compounds, with just one exception: it need only distinguish between *carrier* and *non-carrier* compounds. The ability for an analyzer to measure singular compounds out of a mixture is the fundamental challenge of all analytical instrumentation, and here we see the genius of chromatography: the variable we use to discriminate between different chemical compounds within the sample is *retention time* through the column, and nothing else. Rather than rely on some clever


²²Gas chromatographs are commonly used for industrial analysis on liquid sample streams, by using a heater at the inlet of the chromatograph to vaporize the liquid sample prior to analysis. In such applications the column and sample valve(s) must be maintained in a heated condition as well so that the sample does not condense back into liquid form during the analysis.

²³Stationary phase material used in many hydrocarbon GC's looks much like oily sand.

sensor technology to selectively detect one chemical compound out of a mixture, we are able to use a non-specific²⁴ sensor and let *time* be the discriminating variable.

Appealing to the marathon analogy again, it's as if a platform weigh-scale is situated at the finish line, to weigh the runners as they finish the race. The scale cannot discern the running speed of any runner by their measured weight – indeed, all the scale indicates is that *someone* is crossing the finish line – but it can determine the running speed based on *when* the runner steps on the scale.

If we plot the response of a chromatograph's detector on a graph, we see a pattern of peaks, each one indicating the departure of a species “group” (i.e. a different chemical compound) exiting the column. This graph is typically called a *chromatogram*:



Narrow peaks represent compact bunches of molecules all exiting the column at nearly the same time, analogous to a group of runners crossing the finish line running shoulder-to-shoulder, all stepping on and off the platform scale at the same time. Wide peaks represent more diffuse groupings of similar (or identical) molecules, analogous to a group of equally-fast runners crossing the finish line as a diffuse group. In this chromatogram, you can see that species 4 and 5 are not clearly differentiated over time. Better separation of species may be achieved by altering the sample volume, carrier gas flow rate, carrier gas pressure, type of carrier gas, column packing material, and/or column temperature.

²⁴This is not to say that one *cannot* use a selective sensor as a chromatograph detector. It's just that selectivity between different process compounds is not a *necessary requirement* for a chromatograph detector.

23.3.3 Species identification

Recall that the purpose of any analytical instrument is to discriminately measure *one* variable or component concentration amongst a mix of different chemical components. The fundamental challenge of chemical analysis is how to reliably and accurately detect that one “species” of chemical while being unaffected by the presence of any other chemical species. Most analyzers rely on a sensor that exploits some unique characteristic of the chemical of interest (e.g. its unique ability to pass through a special membrane, its ability to absorb a unique spectrum of light, its tendency to engage in a unique chemical reaction, etc.), but not chromatographs. Chromatographs simply delay the flow rate of different chemical species through a long tube based on those species’ mass and interactions with the stationary phase, using *time* as the discriminating variable. This allows chromatographs to be nearly universal in the range of chemical species they can measure, while using relatively simple and uncomplicated detectors. The detector need only discriminate between carrier gas and anything that isn’t carrier gas in order to serve its purpose in a gas chromatograph. Identifying a suitable detector technology is relatively easy given that fact that the type of carrier gas used in a GC is something *chosen*²⁵ by the designer of the GC. In other words, chromatography gives one the freedom to choose a compound that’s easy for a sensor to detect, rather than try to find a sensor able to detect the specific species one needs to measure in the process. This is important given the fact that there are many more types of chemical compounds in the world than there are chemical sensor types!

In general, chemical compounds having less molecular weight tend to exit the column earlier, while compounds having greater molecular weight exit the column later. The precise sequence of species elution through a column depends greatly on the stationary phase material type, as well as the carrier fluid type. Proper selection of stationary phase and carrier compounds is essential to efficient chromatograph operation, and is usually the domain of trained chemists.

To view a flip-book animation showing how a chromatograph separates a sample into its constituent species, turn to Appendix [A.6](#) beginning on page [3096](#).


Since species identification in a chromatograph is performed with time as the discriminating factor, a chromatograph’s ability to accurately identify chemical compounds depends on its control computer “knowing” when to expect various compounds to exit the column. Chromatographs are calibrated by injecting a sample containing known concentrations of the species of interest (as well as any potentially interfering species), then timing the beginnings and ends of each species peak as each substance exits the column.

²⁵It should be noted that the choice of carrier for any chromatography system, be it manual or automated, is not completely arbitrary. There are some limitations to which carrier fluids may be used, depending on the expected composition of the sample (e.g. you would not want to use a carrier that reacted chemically with any species in the sample so as to alter the sample’s composition!). However, the range of choices afforded to the person designing the chromatograph system lends a unique flexibility to this type of chemical analysis.

23.3.4 Chromatograph detectors

Several different detector designs exist for process gas chromatographs. The two most common are the *flame ionization detector* (FID) and the *thermal conductivity detector* (TCD). Other detector types include the *flame photometric detector* (FPD), *photoionization detector* (PID), *nitrogen-phosphorus detector* (NPD), and *electron capture detector* (ECD). All chromatograph detectors exploit some physical difference between the solutes²⁶ and the carrier gas itself which acts as a gaseous solvent, so that the detector may be able to detect the passage of solute molecules among carrier molecules.

Flame ionization detectors work on the principle of ions liberated in the combustion of the sample species. Here, the assumption is that sample compounds will ionize inside of a flame, whereas the carrier gas will not. A permanent flame (usually fueled by hydrogen gas which produces negligible ions in combustion) serves to ionize any gas molecules exiting the chromatograph column that are not carrier gas. Common carrier gases used with FID sensors are helium and nitrogen, which also produce negligible ions in a flame. Molecules of sample encountering the flame ionize, causing the flame to become more electrically conductive than it was with only hydrogen and carrier gas. This conductivity causes the detector circuit to respond with a measurable electrical signal. A simplified diagram of an FID is shown here:



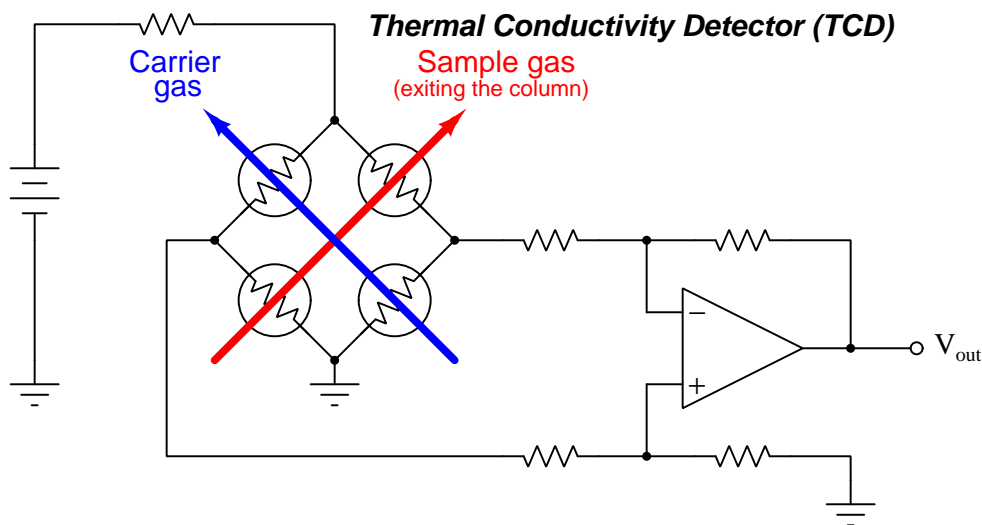
Hydrocarbon molecules happen to easily ionize during combustion, which makes the FID sensor well-suited for GC analysis in the petrochemical industries where hydrocarbon composition is the

²⁶A "solute" being one of the sample species dissolved within the carrier gas

most common form of analytical measurement²⁷. It should be noted, however, that not all carbon-containing compounds significantly ionize in a flame. Examples of non-ionizing organic compounds include carbon monoxide, carbon dioxide, and carbon sulfide. Other gases of common industrial interest such as water, hydrogen sulfide, sulfur dioxide, and ammonia likewise fail to ionize in a flame and thus are undetectable using an FID.

Thermal conductivity detectors work on the principle of heat transfer by convection (gas cooling). Here, the assumption is that sample compounds will have different thermal properties than the carrier gas. Recall the dependence of a thermal mass flowmeter's calibration on the specific heat²⁸ value of the gas being measured. This dependence upon specific heat meant that we needed to know the specific heat value of the gas whose flow we intend to measure, or else the flowmeter's calibration would be in jeopardy. Here, in the context of chromatograph detectors, we exploit the impact specific heat value has on thermal convection, using this principle to detect compositional change for a constant-flow gas rate. The temperature change of a heated RTD or thermistor caused by exposure to a gas mixture with changing specific heat value indicates when a new sample species exits the chromatograph column.

A simplified diagram of a TCD is shown here, with pure carrier gas cooling two of the self-heated thermal sensors and sample gas (mixed with carrier gas, coming off the end of the column) cooling the other two self-heated sensors. Differences in thermal conductivity between gas exiting the column versus pure carrier gas will cause the bridge circuit to unbalance, generating a voltage signal at the output of the operational amplifier circuit:



This type of chromatograph detector works best, of course, when the carrier gas has a significantly different specific heat value than any of the sample compounds. For this reason, hydrogen or helium

²⁷In fact, FID sensors are sometimes referred to as *carbon counters*, since their response is almost directly proportional to the number of carbon atoms passing through the flame.

²⁸See section 22.7.2 beginning on page 1729. The greater the specific heat value of a gas, the more heat energy it can carry away from a hot object through convection, all other factors being equal.

(both gases having very high specific heat values compared to other gases) are the preferred carrier gases for chromatographs using thermal conductivity detectors.

23.3.5 Measuring species concentration

While *time* is the variable used by a chromatograph to discriminate between different species (types) of chemical compounds, the concentration (quantity) of each chemical compound in the sample is inferred from the magnitude of the detector's response as it senses each compound. Specifically, the quantity of each species present in the sample may be determined by applying the calculus technique of *integration* to each chromatogram peak, calculating the area underneath each curve. The vertical axis represents detector signal strength, which is proportional to the rate at which detected molecules are flowing through the sensor at any given time. This means the height of each peak represents mass flow rate of each species (W , in units of micrograms per minute, or some similar units) passing through the detector. The horizontal axis represents time, so therefore the integral (sum of infinitesimal products) of the detector signal over the time interval for any specific peak (time t_1 to t_2) represents a mass quantity that has passed through the column. In simplified terms, a mass flow rate (micrograms per minute) multiplied by a time interval (minutes) equals mass in micrograms:

$$m = \int_{t_1}^{t_2} W dt$$

Where,

m = Mass of sample species in micrograms


W = Instantaneous mass flow rate of sample species in micrograms per minute

t = Time in minutes (t_1 and t_2 are the interval times between which total mass is calculated)

As is the case with all examples of integration, the unit of measurement for the totalized result is the *product* of the units within the integrand: flow rate (W) in units of micrograms per minute multiplied by increments of time (dt) in the unit of minutes, summed together over an interval ($\int_{t_1}^{t_2}$), result in a mass quantity (m) expressed in the unit of micrograms. Integration is really nothing more than the sum of products, with dimensional analysis working as it does with any product of two physical quantities:

$$\left(\frac{[\mu\text{g}]}{[\text{min}]} \right) [\text{min}] = [\mu\text{g}]$$

This mathematical relationship may be seen in graphical form by shading the area underneath the peak of a chromatogram:



Returning to the marathon analogy again, imagine the platform scale at the finish line being the detector, the race course being the chromatograph column, and the runners being individual molecules (each type of molecule running at its own unique speed). If we wish to quantify how many runners of a certain speed were in the race, we would need to integrate the scale's weight reading during the time period in which we expect that class of runners to arrive at the finish line. If a group of runners cross side-by-side such that they all step on the scale and leave the scale simultaneously, the scale's indication will be a tall and narrow peak when plotted on a time-based graph. If an identical-size group of runners arrives at the finish line at some other time, but cross in single-file form rather than side-by-side, the scale will register each runner's weight one at a time with the graphed response being a much shorter and much wider peak. In either group, though, it's the same number of runners finishing. Therefore both the height of the peak (i.e. scale weight) and the duration of the peak need to be taken into consideration when calculating how many runners were in a particular group.

Precise quantitative measurement of species concentration, however, is a bit more complicated than simply integrating detector peaks over time. A problem we must deal with in chromatography is how a detector will respond differently to different chemical compounds. Recall that the singular qualification for a gas chromatograph detector is its ability to detect species different from than the carrier gas: so long as the detector is able to sense the passage of compounds *other* than carrier gas molecules, it is usable as a chromatograph detector. This non-specificity of the detector is what makes chromatography such a versatile method for chemical analysis. However, the simple requirement of being able to detect things other than carrier gas does not mean a detector will detect

all non-carrier substances *equally*. In fact, the opposite is almost always true: each detector type responds differently to different chemical compounds, which means the accumulated area underneath a peak on a chromatogram is not a pure representation of species mass, but rather it is a product of species mass and detector sensitivity to that species.

A flame ionization detector (FID), for example, responds more strongly to a given mass flow rate of butane (C_4H_{10}) than it does for the same mass flow rate of methane (CH_4), due to the greater carbon-per-mass ratio of butane. This means the “raw” chromatogram will reveal a peak of greater area for butane than it will for the same mass concentration of methane, simply because the detector is naturally more sensitive to the former compound than to the latter compound.

Thermal conductivity detectors (TCD) are not immune to this problem either. TCDs exhibit varying responses to species having different specific heat and thermal conductivity coefficients. That is to say, a given mass flow rate of butane through a TCD will yield a different level of response than the same mass flow rate of methane, simply because these two compounds exhibit different thermal properties.

The inconsistent response of a chromatograph detector to different sampled species is not as troubling a problem as one might think, though. Since chromatographs operate on the principle of separating each species from the other over time, the chromatograph’s control computer is “aware” of which chemical species is passing through the detector at any given time. This fact allows us to program the computer with a set of pre-determined “response factors” describing how sensitive the detector is known to be for each species, allowing the computer to scale its interpretation of the detector’s signal during each time period when a known species exits the column. Thus, the chromatograph is able to calculate true mass concentration values based on measured peak areas and pre-programmed response factors.

For example, knowing that a flame ionization detector (FID) is more sensitive to butane than it is to methane means the response factor for butane must be programmed such that it “scales down” the calculated mass based on the detector’s raw signal during the time butane is expected to pass through it. In other words, the chromatograph’s computer “shifts gears” for each species exiting the column, during those times each expected species exits the column.

These detector response factors are determined by passing a mixture of known species proportions (a “calibration gas” sample) through the chromatograph, while programming the chromatograph computer with the known concentrations of this calibration gas. As the chromatograph separates and measures each compound, it records the detector response to each one, calculating response factors to accurately equate detector response to the known concentrations in the calibration gas (usually scaled in percent or parts-per-million). From then on, it will apply these “response factor” values to the raw detector signal, converting the detector’s signal levels into equivalent concentration units.

Most industrial process chromatographs come equipped with self-calibration ability, whereby calibration gas bottles stand connected to solenoid valves so that calibration gas may be directed to the analyzer on a pre-programmed schedule, with the chromatograph’s microprocessor controller managing these calibration cycles all on its own. This allows for frequent re-calibration of the chromatograph, to maintain its measurement accuracy over time with limited technician labor or oversight.

23.3.6 Industrial applications of chromatographs

Since process chromatographs have the ability to independently analyze the quantities of multiple species within a chemical sample, these instruments are inherently *multi-variable*. A single analog output signal (e.g. 4-20 mA) would only be able to transmit information about the concentration of any one species (i.e. any one peak) in the chromatogram. This is perfectly adequate if only one species concentration is of interest in the process²⁹, but some form of multi-channel digital (or multiple analog outputs) transmission is necessary to make full use of a chromatograph's ability. Legacy chromatographs have multiple 4-20 mA analog output channels (one for each compound), while most modern chromatographs provide the option of digital bus communication (e.g. Modbus, FOUNDATION Fieldbus, etc.) to transmit data on multiple species concentrations to indicators, recorders, and/or controllers.

All modern chromatographs are “smart” instruments, containing one or more digital computers executing the calculations necessary to derive precise measurements from chromatogram data. The computational power of modern chromatographs may be used to further analyze the process sample, beyond simple determinations of concentration or quantity. Examples of more abstract analyses include approximate octane value of gasoline (based on the relative concentrations of several species), or the heating value of natural gas (based on the relative concentrations of methane, ethane, propane, butane, carbon dioxide, helium, etc. in a sample of natural gas).

A very common industrial application of chromatographs is the monitoring and control of *separation* processes such as distillation (or “fractionation”) columns. The purpose of any separation process is to take a mixture or a solution and force some of its constituent compounds apart into different fluid streams. The ability of a chromatograph to measure multiple species within a sample makes it ideally suited for the task of quantifying the purity of the separated species exiting the separation process. For example, a chromatograph may be used to analyze the purity of alcohol output by the fractionation column in a distillery, quantifying alcohol concentration, water concentration, and even the concentrations of various aromatic and flavoring compounds within the distillate fluid. That data may be then used to alter some of the controlled parameters of the fractionation process such as feed flow rate, pressure, temperature gradient, etc. to achieve the desired product composition.

Another industrial application of chromatography is the monitoring and control of chemical *reaction* processes. Once again we have a single instrument able to measure the concentration of desired product exiting the reaction, as well as unprocessed reactant species and also undesired products in the same product stream. This data may then be used to control parameters in the chemical reactor to optimize the reaction taking place there.

It should be noted that although *gas* chromatography (GC) is far more prevalent in online industrial process analysis than liquid chromatography (LC), this does not mean the GC technique is limited to the analysis of process fluids existing in the gas phase alone. Gas chromatographs are

²⁹It is not uncommon to find chromatographs used in processes to measure the concentration of a single chemical species, even though the device is capable of measuring the concentrations of multiple species within that process stream. In those cases, chromatography is (or was at the time of installation) the most practical analytical technique to use for quantitative detection of that substance. Why else use an inherently multi-variable analyzer when you could have used a single-variable technology that was simpler? By analogy, it is possible to use a Coriolis flowmeter to measure nothing but fluid density, even though such a device is fully capable of measuring fluid density *and* mass flow rate *and* temperature.

often used to analyze the composition of liquid process samples, by first boiling that liquid sample within the analyzer so it may be analyzed in gaseous form. This means many of the species within the GC must be operated at temperatures exceeding the boiling point of the lowest-boiling-point substance in the sample. While this poses certain technical challenges, it is nevertheless common practice in many industries.

The following photograph shows a gas chromatograph (GC) used to determine the heating value³⁰ of natural gas at a natural gas pipeline compression facility. The entire instrument, from floor level to the top of the black box enclosing the chromatograph's column, is about 6 feet in height:



This particular GC is used by a natural gas distribution company as part of its pricing system. The heating value of the natural gas is used as data to calculate the selling price of the natural gas (dollars per standard cubic foot), so the customers pay only for the actual benefit of the gas (i.e. its ability to function as a fuel) and not just volumetric or mass quantity. No chromatograph can directly measure the heating value of natural gas, but the analytical process of chromatography can determine the relative concentrations of compounds within the natural gas. A computer, taking those concentration measurements and multiplying each one by the respective heating value of each compound, derives the gross heating value of the natural gas.

Although the column cannot be seen in this photograph of the GC, several high-pressure steel “bottles” may be seen in the background holding carrier gas used to wash the natural gas sample through the column.

³⁰Additionally, the data collected by this GC is used to improve the flow-measurement accuracy of their AGA3 honed-run orifice meters. By measuring the concentrations of different compounds in the natural gas, the GC tabulates an average density for the gas, which is then sent to the flow computer to achieve better flow-measuring accuracy than would be possible without this compensating measurement.

A typical gas chromatograph column appears in the next photograph. It is nothing more than a stainless-steel tube packed with an inert, porous filling material:




This particular GC column is 28 feet long, with an *outside* diameter of only 1/8 inch (the tube's inside diameter is even less than that). Column geometry and packing material vary greatly with application. The many choices intrinsic to column design are best left to specialists in the field of chromatography, not the average technician or even the average process engineer.


23.3.7 Chromatograph sample valves

Arguably, the component most critical to measurement accuracy in a gas chromatograph is the sample valve. Its purpose is to inject the exact same sample quantity into the column at the beginning of each cycle. If the sample quantity is not repeatable, the measured quantities exiting the column will change from cycle to cycle even if the sample composition does not change. If the valve's cycle time is not repeatable, species separation efficiency will vary from cycle to cycle. If the sample valve leaks such that a small flow rate of sample continuously enters the column, the result will be an altered "baseline" signal at the detector (at best) and total corruption of the analysis (at worst). Many process chromatograph problems are caused by irregularities in the sample valve(s).

A photograph of the column (the coil of fine tubing about 6 inches in diameter, on the left) and sample valve (stainless-steel cylinder with several tubes entering and exiting, on the right) for a gas chromatograph appears here:




A common form of sample valve uses a rotating element to switch port connections between the sample gas stream, carrier gas stream, and column:



Three slots connect three pairs of ports together. When the rotary valve actuates, the port connections switch, redirecting gas flows.


Connected to a sample stream, carrier stream, and column, the rotary sample valve operates in two different modes. The first mode is a “loading” position where the sample stream flows through a short length of tubing (called a *sample loop*) and exits to a waste discharge port, while the carrier gas flows through the column to wash the last sample through. The second mode is a “sampling” position where the volume of sample gas held in the sample loop tubing gets injected into the column by a flow of carrier gas behind it:



The purpose of the sample loop tube is to act as a holding reservoir for a fixed volume of sample gas. When the sample valve switches to the sample position, the carrier gas will flush the contents of the sample loop toward the column. This valve configuration guarantees that the injected sample volume cannot vary even if the sample valve's actuation is not precise. The sample valve need only remain in the “sampling” position long enough to completely flush the sample loop tube, and the proper volume of injected sample gas is guaranteed. An analogy for the sample loop is that of a

measuring cup held underneath a continuously-spilling stream of water: so long as the cup is held beneath the stream long enough to completely fill, it is guaranteed to deliver a fixed volume of water when removed from the stream and emptied. Like a measuring cup, the sample loop need only be filled completely with sample gas to deliver a fixed and unerring volume of gas to the chromatograph column when switched from the “loading” position to the “sampling” position.

While in the loading position, the stream of gas sampled from the process continuously fills the sample loop and then exits to a waste port. This may seem wasteful but in fact is quite essential for practical sampling operation. The volume of process gas injected into the chromatograph column during each cycle is so small (typically measured in units of *microliters*!) that a continuous flow of sample gas to waste is necessary to purge the impulse line connecting the analyzer to the process and thereby ensure a fresh sample, which in turn is necessary for the analyzer to obtain analyses of current conditions. If it were not for the continuous flow of sample to waste, it would take a *very long time* for a sample of process gas to make its way through the long impulse tube to the analyzer to be sampled, resulting in grossly delayed measurements of process conditions!




23.3.8 Improving chromatograph analysis time

The “Achilles heel” of chromatography is the extraordinary length of time required to perform analyses, compared with many other analytical methods. Cycle times measured in the range of minutes are not uncommon for chromatographs, even continuous “on-line” chromatographs used in industrial process control loops³¹! It is the basic principle of chromatography to separate chemical species using time, and so a certain amount of measurement dead time is inevitable. However, dead time in any measuring instrument is an undesirable quality. Dead time in a feedback control loop is especially bad because enough of it will cause the loop to self-oscillate.

One way to reduce the dead time of a chromatograph is to alter some of its operating parameters during the analysis cycle in such a way that it speeds up the progress of the mobile phase during periods of time where slowness of elution is not as important for fine separation of species. The flow rate of the mobile phase may be altered, the temperature of the column may be ramped up or down, and even different columns may be switched into the mobile phase stream. In chromatography, we refer to this on-line alteration of parameters as *programming*.


Temperature programming is an especially popular feature of process gas chromatographs, due to the direct effect temperature has on the viscosity of a flowing gas³². Carefully altering the operating temperature of a GC column while a sample washes through it is an excellent way to optimize the separation and time delay properties of a column, effectively realizing the high separation properties of a long column with the reduced dead time of a much shorter column:




³¹Laboratory chromatographs may take even longer to complete their analyses.

³²Whereas most liquids decrease in viscosity as temperature rises, gases *increase* in viscosity as they get hotter. In other words, a gas becomes “thicker” as it heats up, thus slowing down its progress through a chromatograph column. Since the flow regime through a chromatograph column is most definitely laminar and not turbulent, viscosity has a great effect on flow rate.


Another way to speed up the analysis time of a chromatograph is to design it with multiple columns and multiple switching valves, timing the valves so that only the fastest species travel through all columns, while slower species bypass later column stages to exit through the detector first. The alternative is to force all species to elute through all columns (or one long column), which means the minimum cycle time will be determined by the slowest species present in the sample. To use the marathon analogy again, it's like having to wait until the very last runner crosses the finish line before we can start another race to challenge faster runners. If, however, we stop the race mid-way to shuttle slow runners to the finish line (because we already know they are slow and will never win the race), we can still let the fastest runners compete the entire distance to determine who among them are the fastest, and thereby end the race sooner so we can move on to the next race:




A sequence for one type of dual-column gas chromatograph begins with the sample valve injecting a precise quantity of sample into the first column. In this illustration, the sample is comprised of 6 species labeled 1 through 6 in the order of their elution speed through the columns:



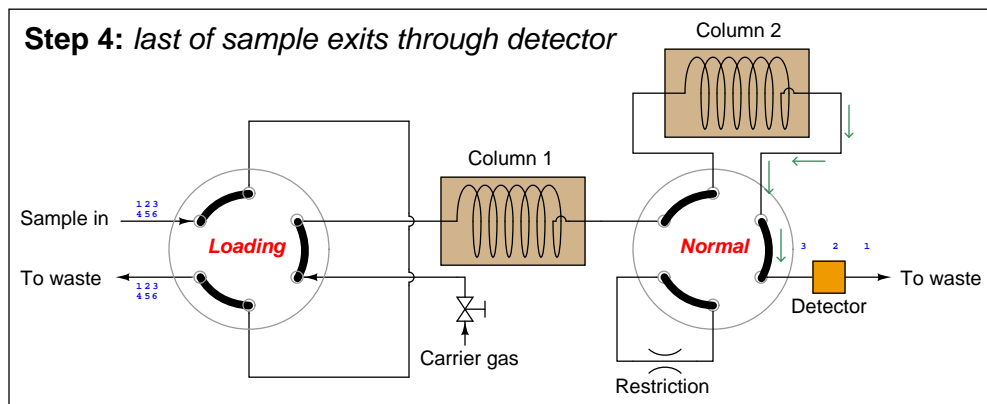
In the next step, the six species elute through column 1, with species 1 through 3 making it into the second column while species 4 through 6 are still working their way through column 1:



At this point in time, the dual-column valve switches into bypass mode, trapping the faster species (1 through 3) inside of column 2 while allowing the slower species (4 through 6) time to exit column 1 and pass through the detector:



In the last step, the dual-column valve switches back to its normal mode, allowing species 1 through 3 to elute through column 2 and pass through the detector:



The dual-column valve's timed switching from normal to bypass and back to normal again permits the slowest species to skip past the second column, while the fastest species must elute through *both* columns for maximum separation. This dual-column switching greatly reduces the total retention time of the sample without sacrificing separation of the fastest species³³.

³³Since the degree of separation between species is roughly proportional to the species' retention time, the slowest species (4, 5, and 6 in this case) do not need to go through two columns to be adequately separated. It is only the fastest species needing more retention time (through an additional column) to separate adequately from one another.

A few noteworthy points must be raised with regard to multi-column chromatographs. First, the example shown in the preceding diagrams is not the only type of multi-column chromatograph. “Trapping” a series of sample compounds inside a column is not the only way to provide different compounds with different column paths for faster separation. Some multi-column chromatographs, for example, use “backflush” valves to reverse flow through one or more columns in an effort to avoid having the slowest species elute through the entire length of those columns. This technique is used in applications where separation among compounds in the “slow” group is not important, since backflushing tends to reverse any separation that took place in the column previously.

The next point regarding multi-column chromatographs is that the dual-column valve timing must be precisely set according to known retention times of the different species inside the different columns. In the example GC shown previously, this means the retention times of the transition species (3 and 4 in this case) through the first column must be precisely known, so the dual-column valve may be switched into bypass mode after species 3 exits the first column but before species 4 exits the first column. The retention time of the slowest species (6) must also be precisely known so that the dual-column valve will not switch back to normal mode too soon and route any of that species into the second column where it would take much more time to leave the system.

A final point regarding multi-column chromatographs is that the order of species progression through the detector will *not* be fastest to slowest as with single-column chromatographs. In the dual-column GC shown previously, the slower group will exit first in order of speed (4, 5, 6), then the fastest group will exit last in order of speed (1, 2, 3).

23.4 Introduction to optical analyses

Light is known to interact with matter in very specific ways, which may be exploited as a means of measuring chemical composition. Either a sample of substance to be analyzed is stimulated into emitting light (optical *emission*), or made to absorb light from an external source (optical *absorption*). The specific frequencies (colors) of light obtained from these analyses serve to identify the chemical elements and/or compounds present in the sample, and the relative intensities of each spectral pattern indicate the concentrations of those elements and/or compounds.

The theoretical basis for optical analysis is the interaction between charged particles of matter and light, which may be modeled both as a particle (called a *photon*) and as an electromagnetic wave possessing a frequency (f) and a wavelength (λ). Thanks to the work of the physicists Max Planck and Albert Einstein at the beginning of the 20th century, we know there is a definite proportionality between the frequency of a light wave and the amount of energy carried by each photon (E). This proportionality is *Planck's constant*, or h :

$$E = hf$$

Where,

E = Energy carried by a single “photon” of light (joules)

h = Planck's constant (6.626×10^{-34} joule-seconds)


f = Frequency of light wave (Hz, or 1/seconds)

If the amount of energy carried by a photon happens to match the energy required to make an atomic electron “jump” from one energy level to another, the photon will be consumed in the work of that task when it strikes the atom. Conversely, when that same electron returns to its original (lower) energy level in the atom, it releases a photon of the same frequency as the original photon that dislodged the electron. Thus, energy is conserved (as always!): the energy received by the atom from the incident photon is later released in the form of another photon carrying the same amount of energy.

Since each element's electron configuration is unique, each element's electrons respond differently to light. Both the colors (frequencies) of light required to boost electron energy levels and the colors (frequencies) of light emitted by those atoms as their electrons fall back to their original energy levels constitute a unique “optical fingerprint” for identifying elements.

Molecular motion (i.e. temperature) is also a source of photon emission. Warm objects radiate energy in the form of electromagnetic radiation, which is why we are able to remotely measure the temperature of an object by the optical radiation it emits.


If we examine the visible light spectrum (a range of wavelengths spanning 700 nm to 400 nm, corresponding to a range of frequencies spanning 4.29×10^{14} Hz to 7.5×10^{14} Hz) emitted by a *blackbody*³⁴ heated to a temperature of 5700 Kelvin, we see a continuous spectrum of color from violet on the left (short wavelength, high frequency, high energy) to red on the right (long wavelength, low frequency, low energy). Here, I am using a computer program called *Spectrum Explorer* (SPEX) to map both the color spectrum and the intensity of radiation across a range of wavelengths:



Unless the light from a heated blackbody passes through some device to separate it into its constituent colors, the human eye blends all the colors together and only sees *white*. Thus, we use the term “white light” to refer to an equal mixture of light frequencies covering the visible spectrum. The grey areas to the far left and far right of the spectrum represent the ultraviolet and infrared regions, respectively, that lie outside of the human vision range. A blackbody heated to 5700 K emits substantial quantities of both ultraviolet and infrared radiation, but this radiation is invisible to the human eye.

³⁴In physics, a “blackbody” is a perfect emitter of electromagnetic radiation (photons) as it is heated. The intensity of light emitted as a function of wavelength (λ) and temperature (T) is $I = \frac{2\pi hc^2 \lambda^{-5}}{e^{hc/\lambda kT} - 1}$.

If we take a sample of pure hydrogen gas and heat it using an electric arc (inside a glass tube), the hydrogen atoms' electrons will be forced into higher energy states by the passage of electric current through the gas. As those electrons fall back to lower original energy levels, they emit photons of characteristic wavelengths (color). These wavelengths do *not* cover the visible spectrum as they do for blackbody objects, but rather reveal themselves as thin “lines” on the visible spectrum range, and as “peaks” on the intensity plot:



Viewed with the unaided human eye, the light emitted from a hydrogen gas discharge tube looks bright red, because that is the predominant wavelength emitted. The other colors tend to be overshadowed by the red, but we can still view them if we pass the light through a prism or through a diffraction grating to split it up into its constituent colors.

This particular set of “lines” is unique for the element hydrogen, and may serve as an identifying “fingerprint” for hydrogen if found in the emission spectrum for any chemical sample generated by the same method.

An alternative to electrically stimulating a quantity of hydrogen gas and thereby forcing an emission of specific wavelengths is to pass white light through a sample of hydrogen gas and then look for which colors are *absorbed* by the gas. As mentioned previously, photons having the necessary energies (frequencies) will become consumed in the work of elevating the hydrogen atoms' electrons to higher energy levels, leaving dark lines in an otherwise unbroken spectrum of colors from violet to red. This is called the *absorption spectrum* for an element, in contrast to the *emission spectrum* obtained by electrically energizing atoms of that element to emit light.

The following illustration shows three different spectra: the *full-color* (white light) spectrum of white light (top), the *emission* spectrum of hydrogen gas (middle), and the *absorption* spectrum of hydrogen gas (bottom). Note how the dark gaps in the absorption spectrum precisely match the positions and colors of the bright lines in the emission spectrum, because the wavelengths of light *absorbed* from white light as it passes through hydrogen gas are the exact same wavelengths *emitted* by hydrogen gas when stimulated by an electric spark in a glass tube:



The dark lines found in the absorption spectrum constitute a distinctive “fingerprint” for the element hydrogen aligned with the colored lines in hydrogen’s emission spectrum, and likewise may be used to detect the presence of hydrogen in gas samples through which white light is passed. These dark spots in the spectrum are akin to “shadows” cast by molecules of hydrogen distributed throughout the sample gas. A solid object casts a shadow whose outline represents the object’s *shape*. A chemical compounds thoroughly dissolved in a solution, whether liquid or gas, casts a shadow whose attenuated wavelengths represent the compound’s *identity*.


Usually in industrial analysis we are more concerned with the quantifiable presence of certain *compounds* in a process sample than we are in the presence of certain elements. Fortunately, molecules also have their own distinctive³⁵ interactions with light. Sometimes, these interactions take the form of molecular electrons being boosted into higher energy levels, much the same as with individual atoms. Other molecular-optical interactions take the form of *vibrations* and *rotations* set up between the atoms of a molecule, usually with photons in the infrared range³⁶. As an infrared photon of the correct wavelength (energy value) strikes an appropriate molecule, its frequency resonates with the bonded atoms, almost as if they acted as miniscule masses connected together by coil springs. This causes a transfer of energy from the photon to the molecule, where the vibration eventually dissipates that energy in the form of heat.

Thus, shining light through a sample of process gas, and analyzing the wavelengths absorbed by that gas sample, can provide quantitative measurements of the concentrations of certain gases types in that sample.

³⁵Molecules typically have much more complex interactions with light than individual atoms. The optical signatures of atoms are principally defined by electron states, light absorbed when electrons are boosted into higher-energy states and light emitted when electrons fall into lower-energy states. Molecules, on the other hand, can absorb and release energy in the inter-atomic bonds as well as in the states of individual electrons. Since molecules have more degrees of freedom with respect to optical interactions, their optical signatures tend to be much broader. This is why molecular absorption spectra consist of broad bands of wavelengths (each band comprised of many discrete lines), while atomic absorption spectra consist of relatively few lines.

³⁶These photons have wavelengths longer than 700 nm, and so have energy values too low to boost electrons into higher levels. However, the attractive bonds *between* atoms in a molecule may be subject to the energy of these infrared photons, and so may dissipate the photons’ energy and thereby attenuate a beam of infrared light.

A few different infrared absorption spectra³⁷ for common industrial compounds are shown here, with the frequency shown in units of *wavenumber* (the number of wavelengths per centimeter³⁸). It should be noted that these absorption spectra are not drawn in scale to each other; rather, they are each drawn to their own scale to better show the relative sizes of the different absorption “dips” across the spectrum for each substance:



Note that the pattern of each absorption spectrum is unique. Each compound tends to absorb infrared light in its own way, and these “signature” absorption patterns provide us with a means to selectively identify the presence of various compounds in a process fluid sample.

Molecule types most effective at absorbing infrared light are those comprised of different atom types, such as carbon monoxide (CO), carbon dioxide (CO₂), sulfur dioxide (SO₂), water vapor (H₂O), and oxides of nitrogen (NO_x). Molecules formed of two identical atoms such as molecular

³⁷In an absorption spectrum diagram, a non-absorbing substance results in a straight line at the 100% mark. Compounds absorbing specific wavelengths of light will produce low “dips” in the graph at those wavelength values, showing how less light (of those wavelengths) is able to pass un-absorbed through the sample to be detected at the other end. By contrast *emission* spectra are usually plotted with the characteristic wavelengths shown as high “peaks” in a graph that normally resides at 0%.

³⁸Wavenumber, being the reciprocal of wavelengths in centimeters, may be thought of in terms of *frequency*: the greater the wavenumber, the higher the frequency of the light wave (the smaller its wavelength). In order to convert wavenumber into wavelength (in microns), reciprocate the wavenumber value and multiply by 10⁴. For example, a wavenumber of 2000 cm⁻¹ is equivalent to a wavelength of 5 microns. In order to convert wavenumber into wavelength (in nanometers), reciprocate the wavenumber value and multiply by 10⁷. For example, a wavenumber of 4000 cm⁻¹ is equivalent to a wavelength of 2500 nm.

oxygen (O₂), nitrogen (N₂), and hydrogen (H₂) exhibit negligible interaction with infrared light. This is a fortuitous quality of infrared analysis, because many process monitoring applications focus specifically on the former compounds to the exclusion of the latter. Monitoring the exhaust emissions of a large combustion system, for example, is an application where the concentration(s) of CO, CO₂, SO₂, and/or NO_x are relevant but the concentration of nitrogen (N₂) is not. As with all chemical analyses, the key to selectivity is to find some property of measurement applicable only to the substance you are interested in measuring and not to any others. This how analytical instruments discriminate between the substance of interest and the other “background” substances.

Between optical emission and optical absorption, absorption analysis seems to be the more popular in modern industrial use, with optical emission analysis limited mostly to laboratory applications. One reason for this is the necessity of heating a sample to a high enough temperature for it to emit light: an energy-intensive and potentially hazardous endeavor. Absorption analyzers need only shine a beam of light through an unheated sample chamber, then measure how much of specific wavelengths were absorbed by the sample. Another important reason for the prevalence of absorption analyzers in industry is the necessity of a sophisticated computer and algorithm to sort the line spectra of substances generated in emission-type analyzers. Inventors have devised clever ways to quantify the absorption spectra of different process substances without resorting to automated pattern-matching of spectra.

In every absorption-type optical analyzer, the fundamental equation relating photon absorption to substance concentration is the *Beer-Lambert Law* (sometimes called the *Lambert-Beer Law*):

$$A = abc = \log \left(\frac{I_0}{I} \right)$$

Where,

A = Absorbance

a = Extinction coefficient for photon-absorbing substance(s)

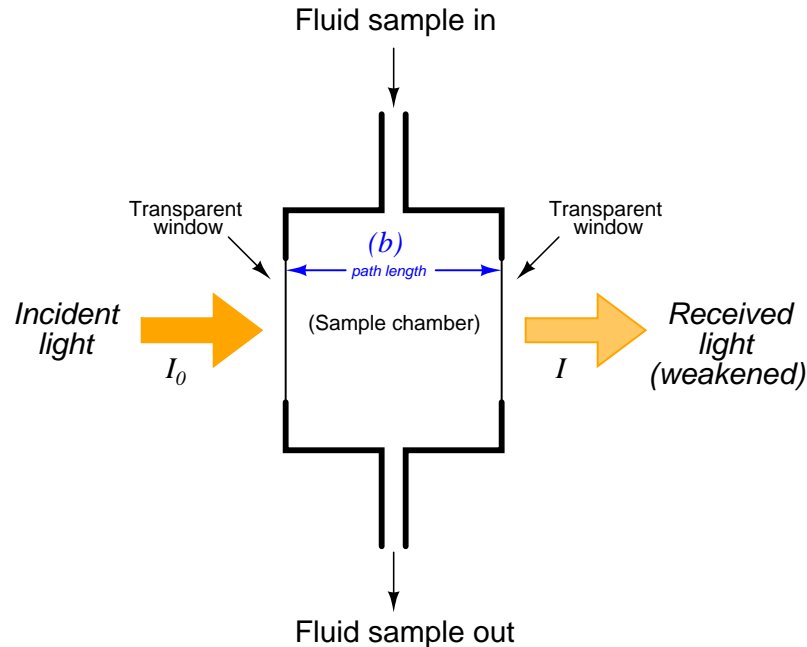
b = Path length of light traveling through the sample

c = Concentration of photon-absorbing substance in the sample

I_0 = Intensity of source (incident) light

I = Intensity of received light after passing through the sample

A typical arrangement for exposing a fluid sample (liquid or air) to light is shown in this diagram:




As indicated by the Beer-Lambert equation, greater sensitivity will be achieved with a longer path length. In some applications where the substance of interest is an atmospheric pollutant, the light beam is simply shot through open air (usually reflecting on a mirror) before returning to the instrument for analysis. If the light source happens to be a laser, the distance may be quite large³⁹.

Once the light has been passed through (or reflected off of) the process sample, it must be analyzed for attenuated wavelengths. Two major types of wavelength analysis exist: *dispersive* (where the light is split up into its constituent wavelengths) and *nondispersive* (where the spectral distribution of the wavelengths is detected without separating colors). These two optical analysis methods form the subject of the next two sections.


³⁹One such analyzer I saw in industry had a path length of a quarter-mile (1320 feet), to better measure extremely low concentrations of a gas! The gas in question was ambient air inside of a large shelter housing a chemical process. The analyzer was mounted on one side of the shelter, aiming a beam of laser light all the way to the opposite wall of the shelter 660 feet away, where a reflector was mounted. The laser beam's path length was therefore twice the length of the shelter, or 1320 feet.

23.5 Dispersive spectroscopy


The dispersion of visible light into its constituent colors goes all the way back to the 17th century with Isaac Newton's experiments, taking a glass *prism* and generating the characteristic "rainbow" of colors:



A modern variation on the theme of a solid glass prism is a thin *diffraction grating*, causing light of different wavelengths to "bend" as they pass through a series of very thin slits:



Some dispersive analyzers use a *reflection grating* instead of a refraction grating. Reflection gratings use fine lines etched on a reflective (mirror) surface to produce an equivalent dispersive effect to a diffraction grating⁴⁰:



In 1814, the German physicist Joseph von Fraunhofer closely analyzed the spectrum of colors obtained from sunlight and noticed the existence of several dark bands in the otherwise uninterrupted spectrum where specific colors seemed to be attenuated. Later that century, experiments by the French physicist Jean Bernard Léon Foucault and the German physicist Gustav Robert Kirchoff confirmed the same effect when white light was passed through a vapor of sodium. They correctly reasoned that the sun's core produced a continuous spectrum⁴¹ of light (all wavelengths) due to its


⁴⁰You may use an old compact disk (CD) as a simple reflection and refraction grating. Holding the CD with the reflective (shiny) surface angled toward you, light reflected from a bright source such as a lamp (avoid using the sun, as you can easily damage your eyes viewing reflected sunlight!) will split into its constituent colors by reflection off the CD's surface. Lines in the plastic of the CD perform the dispersion of wavelengths. You will likely have to experiment with the angle you hold the CD, pointing it more perpendicular to the lamp's direction and more angled to your eyes, before you see the image of the lamp "smeared" as a colorful spectrum. To use the CD as a diffraction grating, you will have to carefully peel the reflective aluminum foil off the front side of the disk. Use a sharp tool to scribe the disk's front surface from center to outer edge (tracing a radius line), then use sticky tape to carefully peel the scribed foil off the plastic disk. When you are finished removing all the foil, you may look *through* the transparent plastic and see spectra from light sources on the other side. Once again, experimentation is in order to find the optimum viewing angle, and be sure to avoid looking at the sun this way!

⁴¹One might wonder why the sun does not produce a line-type emission spectrum of all its constituent elements, instead of the continuous spectrum it does. The answer to this question is that emission spectra are produced only when the "excited" atoms are in relative isolation from each other, such as is the case in a low-pressure gas. In solids, liquids, and high-pressure gases, the close proximity of the atoms to each other creates many different opportunities for electrons to "jump" to lower energy levels. With all those different alternatives, the electrons emit a whole range of different wavelength photons as they seek lower energy levels, not just the few wavelengths associated with the limited energy levels offered by an isolated atom. We see the same effect on Earth when we heat metals: the electrons in a solid or liquid metal sample have so many different optional energy levels to "fall" to, they end up emitting a broad spectrum of wavelengths instead of just a few. In this way, a molten metal is a good approximation of a blackbody photon source.

incredibly high temperature, but that certain gaseous elements (including sodium) in the cooler, outer “atmosphere” of the sun were absorbing some of the wavelengths to cause the *Fraunhofer lines* in the observed spectrum. These scientists noted the same patterns of absorption (dark lines) in the sun’s spectrum that appeared in laboratory absorption tests with sodium. The implication of these scientists’ experiments are truly staggering, as they were able to correctly identify gaseous elements in a sample *93 million miles distant from Earth!*

This sort of spectrographic analysis is called *dispersive*, because it relies on a device such as a prism or diffraction grating to *disperse* the different wavelengths of light from each other so they may be independently measured.

A dispersive analyzer for process fluids would be constructed in this manner, introducing incident light to a windowed sample chamber where some wavelengths of that light would be attenuated by interaction with the process fluid molecules. In this illustration, the sample gas absorbs some of the yellow light wavelengths, resulting in less yellow light reaching the detector array:



The light source need not output white light if the wavelengths of interest do not span the entire visible spectrum. For example, if the absorption spectrum of a particular substance is known to primarily span infrared light and not the visible range, it may be sufficient to use a dispersive analyzer with an infrared light source rather than a “broad spectrum” light source covering both the infrared and visible ranges.

A necessary component of any dispersive analyzer is a computer connected to the detector array with the ability to recognize all expected emission spectra patterns, and quantify them based on the relative strengths of the detected wavelengths. This is a level of sophistication beyond most industrial measuring instruments at the time of this writing (2015), which is one reason dispersive analyzers are not as popular (yet!) for industrial process use. However, once such a computer and necessary software are in place to perform the analyses, measurement of multiple substances from the same absorption spectrum becomes possible. Like chromatographs, dispersive optical analyzers naturally function as multi-component measurement devices.

23.6 Non-dispersive Luft detector spectroscopy

Non-dispersive analysis, while newer in discovery than dispersive analysis (Isaac Newton's 17th-century prism), has actually seen far earlier application as continuous process analyzers. The basic design was developed during the years 1937-1938 by Dr. Luft and Dr. Lehrer in the laboratories of the German chemical company *I.G. Farbenindustrie*. By the end of World War II, over four hundred of these innovative instruments were in service in German chemical plants. Unlike most industrial analyzer technologies which are nothing more than adaptations of laboratory tests previously used by chemists to take manual measurements of substances, the invention of the first non-dispersive process gas analyzer embodied a wholly new analytical technique.

Industrial non-dispersive analyzers typically use either infrared or ultraviolet light sources, because most substances of interest absorb wavelengths in those regions rather than in the visible light spectrum. Non-dispersive spectroscopy using infrared light is usually abbreviated *NDIR*, while non-dispersive spectroscopy using ultraviolet light is abbreviated *NDUV* and non-dispersive spectroscopy using visible light is abbreviated *NDVIS*. Historically, *NDIR* is the more prevalent of the three technologies. Also, *gas* analysis is the more common application of non-dispersive spectroscopy in industry, as opposed to *liquid* analysis, which is why all the examples in this portion of the book assume the analysis of a process gas.

A partial listing of NDIR gas analysis applications at the I.G. Farben synthetic rubber facility in Hüls, Germany at the conclusion of World War II is shown here⁴². Note the impressive diversity of ranges and gases of interest measured by NDIR analyzers at this time in history, less than ten years following the invention of the technique:

Gas of interest	Range	Other gases present in mixture
Carbon monoxide	0 to 0.05% 0 to 0.1%	Hydrogen
Carbon monoxide	0 to 30%	Nitrogen, methane, ethane
Carbon dioxide	0 to 0.1%	Atmospheric air (nitrogen, oxygen, argon)
Carbon dioxide	0 to 0.5%	Hydrogen, methane, ethylene
Carbon dioxide	0 to 2%	Acetylene
Carbon dioxide	0 to 10%	Acetylene, ethylene
Acetylene	0 to 2%	Hydrogen, ethylene, methane, ethane
Acetylene	0 to 5%	Ethylene
Acetylene	0 to 10%	Ethylene, methane, ethane
Acetylene	0 to 40%	Ethylene, propylene, methane, ethane
Acetylene	30 to 80%	Hydrogen, ethylene, methane, ethane
Acetylene	50 to 100%	Ethylene, propylene, methane, ethane Ethylene, ethane, methane Hydrogen, ethylene, methane, ethane
Butadiene	0 to 1%	Atmospheric air (nitrogen, oxygen, argon)
Ethylene	0 to 10%	Acetylene, methane, ethane, propylene, dinitrogen dioxide
Ethylene	0 to 30%	Methane, ethane
Ethylene	0 to 40%	Carbon dioxide, chlorine, propylene, acetylene
Ethylene	40 to 80%	Ethylene, acetylene, dinitrogen dioxide
Ethylene	80 to 100%	Methane, ethane
Methane	75 to 100%	Ethylene, ethane

At a different I.G. Farben facility (in Uerdingen, Germany), an NDIR instrument was used as a safety gas detector for carbon monoxide (0 to 0.1% concentration) in open air. This was in a process area where high concentrations of carbon monoxide gas existed in the lines, and where a leak in a process line or valve posed a considerable safety hazard to personnel.

The challenge of any analytical measurement technology is how to achieve selectivity, where the analyzing instrument responds to the concentration of just one substance (one “species”) and to *no other substance(s) in the mixture*. If the substance of interest exhibits some unique physical property we can readily measure with sensors, the selectivity problem is easy to solve: just measure that one property exclusively, and no other substance will interfere.


⁴²These details taken from pages 93-94 of *Instrumentation and Control in the German Chemical Industry*, a fascinating book detailing the state-of-the-art in process instrumentation in German chemical manufacturing facilities following the war.

In the case of absorption spectrometers such as non-dispersive analyzers, the challenge is to selectively measure the concentration of certain light-absorbing substances amidst the presence of other substances also absorbing certain wavelengths of light. If the substance of interest is the *only* substance present in the mixture capable of absorbing light, selectivity is guaranteed. However, most applications in industry are not this easy, with the mixture containing other light-absorbing substances besides the one of interest. Some of these substances may absorb completely different wavelengths of light, while others may have absorption bands overlapping the absorption bands of the substance of interest (i.e. the interfering substances absorb some of the same wavelengths of light absorbed by the substance of interest, in addition to absorbing some unique wavelengths of their own).

Dispersive spectrographs achieve selectivity by “disassembling” the spectrum into individual wavelengths and measuring them one by one, but a non-dispersive analyzer must somehow distinguish different spectral responses without this “disassembly” of wavelengths. The bulk of this section is devoted to a discussion of exactly how selectivity is accomplished using the NDIR technique.

23.6.1 Single-beam analyzer

Non-dispersive analyzers employ the principle of spectrographic *absorption* to measure how much of a particular substance exists within a sample. NDIR gas analyzers shine light through a windowed sample chamber (typically called a *cell*), through which a fresh flow of process gas continually moves. Certain “species” (compounds) of gas within the sample stream absorb part of the incident light, and therefore the light exiting the cell becomes partially depleted of those wavelengths. A heat-sensitive detector placed behind the cell measures how much infrared light did not get absorbed by the sample gas. If we imagine the concentration of light-absorbing gas increasing over time, more of the infrared light entering the cell will be absorbed by the gas and converted into heat within the cell, leaving less light exiting the cell to generate heat at the detector. The simplest style of non-dispersive analyzer uses a single light source, shining continuously through a single gas cell, and eventually falling on a small thermopile (converting the received infrared light into heat, and then into a voltage signal):




This crude analyzer suffers from multiple problems. First, it is non-selective: *any* light-absorbing gas entering the sample cell reduces heat at the detector (i.e. generates less thermopile voltage), regardless of the species. It might work well enough in an application where the only light-absorbing gas in the process mixture happens to be the one gas we are interested in measuring, but most industrial analyzer applications are not like this. In most cases, our process sample contains multiple species of gases capable of absorbing light within a similar range of wavelengths, but we are only interested in measuring one of them. An example would be the measurement of carbon dioxide (CO_2) concentration in the exhaust gas of a combustion furnace: most of the gases exiting the furnace do not absorb infrared light (nitrogen, oxygen), but CO_2 gas does. However, carbon monoxide (CO), water vapor (H_2O), and sulfur dioxide (SO_2) also absorb infrared light, and are all normally present in the exhaust gas of a furnace to varying degrees. Since our crude NDIR analyzer is non-selective, it cannot differentiate between carbon dioxide and any of the other infrared-absorbing gases present in the exhaust gas.

Another significant problem with this analyzer design is that any variations in the light source’s output cause both a zero shift and a span shift in the instrument’s calibration. Since light sources tend to weaken with age, this flaw necessitates frequent re-calibration of the analyzer.

Finally, since the detector is a thermopile, its output will be affected not just by the light falling on it, but also by ambient temperature, causing the analyzer’s output to vary in ways completely unrelated to sample gas composition.

23.6.2 Dual-beam analyzer

One way to improve on the single-beam analyzer design is to split the light beam into two equal halves, then pass each half-beam through its own cell. Only one of these cells will hold the process gas to be analyzed – the other cell is sealed, containing a “reference” gas such as nitrogen that absorbs no infrared light. At the end of each cell we will place a matched pair of thermopile detectors, connecting these detectors in series-opposing fashion so equal voltages will cancel:



Let us perform some “thought experiments” on this apparatus to explore its behavior. Imagine the sample gas being a non-absorber of infrared light just like the reference gas. In this virtual experiment, the opposed detector pair will generate no voltage signal because each of the two detectors receives the same (full) amount of incident light.

Next, we will alter one of the variables in our “thought experiment” to see what difference that variable makes. Here, we imagine the sample containing some concentration of an infrared-absorbing gas while the reference gas continues to absorb no light. Now, the two thermopile detectors will receive differing intensities of infrared light, causing the series-opposed pair to be out of balance, generating a net voltage signal we can measure as an indication of light-absorbing gas concentration.

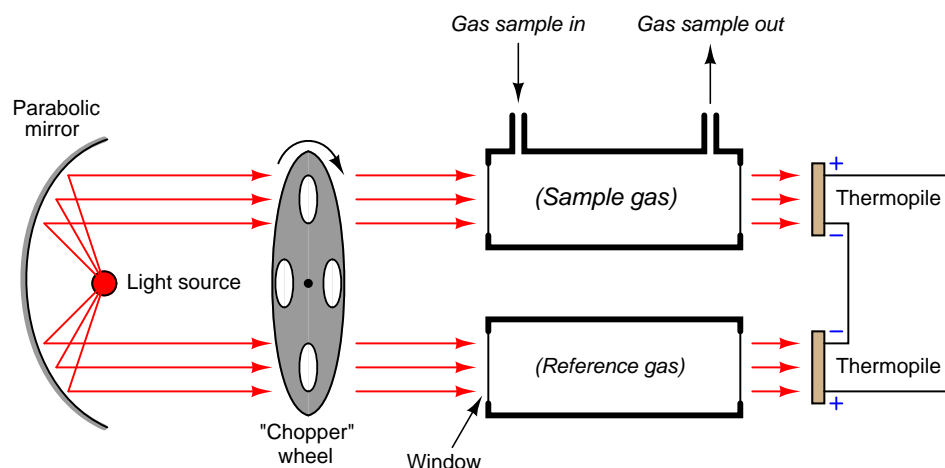
The addition of a reference gas chamber and second thermopile detector completely eliminates the ambient temperature problem seen in the single-detector apparatus. If the analyzer’s temperature happens to rise or fall, the voltages output by *both* thermopiles will rise and fall equally, canceling each other out so that the only voltage produced by the series-opposing pair will be that produced by differences in received light intensity.

The dual-detector design also eliminates the problem of “zero drift” as the light source ages. As time progresses and the light source becomes dimmer, *both* detectors see less light than before. Since the detector pair measures the difference between the two light beam intensities, any degradation common to both beams will be ignored⁴³.

⁴³There will still be a *span* shift resulting from degradation of the light source, but this is inevitable. At least with this design, the zero-shift problem is eliminated.

Another detector problem still remains, in that an imbalance will develop if one detector happens to “drift” in voltage apart from the other, so they are no longer in perfect counter-balance even with the same received light intensities. This might happen if one of the thermopiles experiences greater ambient temperature than the other, perhaps due to convective heat transfer from hot process sample gas in the nearby sample cell and not the reference cell. An ingenious solution to this problem is to insert a spinning metal “chopper” wheel in the path of both light beams, causing the light beams to *pulse* through the sample and reference cells at a low frequency (typically a few pulses per second):


Dual-beam non-dispersive analyzer



The effect of the “chopper” is to make the detector assembly output a *pulsating* (“AC”) voltage signal rather than a steady voltage signal. The peak-to-peak amplitude of this pulsating signal represents the difference in light intensity between the two detectors, but any “drift” will manifest itself as a constant or very slowly-changing (“DC”) bias voltage. The following table illustrates the detector assembly signal for three different gas concentrations (none, little, and much) both with and without a mis-match in detector signals due to thermal drift:

		Absorbing gas concentration (in sample cell)		
		None	Little	Much
Without "drift" (perfectly matched detectors)	0 V			
	0 V			
With "drift"	0 V			
	0 V			

This DC bias voltage is very easy to filter in the amplifier section of the analyzer. All we need is capacitive coupling between the detector assembly and the amplifier, and the amplifier will never “see” the DC bias voltage:




With the detector assembly producing an “AC” (pulsing) signal instead of a “DC” signal, and by using capacitive coupling to the amplifier, the electronic circuit responds only to changes in the amplitude of the AC waveform and not to its DC bias. This means the analyzer will only respond to changes in detector temperature resulting from changes in light absorbance (i.e. gas concentration), and not from any other factor such as ambient temperature drift. In other words, since the amplifier has been built to only amplify pulsing signals, and the only thing pulsing in this instrument is the light, the electronics will only measure the effects generated by the light, rejecting all other stimuli.

Despite the design improvement of the chopper wheel and AC-coupled amplifier circuit, another significant problem remains with this analyzer: it is still a non-selective instrument. *Any* light-absorbing gas entering the sample cell will cause the detector pair to generate a signal regardless of the type of gas, because the thermopile detectors respond to a broad⁴⁴ range of light wavelengths. While this may suffice for some industrial applications, it will not for most where a mixture of light-absorbing gases coexist. What we need is a way to make this instrument *selective* to just one type of gas, in order that it be a useful analyzer in a wider variety of process applications.

⁴⁴In analytical literature, you may read of some detectors having a *catholic* response. This is just a fancy way of saying the detector responds to a wide variety of things. The thermopiles shown in this NDIR instrument could be considered to have a catholic response to incident light. The word “catholic” in this context simply means “universal,” referring to the detector’s non-selectivity. Do not be dismayed if you encounter arcane terms such as “catholic” as you learn more about analytical instruments – the author is probably just trying to impress you with his or her vocabulary!

23.6.3 Luft detectors

A very clever way to achieve selectivity with a non-dispersive optical analyzer is to replace the thermopiles with a detector more sensitive to the wavelengths absorbed by the gas of interest than to the wavelengths absorbed by any other (“interfering”) gas species. Dr. Luft invented just such a detector when developing the NDIR gas analyzer for I.G. Farben in the late 1930’s. His design used two gas chambers and a thin diaphragm to measure the difference in light intensity exiting the sample and reference cells:



As light enters the dual chambers of the detector, the light absorbed by the fill gas causes those gas molecules to increase temperature⁴⁵. This expands the gas, pressing against the diaphragm. If the light intensities are equal, the pressures will be equal and no diaphragm motion will result. If the light intensities are unequal (due to the sample cell absorbing some of the wavelengths), the gas pressure developed inside that half of the Luft detector will be less, causing the thin diaphragm to bow in that direction. A set of fixed metal plates senses the diaphragm’s position using the differential capacitance technique (just like many modern differential pressure sensors). With the “chopper” wheel working to pulsate light through the sample and reference gas cells, the diaphragm motion will likewise pulsate, and the resulting “AC” pulse signal may be filtered and amplified to represent absorbing gas concentration.


What makes the Luft detector selective is that it is filled with a 100% concentration of the gas we are interested in measuring. This means only those wavelengths of light absorbed by the gas of interest will develop heat (and pressure) inside the detector chambers. Different wavelengths of light absorbed by other (“interfering”) gases in the sample will not be absorbed to the same degree (or at all) by the gas inside the Luft detector, and therefore the pressure pulses inside the Luft detector will be primarily a function of our interest-gas concentration and not of the interfering-gas concentration(s).

⁴⁵Recall that the absorption of light by an atom or a molecule causes the photon’s energy to dissipate. An absorbed photon’s energy is transferred to the molecule, often resulting in increased motion (kinetic energy), which as we know is the very definition of temperature. Increased temperature in a gas of confined volume and fixed molecular quantity must result in an increased pressure according to the Ideal Gas Law ($PV = nRT$).

The selectivity gained by a gas-filled Luft detector is not obvious to see at first, and deserves some explanation. We may explore this selective behavior in more detail by performing a set of “thought experiments” whereby we imagine the effects of different gas species on an NDIR analyzer equipped with a Luft detector.

Suppose we have an application where we intend to measure carbon dioxide concentration in a gas mixture also containing ethane. In a simple dual-beam NDIR detector using thermopile detectors, both carbon dioxide and ethane present in the sample chamber will generate a detector response, since both gas species absorb infrared light, and the thermopile detectors respond to *any* change in the amount of infrared light received at the detector. Thus, such a simple analyzer could not tell the difference between a change in carbon dioxide concentration versus a change in ethane concentration. This makes ethane an “interferent” given our goal of *only* measuring carbon dioxide concentration.

While both carbon dioxide and ethane gases absorb infrared light, they do so at different specific wavelengths. The following spectral plots show the unique infrared absorption bands for carbon dioxide and ethane, respectively. As you can see, the wavelengths of infrared light absorbed by each species of gas are unique, and do not overlap:



Let us now imagine replacing the thermopile detectors with a Luft detector, its dual chambers filled with a 100% concentration of carbon dioxide gas. If neither carbon dioxide nor ethane are present in the sample chamber, light will pass from the source undiminished to the Luft detector, causing equal heating of the CO₂ gas in both chambers and therefore zero response. This is precisely what we would expect from any dual-beam NDIR instrument, Luft detector or not.


For our next “thought experiment,” imagine carbon dioxide gas entering the sample chamber. Those carbon dioxide molecules entering the sample chamber will absorb some of the infrared light emitted by the source. Since the carbon dioxide gas molecules inside the Luft detector can only be

heated by those same wavelengths of light absorbed by the molecules in the sample chamber, the sample-side of the Luft detector will now experience less heating than before (while the reference-side experiences the same degree of heating), causing a difference in pressure inside the Luft detector and therefore generating a response. Once again, this is precisely what we would expect from any dual-beam NDIR instrument, Luft detector or not.

However, if we now imagine only *ethane* molecules entering the sample chamber, the Luft detector's response will be different from that of the thermopile detector. Surely, the ethane molecules will absorb some of the infrared light entering that chamber, but these "missing" wavelengths will be of no effect at the Luft detector because they aren't absorbed by the carbon dioxide gas inside the Luft detector chambers anyway, and therefore would not affect the temperature of the detector's carbon dioxide gas whether missing or present. In other words, the gas-filled Luft detector "doesn't care" about any wavelengths of light absorbed by gases in the sample chamber so long as the absorption pattern of the sample gas does not coincide at any point with the absorption pattern of the gas filling the Luft detector. This means the ethane's attenuation of infrared light wavelengths will be ignored by the carbon-dioxide-filled Luft detector, while carbon dioxide's attenuation of infrared light will be sensed by the Luft detector. We may now say that the instrument is "sensitized" to carbon dioxide gas, and that the Luft detector is "selective" to one species of gas over and above all other species.

If a mixture of carbon dioxide and ethane gases enters the sample chamber, each type of gas molecule will absorb its unique pattern of light wavelengths, but only the attenuation of those wavelengths matching the absorption pattern of the Luft detector's fill gas will register in the detector. Thus, the Luft detector is able to selectively measure the concentration of one gas inside the sample chamber to the exclusion of all other gases having different optical absorption patterns.

A modern variation on the Luft detector design replaces Luft's original microphone-style thin diaphragm with a narrow channel and a highly sensitive thermal flow sensor⁴⁶ connecting the two gas-filled chambers. Any difference in expansion between the gases of the two chambers when heated by light causes gas to move past the flow sensor, thus generating a signal:



As the chopper wheel pulses incident light to either chamber of the detector, gas will flow back and forth through the narrow passageway connecting the two chambers, causing an alternating flow response from the flow sensor.

The advantage of a diaphragm-less detector is that it is just as insensitive to mechanical vibration as a thermopile (having no moving parts), but retains the spectral selectivity of the traditional Luft-style detector (being filled with the gas of interest).

While Luft-style detectors greatly enhance the selectivity of non-dispersive spectrographic gas analyzers, there is still room for improvement. Perfect selectivity of measurement is assured by a Luft detector only when the light absorption spectra of the interference gas(es) do not overlap at all with the absorption spectrum of the gas of interest. If there is some overlap, interference will result.

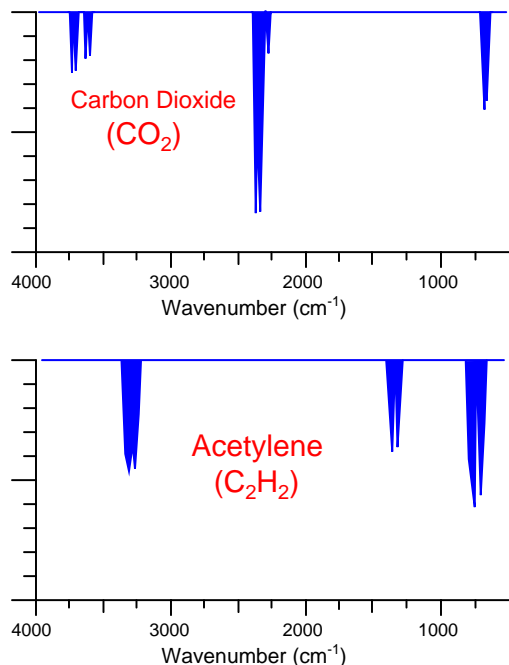
To address this concern, we will explore one more design feature of modern non-dispersive analyzers: *filter cells*.

⁴⁶The flow sensor is similar in design to thermal mass flow sensors discussed in the flow measurement chapter. See section 22.7.2 beginning on page 1726 for more information.

23.6.4 Filter cells


If other species of gas present in the sample do not absorb any of the light wavelengths absorbed by the one gas we're interested in measuring, the selectivity of a Luft detector will be total: the gas-filled detector will *only* respond to the presence of the gas we are interested in. Usually, though, process applications are not this simple. In most applications, the interfering gases have absorption spectra overlapping portions of the interest-gas spectrum. This means changes in interference gas concentration will be sensed by the detector (though not as strongly as changes in the concentration of the gas of interest) because part of the light spectrum absorbed by the interfering gas(es) will have a heating effect on the pure gas inside the detector.

An example of overlapping absorption spectra is found with the combination of carbon dioxide and acetylene gases:



As you can see, there is some common absorption between these two gas species toward the right-hand side of the scale, around 700 cm⁻¹ (approximately 14000 nm wavelength). An NDIR analyzer equipped with a Luft detector filled with 100% of carbon dioxide gas will respond strongly to concentrations of carbon dioxide gas in the sample chamber, and weakly to concentrations of acetylene gas in the sample chamber. Since acetylene does absorb some of the infrared light wavelengths absorbed by carbon dioxide, acetylene gas has the potential to affect the Luft detector and make the analyzer “think” it is measuring slightly more carbon dioxide gas than it actually is.

One more addition to our NDIR instrument helps eliminate this problem: we add two more gas cells in the path of the light beams, each one filled with 100% concentrations of the interfering gas. For this particular example, we would fill each filter cell with a 100% concentration of acetylene gas, and the Luft detector cell with a 100% concentration of carbon dioxide gas:




These *filter cells* purge the light of those wavelengths normally absorbed by the interfering gas (acetylene) inside the sample cell. As a result, the presence of that interfering gas in the sample cell will have negligible effect on the light exiting the sample cell, because those wavelengths have already been severely attenuated by the filter cells. If the filter cells happened to be 100% effective in filtering all wavelengths specific to the interfering gas, there would be absolutely none of the interfering gas's specific wavelengths of light left to be absorbed inside the sample cell, and therefore the interfering gas would have absolutely no effect on the detector's response. In other words, by filtering out all the light wavelengths absorbed by the interfering gas inside the filter cells, the presence of that same gas in the sample cell will have no effect on the detector, and therefore it can no longer interfere with the measurement of our gas of interest.

So long as our gas of interest exhibits absorption wavelengths *not shared by the interfering gas* (i.e. wavelengths of light unique to the gas of interest alone), these wavelengths will still be able to pass through the filter cells and into the sample cell where they will change intensity as the gas of interest varies in concentration. Thus, the detector now *only* responds to the gas of interest (carbon dioxide), and not to the interfering gas (acetylene).

As effective as this filtering technique is, it has the limitation of only working for one interfering gas at a time. If multiple interfering gases exist in the sample stream, we must use multiple filter cells to block those light wavelengths⁴⁷.


⁴⁷And hopefully after all this filtering we still have some (unfiltered) wavelengths unique to the gas of interest we seek to measure. Otherwise, there will be no wavelengths of light remaining to be absorbed by our gas of interest inside the sample cell, which means we will have no means of spectroscopically measuring its concentration!

A photograph showing a dual-beam NDIR analyzer appears here:




What looks like a single gold-colored gas cell is actually two cells (a divider separating the tube lengthwise into two chambers), one for the sample gas and the other for reference. Black-colored hoses pass sample gas through the bottom half of the tube, while the top half is filled with nitrogen gas (the tube connections capped and sealed with black-colored plastic). The light source and chopper assembly appears on the left-hand side of the tube, while the detector resides on the right-hand side.

In this particular instrument (a Rosemount Analytical model X-STREAM X2), the chopper wheel is driven by a stepper motor:



The head of the infrared light source appears just to the right of the chopper wheel motor.

The detector used in the X-STREAM NDIR analyzer is a modern variant of the Luft detector, with a micro-flow sensing element detecting pulses of gas flow between two chambers. In this particular analyzer the detector chambers are filled with carbon monoxide (CO) gas, to sensitize it to that species:




This instrument's maximum detection range happens to be 0 to 1000 ppm of carbon monoxide, with the ability to turn down to a range of 0 to 400 ppm.

23.7 Gas Filter Correlation (GFC) spectroscopy

Using filter cells to eliminate wavelengths associated with interfering gases is called *positive filtering* in the field of spectroscopy. You may think of this as filtering out all the wavelengths the instrument should *not* "care about." In order for positive filtering to be completely effective, the analyzer must filter out *all* wavelengths associated with all interfering species. In some applications, this may require multiple filters stacked in "series," each one filtering out wavelengths for a different interfering gas. Not only is this technique cumbersome when multiple interfering "species" are present in the sample, but it is completely useless when the interfering species are unknown.

A different filtering technique called *negative filtering* does just the opposite: placing a filter cell in the path of the light to absorb all the wavelengths associated with the gas of interest, leaving all other wavelengths unattenuated. One application of this technique is called *Gas Filter Correlation*, or *GFC* spectroscopy. This same technique is alternatively referred to as *Interference Filter Correlation*, or *IFC* spectroscopy.

Gas filter correlation analyzers use a single gas cell rather than dual cells (sample and reference), through which a light beam of alternating spectrum is passed. A rotating filter wheel creates this alternating spectrum:



The filter wheel consists of two transparent halves: one containing a high concentration of the gas of interest, and the other designed to consistently attenuate every light wavelength (i.e. the entire spectrum) emitted by the source. The attenuation factor of the “neutral” half of this filter wheel is precisely adjusted so that the same gross intensity of infrared light enters the sample gas cell at all times, regardless of the filter wheel’s position. The light detector positioned after the sample cell must be designed for non-specific response to light (i.e. not selective to certain wavelengths of light). Unlike a Luft detector, we want this detector to respond to a broad spectrum of light.

If the sample gas chamber contains nothing but non-absorbing gases, the detector will generate a steady (unchanging) signal⁴⁸ because it receives the same total light intensity during each half of the filter wheel’s rotation, albeit at different wavelengths during each half of the wheel’s rotation.

If some of our gas of interest enters the sample cell, it will begin to absorb some of the light during the time when the “neutral” filter aligns in front of the cell. During the other half of the filter’s rotation (when the light must pass through the high gas concentration chamber), our gas of interest inside the sample cell has no effect, because all those wavelengths of light have been eliminated by the filter. The result is a changing signal at the detector⁴⁹, the amplitude of oscillation proportional

⁴⁸Real GFC analyzers also have a chopper wheel preceding the filter wheel to create a pulsating light beam. This causes the detector signal to pulsate as well, allowing the analyzer to electronically filter out sensor “drift” just as in the dual-beam NDIR analyzer design. The chopper wheel has been eliminated from this diagram (and from the discussion) for simplicity. If it were not for the chopper wheel, the GFC analyzer would be prone to measurement errors caused by detector drift.

⁴⁹As previously mentioned, real GFC analyzers have a chopper wheel preceding the filter wheel to make the light beam pulse in addition to changing its spectral composition. This chopper wheel generates multiple light pulses per rotation of the filter wheel. Thus, the signal output by the detector is actually an *amplitude-modulated* waveform, with the “carrier” frequency being the chopper wheel’s pulsing and the slower “modulating” frequency being the filter wheel’s rotation cycle. Hopefully by now you see why I decided to omit the chopper wheel “for simplicity.”

to the concentration of “correlating” gas (matching the absorption spectrum of the rotating filter’s gas) inside the sample cell.

The effect of “interfering” gases in the sample cell depends on the nature of those gases. An “interfering” gas with an absorption spectrum encompassed by the absorption spectrum of the gas of interest would be indistinguishable from the gas of interest by this instrument – we would say this gas has a *positive* interference. Such a gas would absorb wavelengths of light from the beam during the time light passes through the “neutral” filter, and it would absorb no wavelengths during the time light passes through the gas filter, just like the gas of interest. A different “interfering” gas absorbing completely different wavelengths of light than our gas of interest would absorb light at all times regardless of the filter wheel’s position. However, given an equal *percentage* of absorption in a region of the spectrum untouched by the gas filter side of the wheel, but uniformly attenuated by the “neutral” side of the wheel, means that the effect of this gas would be to absorb more light during the gas-filtered part of the wheel’s rotation and less light through the “neutral-filtered” part of the wheel’s rotation – just the opposite of positive interference. Thus, a gas with an absorption spectrum wholly different from our gas of interest will have a *negative* interference effect.

In order to avoid interference of any kind from gases other than the one we are interested in measuring, the effects of positive and negative correlation interference must cancel. Fortunately for this technique, most interfering gases partially overlap spectra with most gases of interest. If the degree of overlap is approximately even, the positive and negative interferences will indeed cancel each other, resulting in little or no interference from the “interfering” gas.

To re-phrase this principle: if the absorption spectrum of a gas perfectly correlates with the spectrum for our gas of interest, the effect will be “positive,” making the analyzer think there is a greater concentration of the gas of interest than there actually is. If the absorption spectrum of a gas is perfectly *anti-correlated* with the spectrum of our gas of interest, the effect will be “negative,” making the analyzer think there is a weaker concentration of our gas of interest than there actually is. However, if the absorption spectrum of any gas is completely *uncorrelated* (i.e. random overlap of spectra) with the spectrum for our gas of interest, there interference will be neutral (little or no effect).

This makes the Gas Filter Correlation (GFC) analyzer ideally suited to distinguish gases whose spectra overlap over the same general ranges but differ in fine detail (i.e. where the individual “peaks” and “dips” in the different spectra randomly intersect). One such practical application for GFC analyzers is combustion exhaust gas analysis for carbon monoxide (CO) in the presence of carbon dioxide (CO₂) and water vapor. Unlike the dual-beam “Luft detector” style of analyzer, the GFC analyzer does not require individual filter cells for each interfering species of gas. This is a major advantage where multiple interfering gases coexist with the gas of interest.

Being a single-beam style of analyzer, GFC instruments are much easier to implement as open-air gas analyzers than dual-beam designs. Dual-beam analyzers require sample and reference cells of equal length and identical construction, in order to draw a comparison between the light passing through sample versus the same light passing through a completely non-absorbing gas. Single-beam analyzers have no need for a reference light path, and so the sample beam may be passed through open air (or through the diameter of an exhaust stack, for example) to sense gases anywhere in that region, rather than be limited to the physical length of any gas-filled cell. Recall from the Beer-Lambert law that absorbance increases in direct proportion to the path length of the light:

$$A = abc = \log \left(\frac{I_0}{I} \right)$$

Where,

A = Absorbance

a = Extinction coefficient for photon-absorbing substance(s)

b = Path length of light traveling through the sample

c = Concentration of photon-absorbing substance in the sample


I_0 = Intensity of source (incident) light

I = Intensity of received light after passing through the sample

The longer the path length, the more light will be absorbed by the gas, all other factors being equal. This increases the analyzer's sensitivity to low concentrations, which is especially important when measuring gas concentrations in the low parts-per-million (ppm) or even parts-per-billion (ppb) range.

An example diagram for a GFC analyzer used to measure gas concentrations in open air is shown here:

GFC spectrometer used in open-air measurement



Light passing through the rotating filter wheel strikes a *beam splitter* (a partially-silvered glass plate angled at 45°) where approximately half the light passes through to the sample space and the other half is lost to reflection. At the far end of the sample space, a full-silvered mirror reflects all the light back to the analyzer, where it strikes the beam splitter again, with approximately half of that light reflecting off at 90° to reach the detector. With this arrangement, the path length (b in the Beer-Lambert Law) is equal to *twice* the distance between the analyzer and the mirror, since light must travel one way to reach the mirror, then return the same distance back to the analyzer. As you might imagine, extremely long path lengths are easy to achieve with this style of open-air analyzer.

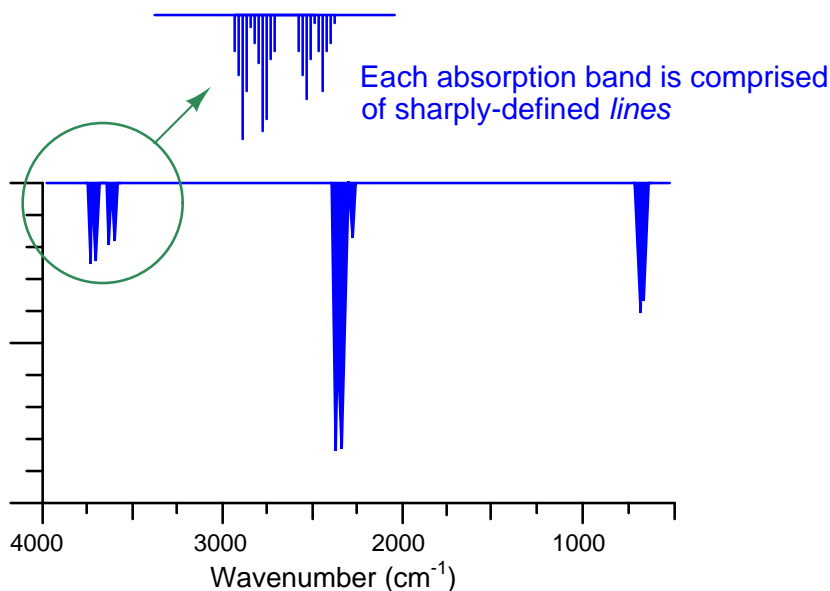
23.8 Laser spectroscopy

A *laser*⁵⁰ is a light source emitting waves of light that not only share the exact same frequency (color), but are also in-phase with each other. A beam of light consisting of just one frequency (color) is called *monochromatic*. A beam of light consisting of waves in perfect synchronization with each other is called *coherent*. Laser light is both monochromatic and coherent⁵¹. Modern solid-state lasers are constructed similarly to light-emitting diodes (LEDs), emitting light when energized with direct current electricity.

Certain types of solid-state lasers are capable of emitting light over a limited range of frequencies. That is to say, the light emitted by one of these laser types will always be of a consistent color at any given moment in time, but that color may be made to change over time. This modern class of solid-state lasers is therefore useful as light sources for absorption spectroscopy because the color of the emitted light may be adjusted to align with the absorption frequencies of certain chemical compounds.

As mentioned in the “Introduction to optical analyses” section, some molecules exhibit optical absorption over a band of frequencies (colors), and these absorptive bands are actually comprised of many fine *lines* clustered together on a spectrum diagram. Each line within an absorption band represents a particular resonant mode which a molecule dissipates incident light energy in the form of heat, and these resonant modes may be highly specific to the molecular species.

The following absorption spectrum illustrates this concept. Magnifying one of the bands found in this absorption spectrum reveals a closely-spaced series of individual absorption lines:



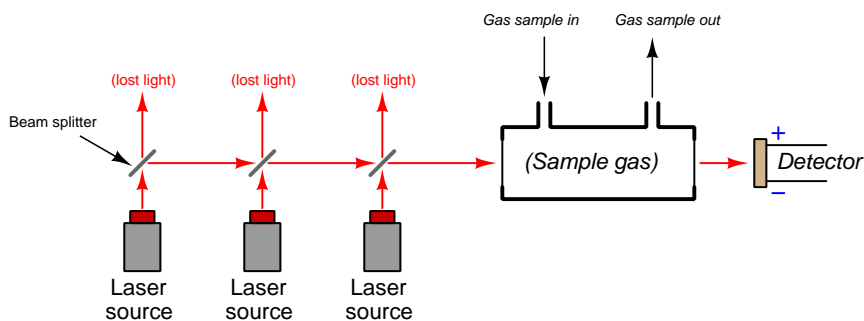
⁵⁰The term “laser” is actually an acronym, standing for **L**ight **A**mplification by **S**timulated **E**mission of **R**adiation.

⁵¹It is this coherence of laser light that enables the beam to remain highly focused, unlike light from other sources which tends to spread.

One of the challenges of absorptive spectroscopy is the problem of distinguishing one molecular species from another when their respective absorption bands happen to overlap. Laser spectroscopy avoids this problem by focusing on just one absorption line at a time, rather than attempting to measure light absorption over a broad spectrum. Even if multiple compounds have overlapping *bands*, a great many of the individual *lines* comprising those bands will be unique to each compound. A diode laser “tuned” to a single absorption line is therefore able to distinguish one compound from another even if other lines within the absorption band are shared by interfering species.

The frequency range of a laser diode is quite narrow, and so any spectroscopic analyzer employing this technology requires a laser light source custom-tuned for the species of interest. If multiple species having widely differing absorption spectra are to be measured, multiple laser sources will be required (each laser tuned to the very narrow range associated with each species’ absorption lines). One design of multi-laser gas analyzer is shown here:

Multi-source laser spectroscopy gas analyzer



In order for the one detector to discriminate between light received from each laser source, the sources must be sequentially energized. A computer controlling the laser sources analyzes the detector’s signal and correlates that against the known frequency emitted by each laser source at any given time in order to determine gas concentration. Like all gas analyzers, this design must be calibrated against known concentrations of gas (i.e. “zero” and “span” calibration gases) in order to establish “baseline” light intensities needed for accurate calculations of gas quantities.

A photograph of a Rosemount model CT5400 stack gas analyzer shows one possible arrangement of multiple laser diode sources, sample chamber, and optical detector. All optical components are mounted to a thick plate-aluminum “bench” maintaining them in precise physical alignment with each other. The laser diode modules are colored black, while the sample chamber is red and the detector is silver. Light is emitted out of the right-hand side of each laser diode module, turning 180 degrees around the right-hand end of the bench by means of mirrors, enters the red sample chamber on its right-hand side and exits on the left-hand-side, turns another 180 degrees around the other end of the bench by means of additional mirrors, and enters the silver detector assembly on its left-hand side. The sample chamber itself contains internal mirrors forcing the light to “zig-zag” through the sample space, thereby simulating a longer sample chamber to maximize light absorption by the gases of interest:



Three laser sources are shown in this photograph (with room for three more!), but one of them is a “demo” model cut open to reveal its beam-splitting mirror⁵². The two active laser sources are marked CO₂ and NO, but their frequency ranges are such that four different gas/vapor species may be measured in total – H₂O, CO₂, N₂O, and NO:



⁵²Such mirrors are *partially* silvered to let some light through while reflecting the rest of the light.

Two similar yet distinct diode laser technologies have been applied as light sources for absorption spectroscopy: *tunable diode lasers* (TDL) and *quantum cascade lasers* (QCL). The output frequency of a tunable diode laser is adjustable by varying the amount of electric current through the laser or by varying the temperature of the laser's diode junction. Typically the temperature is maintained at some constant value and the light frequency is modulated electrically, because this allows for much more rapid adjustments of light frequency. Quantum cascade lasers operate a bit differently: when powered, the internal heating of the laser structure causes its light frequency to naturally "sweep"⁵³ over a certain range without any external modulation. With either technology, the basic measurement principle is the same: shine the light from one of these lasers through a sample chamber containing a gas mixture, and measure the intensity of the laser light received at the far end of the chamber. That light intensity will attenuate every time the laser's frequency aligns with an absorption mode of gas molecules within the light path, and will return to full value when the frequency drifts away from resonance. That time-domain signal is received by a light sensor, then electrically sent to a microprocessor for analysis. The results of that analysis indicate the concentrations of light-absorbing gas molecules within the sample chamber.

⁵³A term often applied to this phenomenon of a QCL's frequency is *chirp*. A "chirp" refers to a burst of signal frequencies either increasing or decreasing along some range.

23.9 Fluorescence

When a high-energy photon strikes an atom, it may eject one of the lower-level electrons from its shell, leaving a vacancy to be filled by one of the electrons already residing in a shell higher than the vacancy but lower than that of the ejected electron. When that medium-level electron falls down to fill the vacancy, it emits a photon of less energy than the one responsible for ejecting the original electron. Thus, a high-energy photon strikes the atom, and in turn the atom releases a low-energy photon. This phenomenon is known as *fluorescence*.

The relationship between a photon's energy and its frequency (and correspondingly, its wavelength) is a well-defined proportionality of Planck's constant h :

$$E = hf \quad \text{or} \quad E = \frac{hc}{\lambda}$$

Where,

E = Energy carried by a single "photon" of light (joules)

h = Planck's constant (6.626×10^{-34} joule-seconds)

f = Frequency of light wave (Hz, or 1/seconds)

c = Speed of light in a vacuum ($\approx 3 \times 10^8$ meters per second)

λ = Wavelength of light (meters)

Therefore, the high-energy photon necessary for ejecting a low-level electron from an atom must be a photon of high frequency (short wavelength), and the low-energy photon emitted by the atom must be one of low frequency (long wavelength).

Photons carrying sufficient energy to eject low-level electrons from atoms typically exist in the ultraviolet range and above. The lower-energy photons emitted by the excited atoms often fall within the visible light spectrum. Thus, what we see here is a mechanism for ultraviolet (invisible) light to cause a substance to glow with visible colors.

Fluorescence is commonly used for entertainment purposes in the form of a *black light*: an electrical bulb designed to emit ultraviolet light. Many different organic compounds readily fluoresce under such a light source, producing an eerie glow. Chemical substances present in white paper, certain inks, and certain types of clothing detergents exhibit strong fluorescent properties, as do many bodily fluids⁵⁴. In fact, the presence of fluorescent compounds in paper, inks, and detergents is often intentional, to enhance appearance when viewed in natural sunlight containing ultraviolet light.

⁵⁴Blood, urine, semen, and various bodily proteins are known to fluoresce in the visible spectrum, making fluorescence a useful tool for crime-scene investigations. It's also useful when purchasing a new house, to check for pet droppings in the carpet. Such analysis is not for the faint of heart.

A variety of common food substances fluoresce. Quinine, an ingredient contained in “tonic water,” glows yellow-green when exposed to ultraviolet light:




Olive oil is another example of a food substance fluorescing easily under ultraviolet light. In this case, the color of the emitted light is amber in tone:



Molasses fluoresces a deep green color when exposed to ultraviolet light:



Chlorophyll is an example of a substance (occurring naturally in the tissues of green plants) capable of fluorescence when exposed to ultraviolet light. The color of its fluorescence is red, as shown in this photograph of a house-plant leaf illuminated by a black light:




Fluorescent dyes are often used as “invisible ink,” marking items in such a way as to be invisible under normal light, but plainly visible when exposed to concentrated ultraviolet light. Such ink is used to mark modern United States currency, such as this \$20 bill. The fluorescent stripe shown in this close-up photograph contains text reading “USA TWENTY”:



Not all substances fluoresce as easily as others. If a substance present within an industrial sample happens to fluoresce, and all other substances in the sample stream do not (or at least do not to any significant degree), we may apply fluorescence as an analytical technique for the selective measurement of that substance.

Sulfur dioxide (SO_2) is an atmospheric pollutant formed by the combustion of fuels containing sulfur. This gas also happens to exhibit fluorescence under ultraviolet light. A photograph of the fluorescence chamber taken from a Thermo Electron model 43 sulfur dioxide analyzer appears here:



A steady flow of sample gas enters and exits the chamber through the black plastic tubes. Ultraviolet light enters the chamber from a special lamp, and a highly sensitive light detector called a *photomultiplier tube* measures the amount of light emitted when SO_2 molecules inside the chamber fluoresce. The greater the concentration of SO_2 molecules in the gas mixture, the more light will be sensed by the photomultiplier tube for any given amount of ultraviolet light.


The incident ultraviolet light from the lamp cannot directly reach the photomultiplier tube, because there is no straight-line path from the lamp to the tube, and the interior walls of the chamber are non-reflective. The only⁵⁵ way for the tube to receive light is if molecules inside the chamber fluoresce when excited by the lamp's ultraviolet light. This ensures the instrument will truly measure fluorescence, and produce a “zero” output when no fluorescent molecules are present.

⁵⁵There is another way that light from the UV lamp could conceivably “take a corner” and reach the detector, and that is if the gas sample happens to contain dust or condensation droplets that would scatter the light. However, since gas samples are always dried and filtered prior to entering the sample chamber, this possibility is eliminated.


A close-up view of the ultraviolet emitter shows it to be a gas-discharge lamp. When an oscillating source of high-voltage electricity energizes the electrodes inside this lamp, an arc forms and emits pulsating rays of ultraviolet light:



The photomultiplier tube is a special vacuum tube operating on the principle of the *photoelectric effect*, whereby an incident photon (light particle) of sufficient energy ejects an electron upon striking a metal surface. Light entering a transparent glass window on the photomultiplier tube causes electrons to be emitted from an electrically-charged metal plate called the *photocathode*. Following the photocathode plate are a series of additional metal plates called *dynodes*, each one at a successively greater positive potential to provide kinetic energy to electrons attracted toward them. Each time electrons strike a dynode plate at high energy, even more electrons are emitted in a process called *secondary emission*. The result of secondary emission is that a multitude of electrons reach the final plate (called the *anode*) for each single photon striking the photocathode: the tube's action effectively *multiplies* the effect of each photon for maximum sensitivity. A relatively strong pulse of electric current measured at the anode signals the tube's reception of each photon:



A simplified photomultiplier tube and power supply circuit are shown here:



In a real instrument, the micro-ammeter would be replaced by an amplifier circuit, producing a strong electrical signal in direct response to received light intensity. In the case of a fluorescence analyzer, the amplifier's output signal becomes a representation of SO_2 molecule concentration inside the chamber.

Like any other type of analyzer technology, we must be aware of potential interfering substances as we use fluorescence to detect the concentration of the species of interest. Not only does sulfur dioxide fluoresce when exposed to ultraviolet light, but so does nitric oxide (NO) and many hydrocarbon components, especially those large hydrocarbon compounds classified as *polynuclear aromatic hydrocarbons* or *PAH*. Unfortunately, both nitric oxide and PAH compounds are produced in some industries where sulfur dioxide is an environmental concern. In order for a fluorescence-based SO_2 analyzer to *only* measure the concentration of sulfur dioxide in a gas stream possibly containing NO and/or PAH compounds, special care must be taken to eliminate the interference.

Fortunately for us, nitric oxide happens to fluoresce at a different wavelength than sulfur dioxide gas. This gives us the ability to de-sensitize the instrument to nitric oxide by placing an appropriate optical filter in front of the photomultiplier tube. This filter blocks light wavelengths emitted by the fluorescence of nitric oxide, so that the photomultiplier tube cannot "see" the fluorescence of NO gas molecules.

The light from hydrocarbon compound fluorescence happens to overlap the spectrum of fluorescent light from SO_2 and therefore the problem of hydrocarbon interference cannot be solved by optical filtering⁵⁶. Instead, this SO_2 analyzer handles the PAH interference problem by *physically* filtering out hydrocarbon gas molecules prior to the sample entering the fluorescence chamber using a device called a *kicker*. The "kicker" is a form of molecular sieve, separating the hydrocarbon molecules from the other molecules in the sample stream so that no hydrocarbon molecules ever enter the instrument's fluorescence chamber.


⁵⁶If one were to install an optical filter in front of the photomultiplier tube designed to block fluorescent light emitted by hydrocarbon molecules, this filter would *also* block the light emitted by fluorescing SO_2 molecules thereby defeating the very purpose of the analyzer: measuring SO_2 concentration by optical fluorescence!

After processing by the electronic circuits of the analyzer, the photomultiplier tube's output signal becomes a representation of SO_2 concentration, displayed on an analog meter movement:



As indicated by the selector switch below the meter face, this instrument has three different display ranges: 0 to 0.5 ppm (*parts per million*), 0 to 1.0 ppm, and 0 to 5.0 ppm. A different selector switch on the left-hand side of the control panel operates solenoid valves allowing either the process sample gas or one of two different calibration gases to enter the analyzer. The “zero” calibration gas contains no sulfur dioxide at all, thus providing a base-line reference for adjusting the 0% point of the analyzer. The “span” calibration gas contains a precise mixture of sulfur dioxide and some non-fluorescing carrier gas, to serve as a chemical reference for some point near the analyzer's upper range limit. These calibration gases are commercially available from chemical laboratories, with instrument technicians commonly referring to them as *zero gas* and *span gas*. Of course, the composition of any “zero” or “span” gas depends entirely on the type of analytical instrument. What may suffice as a span gas for this sulfur dioxide analyzer would certainly not suffice as a span gas for a multi-component chromatograph or for an NDIR analyzer configured to measure carbon monoxide.

Pressure regulators ensure proper gas flow conditioning in and out of the analyzer. A vacuum pump (not shown in any of the photographs) draws sample gas through the analyzer and provides the necessary differential pressure for the hydrocarbon “kicker” to work:




The amount of sample gas pressure is a critically important parameter for this and other optical gas analyzers, because pressure has a direct effect on the density of the gas sample, and therefore on the number of gas molecules per unit volume. If a gas sample of constant species concentration (i.e. a constant parts-per-million proportion of the gas of interest compared to the balance of the gas) happens to increase in pressure, there will now be more molecules of light-absorbing (or light-emitting) gas in that sample, which will register as a higher concentration even though the actual concentration (in percent or ppm or ppb) has not changed. This is why pressure regulation is so important for gas analyzers: an unstable sample gas pressure will result in measurement errors.

23.10 Chemiluminescence

Recall that an *exothermic* chemical reaction is one that releases a net sum of energy, as opposed to an *endothermic* reaction which requires a greater input of energy than it releases. Combustion is a common class of exothermic reactions, with the released energy being very obviously in the forms of heat and light, with heat being the predominant form.

Some exothermic reactions release energy primarily in the form of light rather than heat. The general term for this effect is *chemiluminescence*. A natural example is the “cold” light emitted by North American species of firefly. In this small insect, a chemical reaction intermittently takes place emitting significant amounts of light but insignificant amounts of heat. An artificial example is the light emitted by a “glow stick” when activated. The following photographs show such a light source before activation when the reactants are separated (left) and after activation when the internal barrier is broken and the reactants are allowed to mix (right):



Certain industrial compounds engage in chemiluminescent reactions, and this phenomenon may be used to measure the concentration of those compounds. One such compound is nitric oxide (NO), an atmospheric pollutant formed by high-temperature combustion with air as the oxidizer⁵⁷.


A spontaneous chemical reaction between nitric oxide and *ozone* (an unstable molecule formed of three oxygen atoms: O₃) is known to produce chemiluminescence:



Although this process of generating light is quite inefficient (only a small fraction of the NO₂ molecules formed by this reaction will emit light), it is predictable enough to be used as a quantitative measurement method for nitric oxide gas. Ozone gas is very easy to produce on demand, by exposing air or oxygen to a high-voltage electric discharge.

⁵⁷Combustion is primarily a reaction between carbon and/or hydrogen atoms in fuel, and oxygen atoms in air. However, about 78% of the air (by volume) is nitrogen, and only about 20.9% is oxygen, which means a lot of nitrogen gets pulled in with the oxygen during combustion. Some of these nitrogen atoms combine with oxygen atoms under the high temperature of combustion to form various oxides of nitrogen.

A simplified diagram for a chemiluminescent nitric oxide gas analyzer appears here:




As with many optical analyzers, a photomultiplier tube serves as the light-detecting sensor, generating an electrical signal in proportion to the amount of light observed inside the reaction chamber. The higher the concentration of NO molecules in the sample gas stream, the more light will be emitted inside the reaction chamber, resulting in a stronger electrical signal produced by the photomultiplier tube.

Although this instrument readily measures the concentration of nitric oxide (NO), it is insensitive to other oxides of nitrogen (NO_2 , NO_3 , etc., collectively referred to as NO_x , pronounced “nocks”). Normally, we would consider this selectivity to be a good thing, because it would eliminate interference problems from these other gases. However, as it so happens, these other oxides of nitrogen are every bit as polluting as nitric oxide, and therefore when we measure nitric oxide for pollution monitoring purposes, we usually *also* wish to measure these other oxides⁵⁸ in combination.

⁵⁸The measures used to mitigate nitric oxide emissions are the same measures used to mitigate the other oxides of nitrogen: reduce combustion temperature, and/or reduce the NO_x compounds to elemental nitrogen by mixing the combustion exhaust gases with ammonia (NH_3) in the presence of a catalyst. So here we have a case where we really don't care to distinguish NO from NO_x : we want to measure it *all*.

In order to use chemiluminescence to measure *all* oxides of nitrogen, we must chemically convert the other oxides into nitric oxide (NO) before the sample enters the reaction chamber. This is done in a special module of the analyzer called a *converter*:




A three-way solenoid valve is shown in this diagram, providing a means to bypass the converter so the analyzer only measures nitric oxide content in the sample gas. With the solenoid valve passing all the sample through the converter, the analyzer responds to *all* oxides of nitrogen (NO_x) and not just nitric oxide (NO).

One simple way to achieve the $\text{NO}_x \rightarrow \text{NO}$ chemical conversion is to simply heat the sample gas to a high temperature, around 1300 °F. At this temperature, the molecular structure of NO is favored over more complex oxides such as NO_2 , the result being a release of oxygen from the NO_2 and NO_3 molecules to become NO molecules. A disadvantage of this technique is that those same high temperatures also have a tendency to convert *other* compounds of nitrogen such as ammonia (NH_3) into nitric oxide, thereby creating an unintended interference species⁵⁹.

⁵⁹This particular interference compound is especially problematic if we are using the analyzer to *control* the NO_x concentration in the exhaust of a combustion process, and the manipulated variable for the NO_x control loop is pure ammonia injected into the exhaust. Un-reacted ammonia (commonly called *ammonia slip* in the industry) sampled by the analyzer will be falsely interpreted as NO_x , rendering the measurement meaningless, and therefore making control virtually impossible.

An alternative $\text{NO}_x \rightarrow \text{NO}$ conversion technique is to use a metallic reactant in the converter to remove the extra oxygen atoms from the NO_2 molecules. One such metal that works well for this purpose is molybdenum (Mo) heated to the comparatively low temperature of 750 °F, which is too low to convert ammonia into nitric oxide. The reaction of NO_2 converting to NO is as follows:



Other oxides (such as NO_3) convert in similar fashion, leaving their excess oxygen atoms bound to molybdenum atoms and becoming nitric oxide (NO). The only difference between these reactions and the one shown for NO_2 is the proportional (stoichiometric) ratios between molecules.

As you can see from the reaction, the molybdenum metal is converted into the compound molybdenum trioxide over time, requiring periodic replacement. The rate at which the molybdenum metal depletes inside the converter depends on the sample flow rate and the concentration of NO_2 .


As with other optical gas analyzers, pressure control of the gas sample is critically important for good measurement accuracy. If the pressure of the sampled gas inside the chemiluminescence reaction chamber happens to vary, it will affect the amount of light emitted even if the relative concentration of NO_x gas remains stable. This is because higher pressures will pack gas molecules closer together, resulting in more reactive molecules inside the chamber for any given percentage or ppm concentration. For this reason you will see analyzers such as this equipped with pressure regulation to ensure the gas pressure inside the measurement chamber remains constant.

23.11 Analyzer sample systems

Some analyzers measure the composition of a process stream by directly immersing the sensing element in that stream. This is called *in situ* measurement, which is a Latin phrase meaning “in the place.” A pH probe inserted into a process pipe, an oxygen probe inserted into the stack of a combustion furnace, and a GFC analyzer measuring the concentration of a gaseous pollutant by shooting a light beam across a process room are all examples of *in situ* analyzers: the analytical sensing takes place directly within the process environment.

Alternatively, analytical sensing elements may be located some distance away from the process, in which case a representative sample of that process stream must be conveyed to the analyzer for measurement. A great many industrial analyzers function like this, with a system of tubes, heaters, filters, pumps, regulators, and other components working together to provide the remotely located analyzer with a steady stream of process fluid to sample. It should be noted that quite often problems experienced with process analyzers stem from improperly constructed and/or maintained sample systems.

A very simple sample system for a pH analyzer is shown here, the purpose for it being to allow the pH electrode to operate at atmospheric pressure instead of the high pressure inside the process pipe⁶⁰:



Liquid from the process pipe flows through the isolation and flow-control valves, the rate of flow restricted to a mere trickle, where it continuously fills a sample chamber and overflows into a process sewer designed to transport a constant stream of process liquid. If for some reason a process sewer drain is unsafe, environmentally destructive, or impractical, one may replace the overflow tube with

⁶⁰In-situ pH probes are manufactured for high-pressure applications, but they suffer short lifespans (due to the accelerated erosion of the measurement glass) and decreased sensitivity (due to the extra thickness of the measurement glass) and are substantially more expensive than pH probes designed for atmospheric pressure conditions.

a level control system and pump to re-inject the sample back into the process line. This, of course, adds complexity and expense to the sample system.

Despite the apparent simplicity and necessity of such a sample system, its very existence adds one more layer of fault potential to the measurement system. Imagine, for example, the sample line between the isolation and flow control hand valves becoming plugged with debris from within the process pipe. This would prevent fresh sample from reaching the sample chamber. The pH probe would then measure nothing but old, stale sample liquid rather than a representation of the liquid flowing through the process pipe. The pH analyzer has no way of “knowing” that what it is measuring does not reflect the state of the process, and will report meaningless pH data.

Liquid sample systems may require much more complexity than what is shown here for this simple pH analyzer. To meet the sample needs of some analyzers, we may need to control the flow rate of the sample, the temperature of the sample, and/or filter it for particulates.

Sample systems for gas analyzers tend to be even more complex than sample systems for liquids. Not only is sample pressure control⁶¹, temperature control⁶², and filtering⁶³ standard for most gas analyzers, the trapping and removal of *condensation* from gas sample lines is also important, as is the prevention of gas permeation through sample lines⁶⁴. If liquid and/or particulate matter happen to enter a gas analyzer, the results may range from measurement delay to gross measurement error to damage of critical components within the analyzer, depending on the analyzer type and the concentration of inappropriate matter in the sample. Analytical gas samples, as a general rule, must be impeccably clean and dry.

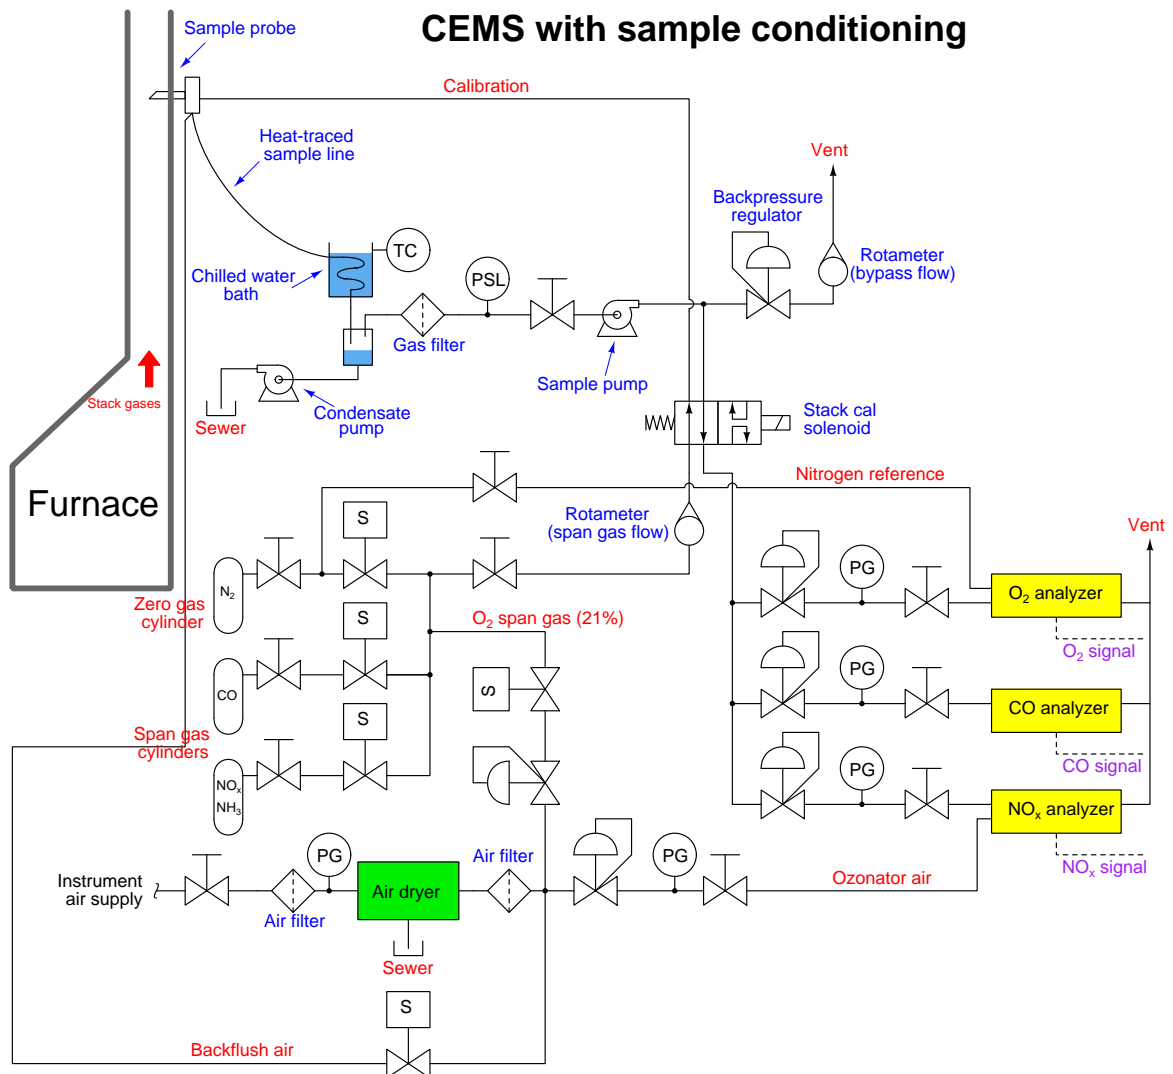
⁶¹Pressure control is important in gas analysis because changes in sample gas pressure will result in different gas densities, thereby directly affecting how many molecules of the gas of interest will be present and therefore detectable inside the analyzer.

⁶²Temperature control is important for similar reasons: the gas species of interest may become more reactive as temperature changes, thereby resulting in a stronger indication even when concentration remains constant.

⁶³It is important to thoroughly filter the gas input to an analyzer so that contaminants do not foul the sensing element(s). This is rather obvious in the case of optical analyzers, where the light to be analyzed must pass through a transparent window of some kind, and that window must be kept clean of dust, condensation, and any other substances that could interfere with the transmission of light.

⁶⁴Some types of plastic sample tubes are permeable to gases, and so represent potential contamination points when the concentrations of interest are in the range of parts per million (ppm) or parts per billion (ppb). In such critical applications, only metal sample tubes (stainless steel, typically) are appropriate.

An example of a typical sample system for a set of gas analyzers used to continuously monitor the emissions from the stack of a combustion furnace (*CEMS*: Continuous Emissions Monitoring System) is shown here:



There is much to comment on in this diagram, just focusing on the sample probe and line. Note how the sample probe is beveled so as to avoid picking up any particulate (solid matter) traveling with the furnace exhaust gas. Note also how the sample tube is heat-traced to prevent condensation from gathering in the tube, and how the entire length of the sample line is sloped downhill to prevent any condensed liquid from forming a "trap" that gases would have to bubble through to reach the analyzer system. Once arriving at the analyzer shelter, the sample line immediately plunges into a

chilled water bath to force any condensible vapors to condense and be drained away, leaving only a dry gas sample to enter the rest of the analyzer tubing system.

Another feature of this CEMS is an *air dryer* (shown as a green rectangle), used to remove moisture from compressed air before that air is used within the system. In this particular CEMS, dry air is used both as a span gas for the oxygen analyzer (ambient oxygen concentration being approximately 20.9% by volumetric or molar concentration) as well as feed for the NO_x analyzer's internal ozonator. A photograph of such an air dryer is shown here, part of the CEMS at a gas-fired power plant:



Accurate and reliable operation of any analyzer depends not just on the analyzer itself, but also on all other equipment supporting its role. With analyzer systems such as CEMS, the amount and complexity of supporting equipment is substantial, which means analyzer systems tend to require much more maintenance than other forms of instrumentation. Fault diagnosis is also more complicated for analyzer systems for the same reason.

Turning our attention now to the rest of the sample system, we see it provides the means for *automatic calibration* of each analyzer, where a computer (most likely a PLC, or a specialized computer sequencing the operation of all three analyzers) periodically shuts off the sample flow to the analyzers and replaces that flow with standard “span gas” to self-check the response of the analyzers. Note how the mixed calibration gas has the option of being injected at the sample probe so as to force the gas mixture to travel down the sample line and through the whole system (with the “stack cal” valve in its default position). While this may seem needless and wasteful, it indeed serves a practical purpose: to test the integrity of the entire *system*, sample conditioning and all. If, for example, there were some source of contamination entering the sample stream by way of the condensate trap, this would be detected during the self-calibration cycle because the calibration gas mixture would have to pass through the same contaminated area as the normal stack gas sample. Another example is a leak in one of the sample lines, allowing ambient gases to leak in and dilute or otherwise contaminate the sample. A self-calibration system bypassing the whole sample line and conditioning system would not detect the presence of any contamination, and therefore would not be testing the integrity of the whole system.

A photograph of three span gas cylinders used in a CEMS self-calibration system appears here:



The three span gas mixtures include one with 30 ppm concentration carbon monoxide gas (CO) and 15% concentration oxygen gas (O₂), another with nitric oxide gas (NO) at a concentration of 25 ppm, and a third with nitrogen dioxide gas (NO₂) at a concentration of 25.6 ppm. The “balance⁶⁵” gas for each of these calibrated mixtures is nitrogen. Each of these span gas mixtures is manufactured at a facility where the concentration quantities are traceable to NIST standards, just like any other calibration-grade measurement standard in instrumentation.

Analytical instruments typically require far more frequent calibration checks than standard instrument types (pressure, level, temperature, flow), and much maintenance labor may be saved by automatic calibration systems. Such routine checks of analyzer calibrations are commonly referred to as *drift tests*, because their purpose is to verify how far the analyzer’s calibration has “drifted” from its specification. Unfortunately, the very same systems intended to reduce routine maintenance effort also introduce more points for failure, and therefore contribute to the non-routine maintenance needs of the system.

⁶⁵The “other” gas in the mixture besides the gas or gases of interest.

23.12 Safety gas analyzers

Process analyzers measure the concentration of specific substances for the purpose of measuring and/or controlling those concentrations in a process stream. Safety analyzers detect the presence of dangerous concentrations of specific substances to warn personnel of threats to life or health. While there is virtually no end to the different types of process analyzers in existence, with new analyzer types invented to meet the needs of process industries, safety analyzers are rather restricted to those substances known to pose health hazards to human beings.

Safety analyzers are designed for fast response, rugged construction, and ease of portability. As such, they are usually not as accurate or as sensitive to slight changes in concentration as process analyzers. The sensing technologies used in safety analyzers are often very different from those used in process analyzers. You will never, for example, see an NDIR instrument with a Luft-style “microphone” detector used for portable safety applications⁶⁶. The high accuracy of a Luft-style NDIR instrument is not necessary for safety, and the bulk (and fragility) of such an instrument makes it completely impractical as a portable device.

An important industrial application of safety gas analyzer technology is testing the safety of the atmosphere inside a *confined space*, such as an industrial vessel that has been emptied in preparation for workers to enter and perform maintenance work. Virtually all pressure vessels are classified as confined spaces from the perspective of internal maintenance work, and many are explicitly labeled as such to remind workers of the danger:



⁶⁶Interestingly, there is a documented case of an NDIR “Luft” analyzer being used as a safety monitor for carbon monoxide, ranged 0 to 0.1% (0-1000 ppm), at one of I.G. Farbenindustrie’s chemical plants in Germany during the 1940’s. This was definitely not a *portable* analyzer, but rather stationary-mounted in a process unit where high concentrations of carbon monoxide gas existed in the pipes and reaction vessels. The relatively fast response and high selectivity of the NDIR technology made it an ideal match for the application, considering the other (more primitive) methods of carbon monoxide gas detection which could be “fooled” by hydrogen, methane, and other gases.

This photograph shows a hand-held safety gas monitor, used to detect four different gases with known hazard levels (oxygen, carbon monoxide, combustibles, and hydrogen sulfide). A technician working in hazardous environments would wear one of these at all times, listening for the audible warning tone generated by the device if any of its pre-set limits is exceeded:



Most portable gas analyzers such as this employ electrochemical sensing “cells” generating a small electric current when exposed to the gas of interest. Such technologies may not always be the most accurate or the most sensitive, but their characteristics are well suited for portable applications where they will be exposed to vibration and must operate on battery power.

A close-up photograph taken of the monitor's alarm thresholds reveals the relative concentrations of four gases monitored by the device. Two distinct alarm levels – “low” (alert) and “high” (danger) – exist to warn the user of threats:



The two most dangerous gases detected by this device – hydrogen sulfide (H₂S) and carbon monoxide (CO) – are measured in units of *parts per million*, or *ppm*. The next two detected gases – oxygen (O₂) and combustibles (“LEL”) – are measured in units of percent (which may be thought of as *parts per hundred*).

Like all safety-related devices, portable gas analyzer require regular “proof-testing” and calibration. In order to accurately check the response of a gas analyzer, it must be exposed to gases of known concentrations. Special mixtures of “test gas” for safety analyzers are available from chemical supply companies, the next photograph showing the certified concentrations of different gases contained inside the pressurized test cylinder:




Note the limited shelf life of this span gas: only one year between its date of manufacture (November 2007) and its expiration (November 2008). Relatively short expiration periods are common with chemical standards used for analytical calibration, owing to slow chemical reactions between the constituent compounds over time altering the composition of the mixture. In this test gas cylinder, for example, the relatively high concentration of oxygen gas (18%) will over time oxidize the hydrogen sulfide, methane, and carbon monoxide gases, resulting in decreased concentrations of all these gases, with the production of water vapor and other oxides (e.g. CO₂) inside the gas cylinder.

23.12.1 Oxygen gas

Most living things require oxygen to survive. The oxygen you breathe combines with nutrients from the food you eat to produce energy in a form usable by your body. If you are deprived of oxygen, your body very quickly shuts down, much like a fire dies when starved of oxygen (and for approximately the same reason). Ambient air is approximately 20.9% oxygen by volume, the majority of air (about 78% by volume) being nitrogen.

The oxygen content of air may be reduced by combustion (which combines oxygen with flammable substances to produce carbon dioxide and water vapor) or by displacement by a denser gas (such as propane) in a low-lying area or by any gas in sufficient quantity filling an enclosed area.

A modern oxygen sensor technology for safety applications is the *micro fuel cell*, generating a measurable electric current in the presence of oxygen by the oxidation of a self-contained fuel source. In many sensors, the fuel is pure lead (Pb), with the resulting chemical reaction producing lead oxide (PbO):



Fuel cell sensors are relatively rugged, accurate, and self-powering, enabling their use in portable oxygen analyzers. Due to their principle of operation, where an internal fuel is slowly oxidized over time, these sensors have a rather limited life and therefore must be periodically replaced.

An interesting and useful technique for testing the operation of an oxygen safety sensor is to exhale on the sensor, watching for a decrease in oxygen content to 15% or below. This testing technique makes use of the fact that your body extracts oxygen from the air, such that your exhaled breath contains less oxygen than it did when inhaled. Therefore, your own body acts as a crude “calibration gas” source for an oxygen safety analyzer.

23.12.2 Lower explosive limit (LEL)

The minimum concentration of a flammable gas in air capable of igniting is called the *Lower Explosive Limit*, or *LEL*. This limit varies with the type of gas and with the oxygen concentration of the air in which the flammable gas is mixed. Sensors designed to detect the dangerous presence of combustible gases are therefore called “LEL sensors.”

LEL monitors are used whenever there is a high probability of explosive gases present in the air. These areas are referred to as *classified* areas in industry, and are precisely defined for safety engineering purposes. Classified areas harboring explosive gases or vapors are deemed *Class I* areas, with different “Group” categories delineating the specific gas or vapor types involved. For more information on classified areas, refer to section [32.1.1](#) beginning on page [2599](#).

Gases and vapors are not the only substances with the potential to explode in sufficient concentration. Certain dusts (such as grain) and fibers (such as cotton) may also present explosion hazards if present in sufficient quantity. Unfortunately, the majority of analytical technologies used to monitor lower explosive limits for safety purposes only function with gases and vapors (Class I), not dusts or fibers (Class II and Class III, respectively).

Popular sensor technologies used to detect the presence of combustibles in air include the following:

- Catalytic bead
- Thermocouple
- Infrared
- Flame ionization

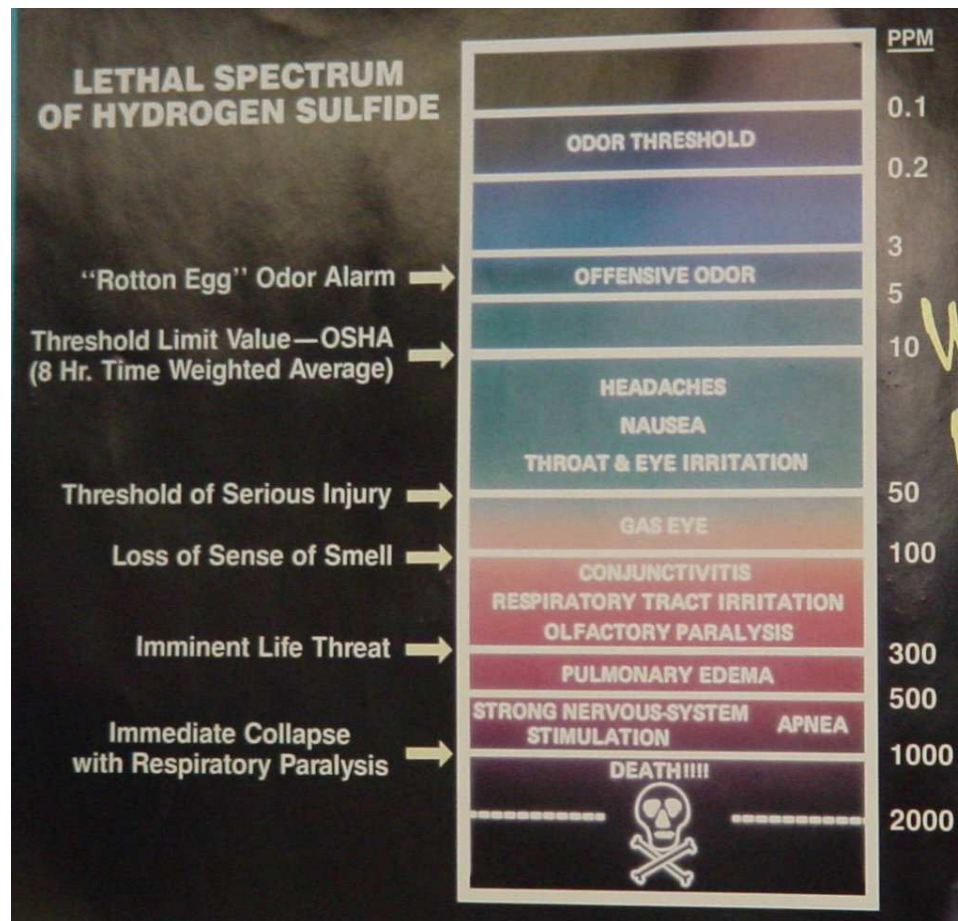
Catalytic bead and thermocouple sensors both function on the principle of heat generated during combustion. Air potentially containing a concentration of flammable gases or vapors passes near a heated element, and any combustion occurring at that point will cause the local temperature to immediately rise. These sensors must be designed in such a way they will *not* initiate an explosion, but merely combust the sample in a safe and measurable manner. Like micro fuel cell oxygen sensors, these sensors may be manufactured in sufficiently small and rugged packages to enable their use as portable LEL sensors.

Infrared analyzers exploit the phenomenon of infrared (IR) light absorption by certain types of flammable gases and vapors. A beam of infrared light passed across a sample of air will diminish in intensity if significant concentrations of the combustible substance exist in that sample. Measuring this attenuation provides an indirect measurement of explosive potential. A major disadvantage of this technique is that many non-flammable gases and vapors also absorb IR light, including carbon dioxide and water vapor. In order to successfully reject these non-flammable substances, the analyzer must use very specific wavelengths of IR light, tuned to the specific substances of interest (and/or wavelengths tuned specifically to the substances of non-interest, as a compensating reference signal for the wavelengths captured by both the substances of interest and the substances of non-interest).

Flame ionization sensors work on the same principle as FIDs for chromatographs: a non-ionizing flame (usually fueled by hydrogen gas) will generate detectable ions only in the presence of air samples containing an ionizing fuel (such as a hydrocarbon gas). Of course, this form of LEL sensor is useless to detect hydrogen gas.

23.12.3 Hydrogen sulfide gas

Hydrogen Sulfide (H_2S) is a highly toxic gas, with a pungent “rotten eggs” odor at low concentrations but no visible color. At higher concentrations, the gas acts as a nerve agent to de-sensitize human smell, so that it seems odorless. Its paralytic effect on smell extends to more important bodily functions such as breathing, causing rapid loss of consciousness and asphyxiation. A photograph of a safety chart (taken at a wastewater treatment facility) shows just how toxic hydrogen sulfide gas is:



Note how concentrations in the *parts per million* range are hazardous, and how little H_2S concentration is required to paralyze one's sense of smell. Hydrogen sulfide also happens to be flammable, its LEL value in air being 4.3%. However, the toxic properties of the gas are generally the more pressing concern when released into the atmosphere. Another hazardous property of hydrogen sulfide is its density: 1.18 times that of air. This means H_2S gas will tend to collect in low areas such as pits, electrical vaults, and empty underground storage vessels.

The principal source of hydrogen sulfide gas is anaerobic (oxygen-less) decomposition of organic matter. Sewage treatment facilities, pulp mills, and oil refineries generate H_2S gas in significant

quantity, and so employees at such facilities must be continually aware of the associated hazards.

One of the most popular analytical sensing technologies for H₂S gas appropriate for portable monitoring includes an electro-chemical reaction cell similar in principle to the micro fuel cells used to detect oxygen concentrations. Hydrogen sulfide gas entering such a cell engages in a specific chemical reaction, creating a small electrical current proportional to the gas concentration. Like oxygen-sensing fuel cells, these chemical cells also have limited lives and must be routinely replaced.

23.12.4 Carbon monoxide gas


Carbon monoxide (CO) gas is a colorless, odorless, and toxic gas principally generated by the incomplete combustion of carbon-based fuels. The mechanism of its toxicity to people and animals is the preferential binding of CO gas molecules to the hemoglobin in blood. At significant concentrations of carbon monoxide gas in air, the hemoglobin in your blood latches on to CO molecules instead of oxygen (O₂) molecules, and remains bound to the hemoglobin, preventing it from transporting oxygen. The result is that your blood rapidly loses its oxygen-carrying capacity, and your body asphyxiates from within. Like hydrogen sulfide, carbon monoxide is also flammable (LEL = 4%), but its toxic properties are generally the larger concern when released into the atmosphere.

Carbon monoxide is not to be confused with carbon *dioxide* (CO₂) gas, which is almost completely inert to the human body. Carbon dioxide is principally produced by *complete* combustion of carbon-based fuels. Its only safety hazard potential is the capacity to displace breathable air in an enclosed area if rapidly released in large volumes.

Combustion burners operating on carbon-based fuels may produce excess carbon monoxide if operating at too rich an air/fuel mixture. Even when adjusted optimally, there will always be some carbon monoxide present in the exhaust. This makes high CO concentrations possible where burners operate in enclosed areas.

Some industrial processes such as catalytic cracking in the petroleum refining industry generate huge amounts of carbon monoxide, but these extremely high concentrations are normally present only within the process piping and vessels, not released to atmosphere. Nevertheless, personnel working near such processes must wear portable CO gas safety monitors at all times to warn of leaks.


Carbon monoxide may be sensed by an electrochemical cell, using iodine pentoxide as the reacting compound. The balanced chemical reaction is as follows:



The strength of the electric current produced by the cell indicates the concentration of carbon monoxide gas.

23.12.5 Chlorine gas

Chlorine (Cl_2) gas is a strong-odored, toxic gas used as a biological disinfectant, bleaching agent, and as an oxidizer in many industrial processes. Colorless in low concentrations, it may appear green in color when mixed in very high concentrations with air. Chlorine is highly reactive, presenting a distinct hazard to mucus membranes (eyes, nose, throat, lungs) by creating hypochlorous acid (HOCl) and hydrochloric acid (HCl) upon contact with water:



The following table correlates levels of chlorine gas concentration in ambient air with degree of hazard. Note the unit of measurement for chlorine concentration in air – *parts per million* (ppm). Bear in mind that one part per million is equivalent to just 0.0001 percent:

Concentration in air	Hazard
1 ppm to 3 ppm	Mild mucus membrane irritation
5 ppm to 15 ppm	Upper respiratory tract irritation
30 ppm	Immediate chest pain, cough, and difficulty breathing
40 ppm to 60 ppm	Toxic pneumonitis and pulmonary edema
430 ppm	Lethal over 30 minutes
1000 ppm (0.1%)	Lethal within a few minutes

Water and wastewater treatment operations frequently⁶⁷ use chlorine for disinfection of water. Pulp mills use either chlorine or chlorine compounds as a bleaching agent to whiten wood pulp.

⁶⁷Some water treatment facilities use powerful ultraviolet lamps to disinfect water without the use of chemicals. Some potable (drinking) water treatment plants use ozone gas (O_3) as a disinfectant, which is generated on-site from atmospheric oxygen. A disadvantage to both chlorine-free approaches for drinking water is that neither one provides lasting disinfection throughout the distribution and storage system to the same degree that chlorine does.

Chlorine may be generated on site by the electrolytic decomposition of salt (sodium chloride – NaCl), or delivered in cylindrical pressure vessels in liquid form as shown here at a large wastewater treatment facility:



23.13 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Conservation of energy:** energy cannot be created or destroyed, only converted between different forms. Relevant to optical analyses, particularly fluorescence where the energy of a photon emitted by fluorescence is always less (never more) than the incident photon, the difference between the two photons' energies being dissipated in the form of heat.
- **Conservation of mass:** mass is an intrinsic property of matter, and as such cannot be created or destroyed. Relevant to chemical reactions, where the total mass of the reaction's products must precisely equal the total mass of the reactants.
- **Kirchhoff's Voltage Law:** the algebraic sum of all voltages in a loop is equal to zero. Relevant to conductivity and pH sensor circuits, as well as bridge circuits used in a variety of analytical instruments.
- **Ideal Gas Law:** $PV = nRT$, describing the relationship between gas pressure, chamber volume, gas quantity (in moles), and gas temperature. Relevant to expansion of gas inside Luft detectors (NDIR instruments).
- **Common logarithms:** used to express measurements spanning a tremendous range. Relevant to pH calculations, where the pH of an aqueous fluid is the negative logarithm of its hydrogen ion activity (molarity) $\text{pH} = -\log[H^+]$. Also used to calculate absorbance of light by different substances (the Beer-Lambert Law $A = \log \frac{I_0}{I}$).
- **Self-balancing opamp circuits:** all self-balancing operational amplifier circuits work on the principle of negative feedback maintaining a nearly zero differential input voltage to the opamp. Making the "simplifying assumption" that the opamp's differential input voltage is exactly zero assists in circuit analysis, as does the assumption that the input terminals draw negligible current.
- **Nernst equation:** $V = \frac{RT}{nF} \ln \left(\frac{a_1}{a_2} \right)$ predicting the amount of voltage developed across an ion-permeable membrane by the exchange of ions across that membrane. Relevant to all forms of potentiometric chemical analysis, where sensor voltage is proportional to the logarithm of concentration quotient across the sensor membrane.
- **Time constant:** (τ), defined as the amount of time it takes a system to change 63.2% of the way from where it began to where it will eventually stabilize. The system will be within 1% of its final value after 5 time constants' worth of time has passed (5τ). Relevant to pH measurement where signal voltage changes are damped by cable capacitance (forming an RC time constant).

- **Integration (calculus):** where one variable is proportional to the accumulation of the product of two others. Integration always results in a multiplication of units. Relevant to calculations of mass passing through a chromatograph detector: total mass of a species passed through the detector (m) equal to the integral of mass flow rate times time: $m = \int W dt$.
- **Quantization of photon energy:** $E = hf$, where the energy carried by each photon (particle of light) is a fixed quantity proportional to the frequency (color) of that photon. Relevant to spectroscopy, where different colors of light interact with substances depending on the energy levels associated with electrons around the atoms of the substance.

References

“Automated Measuring System Technologies”, Best Practice brochure, Cleaner Fossil Fuels Programme, document BPB008, DTI, 2004.

Annino, Raymond; Villalobos, Richard, *Process Gas Chromatography – Fundamentals and Applications*, Instrument Society of America, Research Triangle Park, NC, 1992.

Boylestad, Robert L., *Introductory Circuit Analysis*, 9th Edition, Prentice Hall, Upper Saddle River, NJ, 2000.

Carroll, Grady C., *Industrial Process Measuring Instruments*, McGraw-Hill Book Company, Inc., New York, NY, 1962.

Chu, P.M.; Guenther, F.R.; Rhoderick, G.C.; Lafferty, W.J.; “The NIST Quantitative Infrared Database”, *Journal of Research of the National Institute of Standards and Technology*, Volume 104, Number 1, Gaithersburg, MD, January-February 1999.

Fribance, Austin E., *Industrial Instrumentation Fundamentals*, McGraw-Hill Book Company, New York, NY, 1962.

Gregory, C.H.; Appleton, H.B.; Lowes, A.P.; Whalen, F.C.; *Instrumentation and Control in the German Chemical Industry*, Mapleton House, Brooklyn, NY, 1947.

“Investigation Report – Chlorine Release”, Report number 2002-04-I-MO, U.S. Chemical Safety and Hazard Investigation Board, Washington DC, 2003.

Jernigan, J. Ron, *Chemiluminescence NO_x and GFC NDIR CO Analyzers For Low Level Source Monitoring*, Thermo Environmental Instruments, Franklin, MA.

Kanai, Hideo; Nakamura, Yusuke; Koizumi, Kazuhiro; “ZSS Cross Stack Laser Gas Analyzer”, *Fuji Electric Review*, Volume 54, Number 3.

Kohlmann, Frederick J., *What Is pH, And How Is It Measured?*, Hach Company, 2003.

Kuiken, Tim; Crabbe, Keith; “Tunable Diode Laser Spectroscopy Detection Principles For Monitoring NH₃, HCl & HF”, CEMTEK Environmental, Inc., 2011.

Kume, Hidehiro, *Photomultiplier Tube – Principle to Application*, Hamamatsu Photonics K.K., 1994.

Lavigne, John R., *Instrumentation Applications for the Pulp and Paper Industry*, Miller Freeman Publications, Foxboro, MA, 1979.

Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.

Luft, Von K.F., “Über eine neue Methode der registrierenden Gasanalyse mit Hilfe der Absorption ultraroter Strahlen ohne spektrale Zerlegung”, *Zeitschrift für technische Physik*, Vierundzwanzigster

Jahrgang, Number 5, pp. 97-105, 1943.

“Multiple Component Real Time Impurity & Process Composition Analysis Using Mid IR Process Analytical Tunable Laser Spectroscopy (PATLS)”, document QCL-TWP-Mid-IR-PATLS, Emerson Process Management, October, 2015.

Novak, Joe, *What Is Conductivity, And How Is It Measured?*, Hach Company, 2003.

Pauling, Linus, *General Chemistry*, Dover Publications, Inc., Mineola, NY, 1988.

“pH Electrode Cleaning and Maintenance Guide”, document LIT2781 D12, Hach Company, 2012.

Scott, Raymond P.W., *Gas Chromatography*, Library4Science, LLC, 2003.

Scott, Raymond P.W., *Gas Chromatography Detectors*, Library4Science, LLC, 2003.

Scott, Raymond P.W., *Liquid Chromatography*, Library4Science, LLC, 2003.

Scott, Raymond P.W., *Liquid Chromatography Detectors*, Library4Science, LLC, 2003.

Scott, Raymond P.W., *Principles and Practice of Chromatography*, Library4Science, LLC, 2003.

Sherman, R.E.; Rhodes, L.J., *Analytical Instrumentation: practical guides for measurement and control*, ISA, Research Triangle Park, NC, 1996.

Shinskey, Francis G., *pH and pION Control in Process and Waste Streams*, John Wiley & Sons, New York, NY, 1973.

Skoog, Douglas A.; Holler, F. James; Nieman, Timothy A., *Principles of Instrumental Analysis*, Saunders College Publishing, Harcourt Brace College Publishers, Philadelphia, PA, 1998.

Soleyn, Ken, “Development of a Tunable Diode Laser Absorption Spectroscopy Moisture Analyzer for Natural Gas”, GE Sensing and Inspection Technologies, Rotterdam, 2009.

“Standard Operating Procedures – Thermo Environmental Instruments Model 43C Trace Level Pulsed Fluorescence Sulfur Dioxide Analyzer”, version 2.0, Environmental Protection Agency, Research Triangle Park, NC, 2009.

Theory and Practice of pH Measurement, PN 44-6033, Rosemount Analytical, 1999.

“XSTREAM Gas Analyzer Series Instruction Manual”, document HASX2E-IM-HS, Rosemount Analytical, 2009.

Chapter 24

Machine vibration measurement

Unlike most process measurements, the measurement of a rotating machine's *vibration* is primarily for the benefit of the process equipment rather than the process itself. Vibration monitoring on an ammonia vapor compressor, for instance, may very well be useful in extending the operating life of the compressor, but it offers little benefit to the control of the ammonia vapor.

Nevertheless, the prevalence of machine vibration measurement technology is so widespread in the process industries that it cannot be overlooked by the instrument technician. Rotating machinery equipped with vibration sensors are often controlled by *protection* equipment designed to automatically shut down the machine in the event of excessive vibration. The configuration and maintenance of this protection equipment, and the sensors feeding vibration data to it, is often the domain of instrument technicians.



24.1 Vibration physics

One very convenient feature of waves is that their properties are universal. Waves of water in the ocean, sound waves in air, electronic signal waveforms, and even waves of mechanical vibration may all be expressed in mathematical form using the trigonometric *sine* and *cosine* functions. This means the same tools (both mathematical and technological) may be applied to the analysis of different kinds of waves. A strong example of this is the *Fourier Transform*, used to determine the frequency spectrum of a waveform, which may be applied with equal validity to any kind of wave¹.

¹The “spectrum analyzer” display often seen on high-quality audio reproduction equipment such as stereo equalizers and amplifiers is an example of the Fourier Transform applied to music. This exact same technology may be applied to the analysis of a machine's vibration to indicate sources of vibration, since different components of a machine tend to generate vibratory waves of differing frequencies.

24.1.1 Sinusoidal vibrations

If a rotating wheel is unbalanced by the presence of an off-center mass, the resulting vibration will take the form of a cosine wave as measured by a displacement (position) sensor near the periphery of the object (assuming an angle of zero is defined by the position of the displacement sensor). The displacement sensor measures the air gap between the sensor tip and the rim of the spinning wheel, generating an electronic signal (most likely a voltage) directly proportional to that gap:



Since the wheel's shaft "bows" in the direction of the off-center mass as it spins, the gap between the wheel and the sensor will be at a minimum at 0° , and maximum at 180° .

We may begin to express this phenomenon mathematically using the cosine function:

$$x = D \cos \omega t + b$$

Where,

x = Displacement as measured by sensor at time t

D = Peak displacement amplitude

ω = Angular velocity (typically expressed in units of radians per second)

b = “Bias” air gap measured with no vibration

t = Time (seconds)

Since the cosine function alternates between extreme values of $+1$ and -1 , the constant D is necessary in the formula as a coefficient relating the cosine function to peak displacement. The cosine function’s argument (i.e. the angle given to it) deserves some explanation as well: the product ωt is the multiple of angular velocity and time, angular velocity typically measured in radians per second and time typically measured in seconds. The product ωt , then, has a unit of radians. At time=0 (when the mass is aligned with the sensor), the product ωt is zero and the cosine’s value is $+1$.

For a wheel spinning at 1720 RPM (approximately 180.1 radians per second), the angle between the off-center mass and the sensor will be as follows:

Time	Angle (radians)	Angle (degrees)	$\cos \omega t$
0 ms	0 rad	0°	$+1$
8.721 ms	$\frac{\pi}{2}$ rad	90°	0
17.44 ms	π rad	180°	-1
26.16 ms	$\frac{3\pi}{2}$ rad	270°	0
34.88 ms	0 rad	360° or 0°	$+1$

We know from physics that *velocity* is the time-derivative of displacement. That is, velocity is defined as the rate at which displacement changes over time. Mathematically, we may express this relationship using the calculus notation of the derivative:

$$v = \frac{dx}{dt} \quad \text{or} \quad v = \frac{d}{dt}(x)$$

Where,

v = Velocity of an object

x = Displacement (position) of an object

t = Time

Since we happen to know the equation describing displacement (x) in this system, we may differentiate this equation to arrive at an equation for velocity:

$$v = \frac{dx}{dt} = \frac{d}{dt}(D \cos \omega t + b)$$

Applying the differentiation rule that the derivative of a sum is the sum of the derivatives:

$$v = \frac{d}{dt}(D \cos \omega t) + \frac{d}{dt}b$$

Recall that D , ω , and b are all constants in this equation. The only variable here is t , which we are differentiating with respect to. We know from calculus that the derivative of a simple cosine function is a negative sine ($\frac{d}{dx} \cos x = -\sin x$), and that the presence of a constant multiplier in the cosine's argument results in that multiplier applied to the entire derivative² ($\frac{d}{dx} \cos ax = -a \sin ax$). We also know that the derivative of any constant is simply zero ($\frac{d}{dx}C = 0$), which eliminates the b term:

$$v = -\omega D \sin \omega t$$

What this equation tells us is that for any given amount of peak displacement (D), the velocity of the wheel's "wobble" increases linearly with speed (ω). This should not surprise us, since we know an increase in rotational speed would mean the wheel displaces the same vibrating distance in less time, which would necessitate a higher velocity of vibration.

We may take the process one step further by differentiating the equation again with respect to time in order to arrive at an equation describing the vibrational *acceleration* of the wheel's rim, since we know acceleration is the time-derivative of velocity ($a = \frac{dv}{dt}$):

$$a = \frac{dv}{dt} = \frac{d}{dt}(-\omega D \sin \omega t)$$

From calculus, we know that the derivative of a sine function is a cosine function ($\frac{d}{dx} \sin x = \cos x$), and the same rule regarding constant multipliers in the function's argument applies here as well ($\frac{d}{dx} \sin ax = a \cos ax$):

$$a = -\omega^2 D \cos \omega t$$

What this equation tells us is that for any given amount of peak displacement (D), the acceleration of the wheel's "wobble" increases with the *square* of the speed (ω). This is of great importance to us, since we know the lateral force imparted to the wheel (and shaft) is proportional to the lateral acceleration and also the mass of the wheel, from Newton's Second Law of Motion:

$$F = ma$$

²This rule makes intuitive sense as well: if a sine or cosine wave increases frequency while maintaining a constant peak-to-peak amplitude, the rate of its rise and fall *must* increase as well, since the higher frequency represents less time (shorter period) for the wave to travel the same amplitude. Since the derivative is the *rate of change* of the waveform, this means the derivative of a waveform must increase with that waveform's frequency.


Therefore, the vibrational force experienced by this wheel grows rapidly as rotational speed increases:

$$F = ma = -m\omega^2 D \cos \omega t$$


This is why vibration can be so terribly destructive to high-speed rotating machinery. Even a small amount of lateral displacement caused by a mass imbalance or other effect may generate enormous forces on the rotating part(s), as these forces grow with the square of the rotating speed (e.g. doubling the speed quadruples the force; tripling the speed increases force by *9 times*). Worse yet, these proportions assume a constant displacement (D), which is a best-case scenario. More realistically, we may expect the displacement to actually *increase*, as the centrifugal force generated by the off-center mass bends the rotating shaft to place the mass even farther away from the shaft centerline. Thus, doubling or tripling an imbalanced machine's speed may multiply vibrational forces well in excess of four or nine times, respectively.

In the United States, it is customary to measure vibrational displacement (D) in units of *mils*, with one "mil" being $\frac{1}{1000}$ of an inch (0.001 inch). Vibrational velocity is measured in inches per second, following the displacement unit of the inch. Acceleration, although it could be expressed in units of inches per second squared, is more often represented in the unit of the G : a multiple of Earth's own gravitational acceleration.

To give perspective to these units, it is helpful to consider a real application. Suppose we have a rotating machine vibrating in a sinusoidal (sine- or cosine-shaped) manner with a peak displacement (D) of 2 mils (0.002 inch) at a rotating speed of 1720 RPM (revolutions per minute). The frequency of this rotation is 28.667 Hz (revolutions per *second*), or 180.1 radians per second:



If D is the peak displacement of the sinusoid, then ωD must be the peak velocity (maximum rate-of-change over time) of the sinusoid³. This yields a peak velocity of 0.360 inches per second:



We may apply differentiation once more to obtain the acceleration of this machine's rotating element. If D is the peak displacement of the sinusoid, and ωD the peak velocity, then $\omega^2 D$ will be its peak acceleration.

$$D = \text{Peak displacement} = 0.002 \text{ in}$$

$$\omega D = \text{Peak velocity} = 0.360 \text{ in/s}$$

$$\omega^2 D = \text{Peak acceleration} = 64.9 \text{ in/s}^2$$

The average value of Earth's gravitational acceleration (g) is 32.17 feet per second squared. This equates to about 386 inches per second squared. Since our machine's peak vibrational acceleration is 64.9 inches per second squared, this may be expressed as a "G" ratio to Earth's gravity:

$$\frac{64.9 \text{ in/s}^2}{386 \text{ in/s}^2} = 0.168 \text{ G's of peak acceleration}$$

Using "G's" as a unit of acceleration makes it very easy to calculate forces imparted to the rotating element. If the machine's rotating piece weighs 1200 pounds (in 1 "G" of Earth gravity), then the force imparted to this piece by the vibrational acceleration of 0.168 G's will be 16.8% of its weight, or 201.7 pounds.

³Recall that the derivative of the sinusoidal function $\sin \omega t$ is equal to $\omega \cos \omega t$, and that the second derivative of $\sin \omega t$ is equal to $-\omega^2 \sin \omega t$. With each differentiation, the constant of angular velocity (ω) is applied as a multiplier to the entire function.

24.1.2 Non-sinusoidal vibrations

Normal machine vibrations rarely take the form of perfect sinusoidal waves. Although typical vibration waveforms are periodic (i.e. they repeat a pattern over time), they usually do not resemble sine or cosine waves in their shape:

A periodic, non-sinusoidal waveform



An unfortunate quality of non-sinusoidal waveforms is that they do not lend themselves as readily to mathematical analysis as sinusoidal waves. From the previous discussion on sinusoidal vibrations, we saw how simple it was to take the derivative of a sinusoidal waveform ($\frac{d}{dt} \sin \omega t = \omega \cos \omega t$), and how well this worked to predict velocity and acceleration from a function describing displacement. Most non-sinusoidal waveforms cannot be expressed as simply and neatly as $\sin \omega t$, however, and as such are not as easy to mathematically analyze.

Fortunately, though, there is a way to represent non-sinusoidal waveforms as combinations of sinusoidal waveforms. The French mathematician and physicist Jean Baptiste Joseph Fourier (1768-1830) proved mathematically that *any* periodic waveform, no matter how strange or asymmetrical its shape may be, may be replicated by a specific sum of sine and cosine waveforms of integer-multiple frequencies. That is, any periodic waveform (a periodic function of time, $f(\omega t)$ being the standard mathematical expression) is equivalent to a series of the following form⁴:

$$f(\omega t) = A_1 \cos \omega t + B_1 \sin \omega t + A_2 \cos 2\omega t + B_2 \sin 2\omega t + \cdots A_n \cos n\omega t + B_n \sin n\omega t$$

Here, ω represents the *fundamental* frequency of the waveform, while multiples of ω (e.g. 2ω , 3ω , 4ω , etc.) represent *harmonic* or *overtone* frequencies of that fundamental. The A and B coefficients describe the *amplitudes* (heights) of each sinusoid. We may break down a typical Fourier series in table form, labeling each term according to frequency:

Terms	Harmonic	Overtone
$A_1 \cos \omega t + B_1 \sin \omega t$	1st harmonic	Fundamental
$A_2 \cos 2\omega t + B_2 \sin 2\omega t$	2nd harmonic	1st overtone
$A_3 \cos 3\omega t + B_3 \sin 3\omega t$	3rd harmonic	2nd overtone
$A_4 \cos 4\omega t + B_4 \sin 4\omega t$	4th harmonic	3rd overtone
$A_n \cos n\omega t + B_n \sin n\omega t$	n th harmonic	$(n - 1)$ th overtone


One of the most visually convincing examples of Fourier's theorem is the ability to describe a square wave as a series of sine waves. Intuition would suggest it is impossible to synthesize a sharp-edged waveform such as a square wave using nothing but rounded sinusoids, but it is indeed possible if one combines an *infinite* series of sinusoids of successively higher harmonic frequencies, given just the right combination of harmonic frequencies and amplitudes.

⁴There is an additional term missing in this Fourier series, and that is the "DC" or "bias" term A_0 . Many non-sinusoidal waveforms having peak values centered about zero on a graph or oscilloscope display actually have *average* values that are non-zero, and the A_0 term accounts for this. However, this is usually not relevant in discussions of machine vibration, which is why I have opted to present the simplified Fourier series here.

The Fourier series for a square wave is as follows:


$$\text{Square wave} = 1 \sin \omega t + \frac{1}{3} \sin 3\omega t + \frac{1}{5} \sin 5\omega t + \frac{1}{7} \sin 7\omega t + \dots$$

Such a series would be impossible to numerically calculate, but we may approximate it by adding several of the first (largest) harmonics together to see the resulting shape. In each of the following plots, we see the individual harmonic waveforms plotted in red, with the sum plotted in blue:




If we continue this pattern up to the 13th harmonic (following the same pattern of diminishing reciprocal amplitudes shown in the Fourier series for a square wave), we see the resultant sum looking more like a square wave:

1st + 3rd + 5th + 7th + 9th + 11th + 13th



Continuing on to the 35th harmonic, the resultant sum looks like a square wave with ripples at each rising and falling edge:


All odd-numbered harmonics up to the 35th



If we were to continue adding successive terms in this infinite series, the resulting superposition of sinusoids would look more and more like a perfect square wave.

The only real question in any practical application is, “What are the A , B , and ω coefficient values necessary to describe a particular non-periodic waveform using a Fourier series?” Fourier’s theorem tells us we should be able to represent *any* periodic waveform – no matter what its shape – by summing together a particular series of sinusoids of just the right amplitudes and frequencies, but actually determining those amplitudes and frequencies is a another matter entirely. Fortunately, modern computational techniques such as the *Fast Fourier Transform* (or *FFT*) algorithm make it very easy to sample any periodic waveform and have a digital computer calculate the relative amplitudes and frequencies of its constituent harmonics. The result of a FFT analysis is a summary of the amplitudes, frequencies, and (in some cases) the phase angle of each harmonic.

To illustrate the relationship between a waveform plotted with respect to time versus a Fourier analysis showing component frequencies, I will show a pair of Fourier spectrum plots for two waveforms – one a perfect sinusoid and the other a non-sinusoidal waveform. First, the perfect sinusoid:




Fourier spectra are often referred to as *frequency-domain* plots because the x-axis (the “domain” in mathematical lingo) is frequency. A standard oscilloscope-type plot is called a *time-domain* plot because the x-axis is time. In this first set of plots, we see a perfect sine wave reduced to a single peak on the Fourier spectrum, showing a signal with only one frequency (the fundamental, or 1st harmonic). Here, the Fourier spectrum is very plain because there is only one frequency to display. In other words, the Fourier series for this perfect sinusoid would be:

$$f(\omega t) = 0 \cos \omega t + 1 \sin \omega t + 0 \cos 2\omega t + 0 \sin 2\omega t + \dots + 0 \cos n\omega t + 0 \sin n\omega t$$

Only the B_1 coefficient has a non-zero value. All other coefficients are zero because it only takes one sinusoid to perfectly represent this waveform.

Next, we will examine the Fourier analysis of a non-sinusoidal waveform:



In this second set of plots, we see the waveform is similar to a sine wave, except that it appears “clipped” at the peaks. This waveform is obviously not a perfect sinusoid, and therefore cannot be described by just one of the terms ($\sin \omega t$) in a Fourier series. It can, however, be described as equivalent to a *series* of perfect sinusoids summed together. In this case, the Fourier spectrum shows one sinusoid at the fundamental frequency, plus another (smaller) sinusoid at three times the fundamental frequency (3ω), plus another (yet smaller) sinusoid at the 5th harmonic and another (smaller still!) at the 7th: a series of *odd-numbered* harmonics.

If each of these harmonics is in phase with each other⁵, we could write the Fourier series as a set of sine terms:

$$f(\omega t) = (0 \text{ dB}) \sin \omega t + (-65 \text{ dB}) \sin 3\omega t + (-95 \text{ dB}) \sin 5\omega t + (-115 \text{ dB}) \sin 7\omega t$$

Translating the decibel amplitude values into simple coefficients, we can see just how small these harmonic sinusoids are in comparison to the fundamental:

$$f(\omega t) = 1 \sin \omega t + 0.000562 \sin 3\omega t + 0.0000178 \sin 5\omega t + 0.00000178 \sin 7\omega t$$

If the waveform deviated even further from a perfect sinusoid, we would see a Fourier spectrum with taller harmonic peaks, and perhaps more of them (possibly including some even-numbered harmonics, not just odd-numbered), representing a harmonically “richer” spectrum.

Within the technical discipline of machine vibration analysis, harmonic vibrations are often referred to by labels such as $1X$, $2X$, and $3X$, the integer number corresponding to the harmonic order of the vibration. The fundamental, or first harmonic, frequency of a vibration would be represented by “ $1X$ ” while “ $2X$ ” and “ $3X$ ” represent the second- and third-order harmonic frequencies, respectively.

⁵We have no way of knowing this from the Fourier spectrum plot, since that only shows us amplitude (height) and frequency (position on the x-axis).

On a practical note, the Fourier analysis of a machine's vibration waveform holds clues to the successful balancing of that machine. A first-harmonic vibration may be countered by placing an off-center mass on the rotating element 180 degrees out of phase with the offending sinusoid. Given the proper phase (180° – exactly opposed) and magnitude, any harmonic may be counterbalanced by an off-center mass rotating at the same frequency. In other words, we may cancel any particular harmonic vibration with an equal and opposite harmonic vibration.

If you examine the “crankshaft” of a piston engine, for example, you will notice counterweights with blind holes drilled in specific locations for balancing. These precisely-trimmed counterweights compensate for first-harmonic (fundamental) frequency vibrations resulting from the up-and-down oscillations of the pistons within the cylinders. However, in some engine designs such as inline 4-cylinder arrangements, there are significant harmonic vibrations of greater order than the fundamental, which *cannot* be counterbalanced by any amount of weight, in any location, on the rotating crankshaft. The reciprocating motion of the pistons and connecting rods produce periodic vibrations that are non-sinusoidal, and these vibrations (like all periodic, non-sinusoidal waveforms) are equivalent to a series of harmonically-related sinusoidal vibrations.


Any weight attached to the crankshaft will produce a first-order (fundamental) sinusoidal vibration, and that is all. In order to counteract harmonic vibrations of higher order, the engine requires counterbalance shafts spinning at speeds corresponding to those higher orders. This is why many high-performance inline 4-cylinder engines employ counterbalance shafts spinning at *twice* the crankshaft speed: to counteract the second-harmonic vibrations created by the reciprocating parts. If an engine designer were so inclined, he or she could include several counterbalance shafts, each one spinning at a different multiple of the crankshaft speed, to counteract as many harmonics as possible. At some point, however, the inclusion of all these shafts and the gearing necessary to ensure their precise speeds and phase shifts would interfere with the more basic design features of the engine, which is why you do not typically see an engine with multiple counterbalance shafts.

The harmonic content of a machine's vibration signal in and of itself tells us little about the health or balance of that machine. It may be perfectly normal for a machine to have a very “rich” harmonic signature due to convoluted motions of its parts⁶. However, Fourier analysis provides a simple way to quantify complex vibrations and to archive them for future reference. For example, we might gather vibration data on a new machine immediately after installation (including its Fourier spectra on all vibration measurement points) and save this data for safe keeping in the maintenance archives. Later, if and when we suspect a vibration-related problem with this machine, we may gather new vibration data and compare it against the original “signature” spectra to see if anything substantial has changed. Changes in harmonic amplitudes and/or the appearance of new harmonics may point to specific problems inside the machine. Expert knowledge is usually required to interpret the spectral changes and discern what those specific problem(s) might be, but at least this technique does have diagnostic value in the right hands.

⁶Machines with reciprocating components, such as pistons, cam followers, poppet valves, and such are notorious for generating vibration signatures which are anything but sinusoidal even under normal operating circumstances!

24.2 Vibration sensors

Sensors used to measure vibration come in three basic types: *displacement*, *velocity*, and *acceleration*. Displacement sensors measure changes in distance between a machine's rotating element and its stationary housing (frame). Displacement sensors come in the form of a probe that threads into a hole drilled and tapped in the machine's frame, just above the surface of a rotating shaft. Velocity and acceleration sensors, by contrast, measure the velocity or acceleration of whatever element the sensor is attached to, which is usually some external part of the machine frame⁷.




A design of displacement sensor manufactured by the Bently-Nevada corporation uses electromagnetic *eddy current* technology to sense the distance between the probe tip and the rotating machine shaft. The sensor itself is an encapsulated coil of wire, energized with high-frequency

⁷From the perspective of measurement, it would be ideal to affix a velocimeter or accelerometer sensor directly to the rotating element of the machine, but this leads to the problem of electrically connecting the (now rotating!) sensor to stationary analysis equipment. Unless the velocity or acceleration sensor is wireless, the only practical mounting location is on the stationary frame of the machine.

alternating current (AC). The magnetic field produced by the coil induces eddy currents in the metal shaft of the machine, as though the metal piece were a short-circuited secondary coil of a transformer (with the probe's coil as the transformer primary winding). The closer the shaft moves toward the sensor tip, the tighter the magnetic coupling between the shaft and the sensor coil, and the stronger the eddy currents.

The high-frequency oscillator circuit providing the sensor coil's excitation signal becomes loaded by the induced eddy currents. Therefore, the oscillator's load becomes a direct indication of how close the probe tip is to the metal shaft. This is not unlike the operation of a metal detector: measuring the proximity of a wire coil to any metal object by the degree of loading caused by eddy current induction.


In the Bently-Nevada design, the oscillator circuit providing sensor coil excitation is called a *proximitor*. The proximitor module is powered by an external DC power source, and drives the sensor coil through a coaxial cable. Proximity to the metal shaft is represented by a DC voltage output from the proximitor module, with 200 millivolts per mil ($1 \text{ mil} = \frac{1}{1000} \text{ inch}$) of motion being the standard calibration.




Since the proximitor's output voltage is a direct representation of distance between the probe's tip and the shaft's surface, a "quiet" signal (no vibration) will be a pure DC voltage. The probe is adjusted by a technician such that this quiescent voltage will lie between the proximitor's output voltage range limits. Any vibration of the shaft will cause the proximitor's output voltage to vary in precise step. A shaft vibration of 28.67 Hz, for instance, will cause the proximitor output signal to be a 28.67 Hz waveform superimposed on the DC "bias" voltage set by the initial probe/shaft gap.

An oscilloscope connected to this output signal will show a direct representation of shaft vibration, as measured in the axis of the probe. In fact, *any* electronic test equipment capable of analyzing the voltage signal output by the proximator may be used to analyze the machine's vibration: oscilloscopes, spectrum analyzers, peak-indicating voltmeters, RMS-indicating voltmeters, etc.


Photographs of a Bently-Nevada displacement sensor (sensing axial vibration on a “ring” style air compressor) and two proximator modules are shown here:




It is customary to arrange a set of three displacement probes at the end of a machine shaft to measure vibration: two *radial* probes and one *axial* (or *thrust*) probe. The purpose of this *triaxial* probe configuration is to measure shaft vibration (and/or shaft displacement) in all three dimensions:



The following photograph shows two displacement probes sensing vibration in the X and Y radial axes for a large vertical-shaft hydroelectric power plant turbine at Grand Coulee Dam:



It is also common to see one *phase reference* probe installed on the machine shaft, positioned in such a way that it detects the periodic passing of a keyway or other irregular feature on the shaft. The “keyphasor” signal will consist of one large pulse per revolution:



The purpose of a keyphasor signal is two-fold: to provide a reference point in the machine’s rotation to correlate other vibration signals against, and to provide a simple means of measuring shaft speed. The location in time of the pulse represents shaft position, while the frequency of that pulse signal represents shaft speed.

For instance, if one of the radial displacement sensors indicates a high vibration at the same frequency as the shaft rotation (i.e. the shaft is bowed in one direction, like a banana spinning on its long axis), the phase shift between the vibration’s sinusoidal peak and the phase reference pulse will indicate to maintenance machinists where the machine is out of balance. This is not unlike automatic tire-balancing machines designed to measure imbalance in automobile tire and wheel assemblies: the machine must have some way of indicating to the human operator *where* a balancing weight should be placed, not just how far out of balance the tire is. In the case of machine vibration monitoring equipment, the keyphasor signal and one of the axial displacement signals may be simultaneously plotted on a dual-trace oscilloscope for the purposes of determining the position of the imbalance on the machine shaft.


24.3 Monitoring hardware

The following photograph shows a large air blower in a wastewater treatment facility equipped with a Bently-Nevada model 3300 vibration monitoring rack (located left-center on the foreground panel):




Five vibration measurement and display cards are installed in this rack, each card capable of processing up to two displacement sensor signals. A six-channel temperature monitor card is also installed in the rack, used to display bearing and other machine component temperatures. Like the vibration cards, the temperature card is capable of generating both “alert” and “trip” signals, monitoring the presence of slightly abnormal conditions and taking automatic shut-down action in the event of excessively abnormal conditions, respectively.

A closer view of a different Bently-Nevada model 3300 vibration monitoring rack is shown in this photograph:



Each “card” inserted into this rack performs a different measurement function.

The following photographs show even closer views of the cards, revealing the display bargraphs and the units of measurement. From left to right; thrust measurement, vibration measurement, temperature measurement (6 channels), and speed measurement:



BNC-style cable connectors on the front of the cards provide convenient connection points for electronic test equipment such as oscilloscopes or spectrum analyzers. This eliminates the need to un-do wire connections on the proximator units in order to take diagnostic measurements. Each card also provides “alert” and “danger” levels for their respective measurements, generating a contact-closure signal which may be connected to an automatic shutdown (“trip”) system to take protective action if vibration or thrust displacement ever exceeds pre-set limits.

Another variety of vibration monitoring hardware is the Bently-Nevada 1701 FieldMonitor. This hardware lacks the convenient front-panel displays of the model 3300, opting instead to communicate vibration data in digital form to an Allen-Bradley programmable logic controller (PLC). Not only does this make it possible to display the vibration data remotely through HMI (Human-Machine Interface) panels, but it also enables vibration data to engage automatic “trip” logic programming in the PLC to shut the machine down in the event of excessive vibration. This next photograph shows several FieldMonitor modules plugged into a rack, acquiring displacement data from eight proximity probes (X and Y axis radial measurements at three machine bearing locations, plus one axial (thrust) measurement and one phase reference measurement):





24.4 Mechanical vibration switches

A much simpler alternative to continuous vibration sensors (displacement or acceleration) and monitoring equipment suitable for less critical applications is a simple mechanical switch actuated by a machine's vibration. These switches cannot, of course, quantitatively analyze machine vibrations, but they do serve as qualitative indicators of gross vibration.

The following photograph shows a Robertshaw "Vibraswitch" unit:




This switch works on the principle of a weighted lever generating a force when shaken. A pair of magnets located at the weighted end of the lever hold it in either the "reset" (normal) or "tripped" position:



When reset, the lever is pre-loaded by spring tension to flip to the "tripped" position. All it needs to make that transition is enough acceleration to generate the "breakaway" force necessary to pull away from the holding magnet. Once the acceleration force exceeds that threshold, the lever moves toward the other magnet, which holds it securely in position so that switch will not "reset" itself with additional vibration.

This pre-loading spring is adjustable by a small screw, making it possible to easily vary the sensitivity of the switch:



24.5 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Newton's Second Law of motion:** $F = ma$, describing how the acceleration of an object is directly proportional to the amount of applied (resultant) force and inversely proportional to its mass. Relevant to the calculation of force developed on a machine part from the acceleration and deceleration of vibration.
- **Differentiation (calculus):** where one variable is proportional to the rate-of-change of two others. Differentiation always results in a division (quotient) of units. Relevant to conversion from position to velocity, and from velocity to acceleration: $v = \frac{dx}{dt}$ and $a = \frac{dv}{dt}$.
- **Integration (calculus):** where one variable is proportional to the accumulation of the product of two others. Integration always results in a multiplication of units. Relevant to conversion from acceleration to velocity, and from velocity to position: $v = \int a dt$ and $x = \int v dt$.
- **Fourier transforms:** any repetitive waveform, no matter what its shape, is mathematically equivalent to a series of sinusoidal (sine and cosine) waves of different frequencies, amplitudes, and phase shifts added together. The frequencies of these sinusoids are all integer multiples, called harmonics. Relevant to decomposing vibrational wave signals into their constituent harmonic frequencies, to determine which parts of a machine are vibrating most.

References

Kaplan, Wilfred, *Advanced Mathematics for Engineers*, Addison-Wesley Publishing Company, Reading, MA, 1981.

Smith, Steven W., *The Scientist and Engineer's Guide to Digital Signal Processing*, California Technical Publishing, San Diego, CA, 1997.

White, Glenn D., *Introduction to Machine Vibration*, version 1.76, part number 8569, DLI Engineering Corp., Bainbridge Island, WA, 1995.

Chapter 25

Electric power measurement and control

Electrical power is a commodity in the modern world, bought and sold on the open market like any other. Thus, it is important to be able to measure and control electricity, not only for reasons of efficiency but also for sale, taxation, safety, equipment protection, and reliability of service.


As with any other quantity we wish to measure and control, the systems designed for these purposes may be divided into three general categories: *sensors* to measure, *final control elements* to exert influence, and *controllers* to make the necessary control decisions automatically. This chapter will discuss all three of these categories as they relate to measurement and protection subsystems found in modern electrical power grids.

This chapter cannot in any way do justice to the scope and complexity of electrical power grids. What it aims to do, however, is focus on the monitoring and protective functions subsystems essential to any functional power grid – the *instrumentation* within an electrical power grid, as it were – touching on the function of various pieces of electrical equipment as necessary to understand the purpose and application of those monitoring and protective subsystems.

25.1 Introduction to power system automation

Those familiar with industrial instrumentation will find much within the electric power industry remarkably familiar in concept. In industrial instrumentation we apply principles of physics, electricity, and chemistry to the measurement and automation of a wide range of “processes”. In the electric power industry the main “process” is the flow of electrical energy across long distances, but within that main process are a multitude of smaller processes with their own sensors, final control elements, and computation/control devices.


Within each of those smaller processes in a large electrical power system there exist automatic monitoring and control systems very similar to industrial process controls. A general block diagram showing the essential components of a feedback control system (used elsewhere in this book) applies to electrical power system automation as well:



Measurement devices in an electrical power system usually take the form of *instrument transformers* designed to represent high voltage and high current quantities as smaller, proportionate electrical signals. Controllers take the form of *protective relays* and other control systems designed to display and record the measured quantities, as well as take automatic control action. Final control is generally realized in the form of *circuit breakers* designed to redirect power flow and/or isolate sections of the power system.

Modern electrical power automation systems, like industrial automation, also employ sophisticated digital communication subsystems to exchange critical data such as power flow and fault diagnosis across wide regions.

Let us examine electric power *substations* as an example of automation. A “substation” is to an electrical power system what an intersection is to a system of highways and streets: a place where multiple paths intersect and flows are directed to their intended destinations. Just as road maps are used to graphically represent roads and intersections, *single-line diagrams*¹ are used to represent power lines and substation components. An example of a single-line diagram showing multiple substations appears here:



Electrical *generators* appear as circles with the letter “G” inside. *Loads* also appear as circles, but labeled uniquely. *Circuit breakers* (used to interrupt the flow of power during full-load and fault conditions) appear as squares, shown here with color-coded states². *Disconnect switches* (used


¹Single-line electrical diagrams are similar to Process Flow Diagrams (PFDs) used in industrial instrumentation, concentrating on the process flows more than the monitoring and control equipment. It is important to note that single-line diagrams are not the same as electrical schematics: in a single-line diagram, each line represents a *set* of power conductors (typically three or four conductors if the power system is 3-phase, which most large-scale AC power systems are). For this reason, we must interpret a single-line diagram much more like a *pipeline* system than an electrical circuit, in that the electrical power flows in one direction at any given time through these single lines, never making a complete loop as is the case in real life and in an electrical schematic diagram.

²In the electrical power industry, the color red universally represents an energized (closed breaker) condition while the color green represents a de-energized (open breaker) condition.

to isolate components from power during maintenance operations) appear as standard schematic switch symbols: a line broken by a diagonal line segment. Short line segments joining circuit breakers and disconnects with other devices in each station represent *busses*, which are sets of rigid metal conductors suspended by insulators. Longer lines connecting stations to each other represent transmission or distribution *power lines*. *Transformers*, used to step voltage up and current down for efficient long-distance transmission, or vice-versa for distribution and end-use, appear as standard schematic winding symbols. All these devices appear on single-line diagrams with single lines showing the route for power into and/or out of the device, rather than showing all electrical conductors connecting with the actual devices. This simplification is similar to the way road maps show streets and highways as single lines but generally do not show the number of lanes within each road.

Within each of these substations you can see circuit breakers and disconnect switches used to route the flow of electricity from sources to loads. These devices are analogous to control valves and block valves used to control fluid motion in industrial processes. Each circuit breaker, as a “final control element” in an automated system, may be commanded to open (trip) and/or close either by human action or by automatic action through special controllers called *protective relays* designed to protect the power system against damage caused by faults such as downed power lines, lightning strikes, and insulator breakdown. These protective relays sense voltage and current conditions through *instrument transformers* stepping high voltage down to safe sensing levels (Potential Transformers, or PTs) and stepping line current down to safe levels (Current Transformers, or CTs).


An example of a single-line diagram showing such an automated protection system for one of the power transformers in this system appears here:



Each protective relay function appears as a small circle enclosing a number, representing an industry-standardized code for that protective function (e.g. 50 = instantaneous overcurrent, 51 = time overcurrent, 86 = lockout, 87 = differential current). Solid lines show power and analog signal wiring, while dashed lines show control (relay output) wiring which are typically discrete (on/off) signals. In this particular case, any condition of overcurrent or current imbalance for this transformer causes the lockout relay (86) to trip, which in turn commands both line and load circuit breakers to trip, isolating the power transformer and thereby protecting it from harm. A single

potential transformer array (PT) steps down the high line voltage to a safe level (typically 120 volts nominal) for the wattmeter to read. A set of current transformers (CTs) step line current down to safe levels (typically 5 amps at full load) for the wattmeter and protective relays to read. As you can see, the disconnect switches have no connection to the automated system because they are manually-controlled devices, analogous to manual block valves flanking an automatic control valve in a process pipe.

Now that we know the functions of instrument transformers, protective relays, circuit breakers, and disconnect switches, we may examine some photographs of these power system components. First, we will examine some potential transformers (PTs), sometimes referred to as *voltage transformers* (VTs). The left-hand photograph shows a set of three PTs, each one used to sense phase-to-ground voltage in a 115 kV 3-phase power bus within a substation. The right-hand photograph shows a single PT sensing phase-to-phase voltage (i.e. line voltage) for a 13.8 kV bus within a substation:



Each of these PTs step high voltage down to a nominal value of 120 volts for direct meter indication and/or protective relay input signals. This is analogous to the industrial instrumentation signal standard of 4-20 mA representing such things as pressure, flow, and temperature: a relatively small electrical signal is used as a representation of some other real-world measurement. As the power system voltage rises as falls, these PTs' voltage output signals will rise and fall proportionate to the turns ratio³ of each transformer.

³For example, a potential transformer (PT) constructed to step 13.8 kilovolts down to 120 volts for safe monitoring of that line voltage must have a turns ratio equivalent to 13800:120, or 115:1.

Next we will examine some current transformers (CTs). The left-hand photograph shows a current transformer with a 400:5 amp ratio, which means a line current of 400 amps AC passing through the horizontal metal bar will induce a secondary winding current of 5 amps AC available at the screw terminals on top of this CT. The middle photograph is another style of CT, often called a “donut” or “window” CT because it has a large hole in the center through which the power conductor is routed (as a single “turn” primary winding). The right-hand photograph shows a set of CTs used to measure current in a 500 kV substation, the CT being located at the very top inside the box, while a long insulator supports the CT and holds it several feet above ground level for safety (since 500 kilo-volts can “jump” a fair distance through air and therefore must be separated from the earth):



All of these CTs output a nominal current of 5 amps AC at full load rating, which is a common CT signal standard within the electrical power industry, like 120 volts is for PT output signals. As line current rises and falls, these CTs' signals will proportionately rise and fall according to the turns ratio⁴ of each CT. These 0 to 5 amp AC signals are wired to measuring and/or protection instruments located in the substation control building.

Together, PTs and CTs constitute the *primary sensing elements* of electrical power measurement, control, and protection systems. One of the tasks of metering and protection technicians in the electric power industry is to periodically check the accuracy and performance of these instrument transformers, just as an industrial instrument technician periodically checks the calibration of process sensing elements and transmitters.

⁴For example, a current transformer (CT) constructed to step 400 amps down to 5 amps for safe monitoring of that line current must have a turns ratio equivalent to 400:5, or 80:1. This means the single “turn” of the power conductor through the center of the CT is flanked by exactly 80 turns of wire wrapped around the toroidal iron core of the CT.


Next we will examine some of the panel-mounted instruments receiving signals from PTs and CTs. First are simple meters, designed to display system measurements to human operators. These instruments are labeled with high ranges despite the fact that their actual driving signals are relatively small (e.g. 0 to 120 volts for voltage instruments and 0 to 5 amps for current instruments). This next photograph shows a portion of an analog control panel with several meters registering voltage, current, and power factor⁵ values for a substation bus. Several indicating lamps (red and green) show the statuses of various circuit breakers in the substation yard, with L-handle control switches providing remote trip/close operation of circuit breakers and red plastic lines representing the single-line diagram of the substation:



⁵To review, the *power factor* of an AC circuit is the cosine of the phase angle between total (source) voltage and total (source) current. Power factor represents how much of the line current goes toward doing useful work. Reactive loads do not transform electrical energy into work, but rather alternately store and release electrical energy. Current at a purely reactive load, therefore, is not as useful as current at a purely resistive load. However, reactive current still “occupies” ampacity on a power line, and so the existence of a low power factor means the system is not delivering as much power as it could.

Modern SCADA (Supervisory Control And Data Acquisition) hardware designed for power systems also input PT and CT signals, displaying those values on computer monitors instead of analog meter movements. In addition to analog voltage and current signals, SCADA systems also input discrete signals from circuit breaker auxiliary contacts, disconnect switch status contacts, pressure switches, and other on/off sensing devices located near the high-voltage power conductors. This provides operators with remote viewing of device status, which is then displayed as different colors (red or green) on a graphic single-line diagram of the power system.

The following photograph shows a SCADA display screen of a large public utility power grid. The scale of this particular display is such that individual circuit breakers are not represented, showing entire substations as single colored squares. However, more detailed diagrams are viewable by selecting a particular substation on this screen, these detailed displays showing individual circuit breakers and other associated equipment within that substation:



SCADA systems utilize a variety of telecommunication pathways to distribute power system data over long distances, including microwave (radio), optical fiber, leased telephone lines, and sometimes even high-frequency AC signals superimposed⁶ on power line conductors.

⁶This legacy technology is called *Power Line Carrier*, or *PLC* which is unfortunately confusing because it has nothing to do with Programmable Logic Controllers (also abbreviated PLC). The concept is not unlike the HART analog-digital hybrid system used to communicate digital information to process transmitters over 4-20 mA analog signal lines, except in the case of power-line carrier systems the signal frequencies are much higher and the challenge of safely coupling these signals to high-voltage power line conductors is much greater.

Protective relays have been described as the “silent sentinels” of electric power systems, quietly monitoring voltage and/or current conditions, ready to spring into action to protect the system against damage from faults. These automatic control devices have existed in one form or another for over a century, beginning with crude electromechanical designs and now culminating in state-of-the-art microprocessor-based computing machines. Relay functions are commonly designated by numerical codes standardized by *ANSI*, some of which will be listed in this section.

A series of electromechanical protective relays appears in the following photographs, taken at a large substation. The left-hand photograph shows a pair of *distance* relays (ANSI code 21) designed to sense the electrical impedance⁷ of a long power line and its load, tripping the circuit breaker(s) supplying power to that line if a fault reduces the impedance to a value equal to or less than that of the line itself. The middle photograph shows a set of *transformer differential current* relays (ANSI code 87) designed to compare the amount of current in the primary and secondary windings of a transformer, tripping circuit breakers on both sides of the transformer in the event a transformer fault is detected (i.e. if the amount of current exiting the transformer does not proportionately match the amount of current entering it). The right-hand photograph shows a set of *overpressure* relays (ANSI code 63) designed to trip circuit breakers feeding power to a device if the pressure inside that device rises to unacceptable levels:



A routine task for relay technicians working on electromechanical relays is periodic *recalibration* of these devices. Since they contain potentiometers, magnets, inductive coils, and moving parts they are susceptible to calibration drift just like any other analog electronic or mechanical device.

⁷To review, *impedance* is the sum total opposition to electric current in a circuit, consisting of resistance and/or reactance. Impedance is measured in ohms, and so a distance relay (21) is set to “pick up” a fault in a power line if the measured impedance of that line falls below a threshold value based on the length of that line.

Modern digital electronic protective relays are also panel-mounted, but of course contain no moving parts and are much more capable in terms of their ability to discriminate between normal operating conditions and faulted conditions meriting the tripping of circuit breakers. The following photographs show some examples of these devices. The left-hand photograph shows a transformer protection relay, incorporating the *differential current* (ANSI code 87) protection of the previous electromechanical relay plus a number of other features including *instantaneous* and *time-overcurrent* functions. The right-hand photograph shows a pair of digital relays, the upper one providing *instantaneous overcurrent* (ANSI code 50) plus *time-overcurrent* (ANSI code 51) plus circuit breaker *reclosing* (ANSI code 79) functionality, tripping the circuit breaker in the event of excessive current⁸, and then re-closing that same circuit breaker a short moment after to check if the fault has cleared. The right-hand photograph shows a *directional overcurrent* relay (ANSI code 67) designed to sense excessive line current in one particular direction⁹ along the line, tripping the circuit breaker(s) feeding power to that line if a fault is detected:



One of the benefits of digital protective relays is their remarkable stability compared to electromechanical relays, being virtually immune to calibration drift. This translates to less routine maintenance for relay technicians.

Not only do modern protective relays perform their basic system protection functions, but they also *record* data for later retrieval and analysis by relay technicians and protection engineers. These relays, being microprocessor based, may also be interconnected using high-speed data networks to exchange data with each other as part of certain protection strategies. These new capabilities, coupled with the need to maintain accurate archives of digital relay configuration files, means the

⁸The difference between an instantaneous overcurrent (50) function and a time-overcurrent (51) function is the amount of time delay between the detection of an overcurrent event and the relay's command to trip the circuit breaker. Any detected level of line current in excess of the instantaneous overcurrent "pickup" threshold will *immediately* issue a trip command, while the level of line current in excess of the time-overcurrent "pickup" threshold will determine the amount of time delay before the issuance of a trip command.

⁹Directional relays are useful for protecting electrical generators susceptible to acting as a motor and drawing power from the network rather than delivering power to the network. Generators driven by wind turbines are an example of this class: even a relatively small amount of power flowing in reverse direction (from the grid to the generator, "motoring" the generator) is undesirable, and so it is wise to isolate a "motoring" generator based on a much lower current than what would be considered unacceptable in the generating direction. A regular 50 or 51 overcurrent relay cannot discriminate between the two directions of power flow, but a 67 overcurrent relay can.

job of the relay technician has evolved: there is less routine calibration work, but more routine record-keeping and high-level diagnostic work.

Finally, we have the final control elements of the electric power industry: *circuit breakers* and *disconnects*. These two types of devices are common in that they both serve to connect and disconnect portions of a power system. They differ in their ability to interrupt current: circuit breakers are built with very rugged electrical contacts capable of safely and reliably interrupting huge magnitudes of electric current (including currents arising from short-circuit faults in the power system), whereas disconnects are switches that cannot make or break such large currents, and are intended to be operated only when the series-connected circuit breaker is open (tripped).

The following photographs show sets of three-phase 115 kV disconnects, the left-hand photograph showing a set in the closed position and the right-hand photograph showing a set in the open position:



As you can see, a high-voltage disconnect is nothing more than an open-air knife switch. Some are manually operated (by a lever or a hand crank) while others use an electric motor for remote operation by a SCADA system or by an operator in a substation control room.

Medium-voltage and high-voltage circuit breakers come in a variety of shapes and sizes. Perhaps the most significant difference between them is the method(s) employed to extinguish the electric arc formed when the contacts separate to interrupt line current. Some circuit breakers have their contacts immersed in a tank of dielectric oil, while others use contacts sealed inside of vacuum chambers (where there is no gas at all to ionize and create an arc), or enclose their contacts in chambers filled with a special dielectric gas such as sulfur hexafluoride (SF_6), or use high-pressure jets of air to “blow out” the arc.

The photograph on the left shows a legacy oil-tank circuit breaker in a 115 kV substation yard, consisting of three separate tanks containing contacts to interrupt one phase each. All three contacts operate simultaneously by the same mechanism. The right-hand photograph shows a more modern circuit breaker design, this one in a 250 kV substation yard, enclosing its three contact sets in pressurized SF_6 gas:

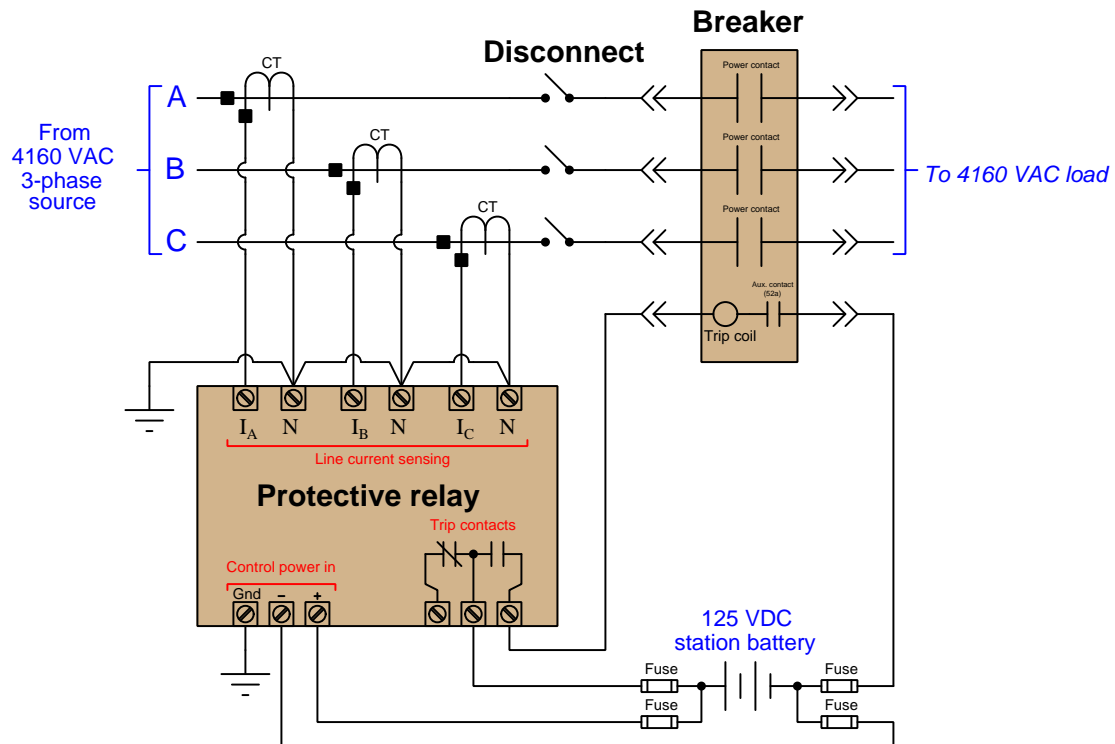


Safe and effective interruption of electric current at these elevated potentials demands quick contact action, and this is possible only with some form of stored-energy mechanism inside the circuit breaker. Some large circuit breakers use reservoirs filled with compressed air as the actuating medium, the reservoir maintained in a state of high pressure by an electric air compressor. Other circuit breakers use mechanical springs pre-charged by an electric motor and gear mechanism¹⁰. In all designs, though, the energy required to quickly close and open (trip) the circuit breaker contacts is

¹⁰These mechanisms are similar in principle to the trigger, spring, and hammer of a firearm: the mechanical energy necessary to ignite the primer of a cartridge comes from a spring that has been “charged” either by manual operation or by the action of the gun during the last firing cycle. This spring energy is released by a sensitive *sear* mechanism driven by the finger-operated trigger, requiring very little energy to operate. In a similar manner, the operating springs of large circuit breakers are “charged” by an electric motor whenever a relaxed state is detected. That mechanical energy is then released by a relatively sensitive mechanism driven by an electric solenoid, allowing a small electrical signal to rapidly operate the large contact mechanism.

provided by some energy-storage mechanism, that energy released on command by electric solenoid coils which may be remotely operated by human action and/or by protective relays.

Putting all these devices together and representing them in schematic form, we see the following example whereby a protective relay senses line current and issues a “trip” command to the medium-voltage (4160 volt AC) circuit breaker in the event of an overcurrent (fault) condition:



125 volt DC “station” power is used in this particular system for high reliability, ensuring the protective system will still be able to function even if an interruption occurs in AC power to the station. This 125 volt battery bank is maintained in a continuous state of charge by an AC-DC battery charger fed from the AC power source (not shown).

In summary, electric power systems employ automation to measure power conditions and take protective action when needed in the event of major line or device faults. These automated systems resemble industrial process control and safety systems in their three-part division (sensing, control, and final action) as well as in their graphical representation, calibration, and other maintenance.

25.2 Electrical power grids

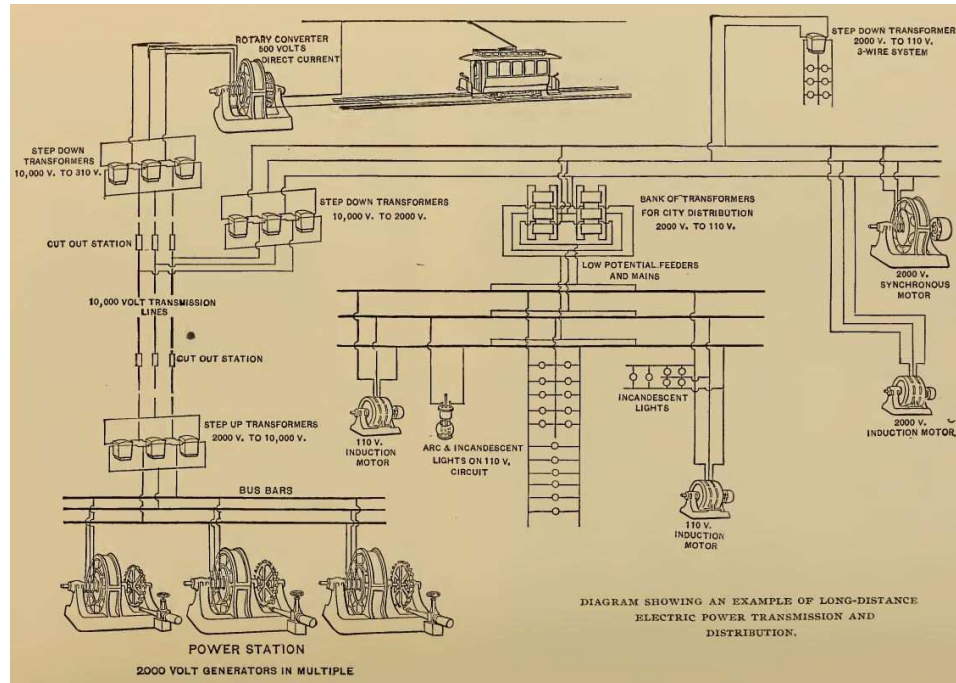
The term “grid” refers to the conductors and equipment interconnecting power sources to power loads in a wide-spread electrical system. Generating stations (i.e. “power plants”) convert various forms of energy such as fossil fuel, solar, wind, elevated water, and nuclear into electrical power; which is then sent through step-up transformers to raise the voltage and reduce current¹¹; conveyed long distances over “transmission lines” at voltages ranging in the hundreds of kilovolts; received by “substations” which serve as interconnecting hubs for generating stations and loads, using step-down transformers to reduce transmission line voltage and increase current for distribution to local loads; conveyed to industrial, commercial, and residential points of use over “distribution lines” at tens of kilovolts (or less); and then finally stepped down in voltage once more at points of use by transformers located at power customer sites. A myriad of circuit breakers, disconnect switches, fuses, and other devices serve to disconnect and re-route power as needed within the grid. Sensing instruments monitor the flow of power throughout the grid, for regulatory (control), billing (metering), and protection (shutdown) purposes.

The innovation of alternating current (AC) is what made large-scale electrical power grids possible. DC power – at least when implemented with 19th century technology – is prohibitively expensive¹² to transport over long distances due to the inability to easily transform voltage and current levels. AC by contrast allows the use of electromagnetic *transformers*, allowing voltage to be stepped up and current stepped down for economical transmission (i.e. small-gauge conductors suspended by long insulators), and then allowing voltage to be stepped down (for safety) with a proportional increase in current for driving heavy loads at points of use.

¹¹The sole purpose of transforming voltage and current levels in a power grid is to minimize power losses due to the electrical resistance of the conductors. Recall from basic DC electrical theory that the amount of power dissipated by a current-carrying resistance is $P = I^2R$. This means doubling the current through a resistive conductor will increase that conductor’s power dissipation four-fold, all other factors being equal. Metal wire is expensive, especially when thousands of miles of it must be run to form a power grid. In the interest of reducing this expense, transformers are used to maintain long-distance power line voltages high and currents low, permitting the use of smaller-gauge conductors to carry that current.

¹²Thomas Edison’s original DC-based power grid was limited in radius to the size of a city, because all components operated at one voltage level (about 110 VDC). Large copper busbars served as distribution lines from coal-fired generating stations to points throughout the city, the sheer mass of these copper bars necessitating their installation in underground trenches rather than as overhead lines. Voltage losses from the generating station to points at the furthest reaches of the DC grid were significant, meaning customers at the “end of the line” had to tolerate dimmer lamps than customers located nearer the generating station.

The following illustration shows a simplified schematic diagram of a polyphase AC power system dating from the year 1895¹³, showcasing an example of a Westinghouse power transmission and distribution system designed to generate power at Niagara Falls and send it some 60 miles distant:




This system was typical of its era, with one centralized generating station providing electrical power to all portions of the city. Electricity at the generators was stepped up in voltage to 10 kV to facilitate long-distance transmission over economically thin wires, then stepped back down through a series of transformers to voltages appropriate for direct use. Certain applications such as traction motors for electric streetcars required DC power rather than AC power (as efficient speed control for AC motors did not exist at that time), and given the lack of modern solid-state rectifier (diode) circuits the only practical solution was to employ motor-generator machines called *rotary converters* to convert AC into DC.

Modern electric power grids link dozens of large electric generating stations together with hundreds of cities to provide electricity across much larger geographic spans than the Westinghouse system of 1895. Modern power electronics has all but replaced rotary converters and made conversion between AC and DC relatively easy and efficient, allowing DC to now play a role in high-voltage power transmission. Some electrical power plants located in remote regions output DC along two-conductor transmission lines, which is later converted to AC and synchronized with AC generating stations on the grid.

¹³The source for this historical illustration is *Cassier's Magazine*, which was an engineering periodical published in the late 1800's and early 1900's out of London, England. The Smithsonian Institute maintains online archives of *Cassier's* spanning many years, and it is a treasure-trove for those interested in the history of mechanical, electrical, chemical, and civil engineering.

The following photograph shows a computer monitor plotting a trend over time of power sourced to and drawn from one electric utility company's service region, in this particular case it is the one operated by Puget Sound Energy in northwest Washington state:



Lines of differing color reveal power flow to and from this grid. The black trend represents total *load* on the system from all consumers: industrial, commercial, and residential. Each of the other colored trends represent different sources of power to the grid. Red represents thermal-based power plants, split between natural gas and coal energy sources. Yellow represents wind turbine power generating stations, with zero output between the times of 1:00 AM and 8:00 AM on the graph (i.e. negligible wind blowing during that timespan). Blue represents hydroelectric generators. Finally, green represents power purchased from other electric utilities (e.g. Bonneville Power Administration and Seattle City Light).

These trend graph reveals wind power to be a relatively small percentage of the total power generated during the timespan shown, with coal and hydroelectric sources remaining fairly constant over that time and the balance made up by power purchased from other utilities.

With no large-scale means to store excess electrical energy, the stability of an electrical grid depends on generating stations' ability to alter their power output to meet a continuously changing demand¹⁴. Some types of generating stations lend themselves better to rapid changes in power

¹⁴Other options may exist for some grids. For example, large-scale industrial customers may be requested to curtail their power consumption at certain times in order to offset a deficit in supply. An example of this might be an aluminum smelter (which uses hundreds of megawatts of electricity to reduce alumina powder to molten aluminum metal) operating as a sheddable load while the same grid employs a nuclear fission power plant as one of its sources. If the nuclear generator's reactor happens to "scram" (shut down for any reason), that reactor's power output will

output than others: single-cycle gas turbines, for example, excel at their ability to ramp power output up or down in relatively little time. Steam-based thermal generators such as coal-fired power plants are not so tolerant. Wind and solar generators operate at the whim of natural forces.


Needless to say, modern electrical power grids are highly complex systems. The critically important nature of a functioning power grid is difficult to overstate for any first-world nation, as so much depends on the uninterrupted flow of power from generating stations to customers. The complexity of power grids and their control systems will only increase over time as more renewable generation capacity (e.g. wind and solar) is brought on-line to reduce overall dependency on non-renewable energy sources. Large-scale energy storage systems will become a necessity for maintaining the stability of power grids dependent on power sources such as wind and solar that cannot be arbitrarily controlled.

drop off the grid immediately, which may constitute hundreds of megawatts of lost generation. In such an event, the grid dispatch system may issue a “load shed” command to the aluminum smelter to drop a substantial portion of its consumption, as it may not be practical to immediately bring that much extra power on-line from some other source.

25.3 Interconnected generators

Any power grid large enough to meet the demand of an entire nation must have multiple generators drawing from multiple energy sources supplying that power. Connecting these generators together so as to equitably share loads in the grid is no trivial task. The basic concepts involved with interconnecting generators, however, are independent of the distance separating those generators. Thus, the examples given in this section are as relevant to generators located adjacent to each other as they are to generators located hundreds of miles apart.


Let's begin with a simple electromechanical DC generator application, where one generator supplies power to one load. Although not shown in this electrical schematic, there will be some form of “prime mover” such as an engine or a turbine turning the mechanical shaft of the generator:




This system is almost too simple to warrant comment: the generator outputs a constant voltage, with load current being a function of load resistance in accordance with Ohm's Law ($I = \frac{V}{R}$). If load resistance happens to decrease and generator voltage remains constant, load current will increase as a result. This additional load current has the effect of making the electromechanical generator harder to turn at the same shaft speed: demanding greater mechanical power input to the generator in order¹⁵ to deliver greater electrical power to the load. In this way the generator naturally “senses” the power demanded by the load.

¹⁵This phenomenon is just one more application of the *Law of Energy Conservation*, which states energy cannot be created or destroyed, but must be accounted for in all processes. Every joule of energy delivered to the load in this example circuit must be supplied by the generator, which in turn draws (at least) one joule of energy from the prime mover (e.g. engine, turbine). Since the power “grid” shown in this diagram has no means of storing energy for future use, the load's demand must be instantaneously met by the generator, and in turn by the prime mover. Thus, sudden changes in load resistance result in instantaneous changes in power drawn from the prime mover, all in accordance with the Law of Energy Conservation.

Now, let us consider a second generator added in parallel to this simple power “grid”. In this configuration each of the two generators should contribute current to the grid, helping each other power the one load (resistor) shown on the right-hand side of the diagram:



We know from basic circuit theory that parallel-connected components share the same voltage. From this principle we may conclude that the two generators V_1 and V_2 will need to output the same amount of voltage in order to be compatible with each other in this circuit configuration. To further explore this concept, we will consider what would happen if the two generator voltages were unequal to each other, including the resistance of the line connecting the two generators together:



If the voltage of generator V_1 exceeds the voltage of generator V_2 , the difference of those two generators’ output voltages will be dropped along the length of the conductor R_{line} . Ohm’s Law allows us to predict the amount of current flowing through that line arising from the generators’ differing output voltages ($I = \frac{V_1 - V_2}{R_{line}}$). If the voltage difference is substantial and the line resistance is minimal, this current will be quite large. If this amount of current exceeds the amount drawn by the load, Kirchhoff’s Current Law tells us the current through generator V_2 will be going the wrong direction: *down* rather than *up*.

This situation is undesirable because it means generator V_2 will actually be functioning as a *load* rather than as the *source* it should be. Not only will generator V_2 not be contributing any power to the grid, but it will actually *draw* power away from generator V_1 that could otherwise go to the load. If generator V_2 is an electromechanical machine, it will operate as a motor as it draws power from generator V_1 : “motor” V_2 increasing speed and generator V_1 slowing down from the additional loading.


To summarize: parallel-connected DC generators must output the same voltage in order to equitably share the burden of powering loads. If one generator outputs less voltage than another, it will contribute less power. If this disparity is great enough, the weaker generator will actually become a load and begin to function as an electric motor rather than the generator (source) it should be.

Connecting AC generators in a power grid presents the same fundamental problem: all parallel-connected generators must output the same voltage to the “grid” in order to equitably share the load. What makes AC generator interconnections more complex than DC generator interconnections is the fact that the voltage output by an AC generator is not a static quantity but rather is oscillating sinusoidally¹⁶. This means paralleled AC generators must closely match one another’s output voltage *at every point along their sine-wave cycles* in order for them to productively work together on the grid. The only way two or more AC generators may continuously match one another’s output voltage is if their peak voltages are the same, their frequencies are the same, and they remain in-phase with each other.

AC generator frequency is a direct function of shaft speed. AC generator voltage is a direct function of shaft speed and rotor excitation current. Thus, in order to connect two or more AC generators together, these two parameters must be precisely controlled.

The process of ensuring an AC generator is ready to be connected to a live grid is called *synchronization*. This may be done manually by human operators, or automatically by synchronization relays. However it is done, the principle is the same: the voltage output by the un-synchronized generator is compared against the voltage of the grid, and the disconnecting switch or circuit breaker is not closed until the difference between those two voltages is nearly zero.

A simplified demonstration circuit serves to illustrate this process:




Imagine a condition where generator V_1 is operating and powering the load, but generator V_2 is stopped with its disconnecting switch in the open position. In this condition the lamp will glow steadily, operating on the difference of potential between the power line (full AC voltage) and the idle generator V_2 (0 volts).

If we bring generator V_2 up to speed while keeping the disconnect switch in its open position, regulating the generator V_2 ’s output voltage and frequency to be equal to generator V_1 , the lamp will experience a voltage strictly dependent on the degree of *phase shift* between the two generators. If, at this correct voltage and frequency, the phase of generator V_2 precisely matches the phase of generator V_1 (i.e. the sine-wave outputs of these two generators are in perfect lock-step with each other), the lamp will experience zero voltage and will remain dark. If the two generators happen to be exactly 180 degrees out of phase with each other while maintaining equal voltage and frequency, the lamp will experience a sinusoidal voltage at that same frequency having a peak value *twice* that of either generator, and will therefore glow at maximum brightness. If the amount of phase shift between these two generators is any value between 0° and 180° , the voltage experienced by the lamp will vary proportionately.

¹⁶The standard frequency for a power grid is typically 50 Hz or 60 Hz, depending on which part of the world you are in. North American power grids typically operate at 60 Hz, while 50 Hz is more common in Europe.

If you imagine the two generators' voltages being phasors joined at the tails, the amount of voltage seen by the lamp will be equivalent to the distance between those two phasors' tips¹⁷:



If the two generators output different frequencies, the phase shift angle between them will not be constant. Instead, the two generators will fall in and out of phase with each other at a frequency equal to the difference between the individual generator frequencies. For example, if V_1 outputs 240 volts peak at 60 Hz and V_2 outputs 240 volts peak at 59 Hz, the two generators will roll in and out of phase with each other once per second ($60 \text{ Hz} - 59 \text{ Hz} = 1 \text{ Hz}$). The lamp will thus experience an AC voltage that varies from 0 volts (when the two generators happen to be perfectly in phase) to 480 volts peak (when the two generators happen to be 180° out of phase). In visual terms this means the lamp will alternate from complete darkness to full brightness and back again once per second. Thus, the lamp's *oscillation* serves to indicate the difference in frequencies between the two generators.

If the two generators output different amounts of voltage *and* at different frequencies, the lamp will oscillate from bright to dim, never reaching full brightness or going completely dark. Thus, the amount of *variance in lamp intensity* serves to indicate the difference in voltage between the two generators.

The purpose of the lamp in this circuit, of course, is to indicate when it is safe to close the disconnect switch and tie generator V_2 to the power grid. Since we know paralleled generators are electrically compatible only when their output voltages match at all times, we look for a condition of complete lamp darkness before closing the disconnect switch.

Once the generators have been synchronized and the disconnect switch closed, an interesting phenomenon occurs: the two generators now behave as if their shafts were mechanically coupled¹⁸ together. This is similar to the phenomenon experienced with the two parallel-connected DC generators shown earlier, where an under-performing generator would begin to “motor” and draw power from the stronger generator if their voltages were sufficiently different. In the case of paralleled AC generators, a generator that begins to lag behind the other(s) in speed will act as a synchronous motor and draw power from the grid to match the speed of the other generator(s), staying in lock-step with the grid frequency so long as it is connected to the grid.

¹⁷A common analogy for this is two children swinging on adjacent swings in a playground. Imagine the distance between the children being the amount of voltage difference between the two generators at any given point in time, with the amplitude of each child's swing representing the peak voltage of each generator and the pace of each child's oscillation being the frequency of each generator. When two children are swinging in perfect synchronization, the distance between them remains minimal at all times. When they swing 180° out of phase each other, the distance between them varies from minimal to maximal at a pace equal to the difference in their individual swinging rates.

¹⁸This “coupling” is not perfectly rigid, but does allow for some degree of phase difference between the generator and the grid. A more accurate analogy would be to say the generators act as if their shafts were linked by a *flexible* coupling.

If an AC generator is synchronized and connected to the grid, and then its prime mover's power is increased in an attempt to increase the shaft speed, that generator will in effect be trying to force all the other AC generators on that grid to a faster speed. If the generator in question represents a small fraction of the grid's total power generating capacity as is typically the case, increasing or decreasing its speed becomes impossible¹⁹ due to this coupling effect.

25.4 Single-line electrical diagrams


Electrical power grids primarily consist of *three-phase* AC circuits. This means most power lines (transmission and distribution) have at least three conductors, and power transformers are either three-phase units or banks of single-phase transformers connected in Delta and/or Wye primary and secondary winding configurations. Of course, diagrams must be drawn to document how all these conductors and power components interconnect, and standard electrical schematics serve that purpose well at the equipment level. When analyzing power grids on the transmission or distribution scale, however, showing each and every conductor in electrical schematic form would make the system diagram needlessly complex.

For this reason electrical power grids are most commonly represented in a *single-line diagram* format. This means each transmission or distribution power line appears as a single line on the page, rather than as three (or four) lines showing individual conductors in a three-phase AC circuit. Single-line diagrams work well to analyze the general flow of electrical power from sources to loads.

¹⁹It should be noted that a grid-connected AC generator can in fact be over-spun with sufficient mechanical power input, but only if it “slips a pole” and falls out of synchronization as a result. Such an event can be catastrophically to the offending generator unless it is immediately disconnected from the grid to avoid damage from overcurrent.

The following schematic diagram represents a segment of an industrial power distribution system containing generators, power transformers, busses (sets of conductors used to connect multiple loads and/or sources in parallel with each other), instrument transformers²⁰ and meters, circuit breakers, motors, and motor-starting switches:


Schematic diagram representation:



²⁰In this example, three *current transformers*, or CTs, are shown stepping down the bus line current to levels safely measured by panel-mounted ammeters. Current transformers typically step down line current to a nominal value of 5 amps to drive meters, relays, and other monitoring instruments.


A single-line diagram of the same industrial power distribution system shows all the same components:


Single-line diagram representation:



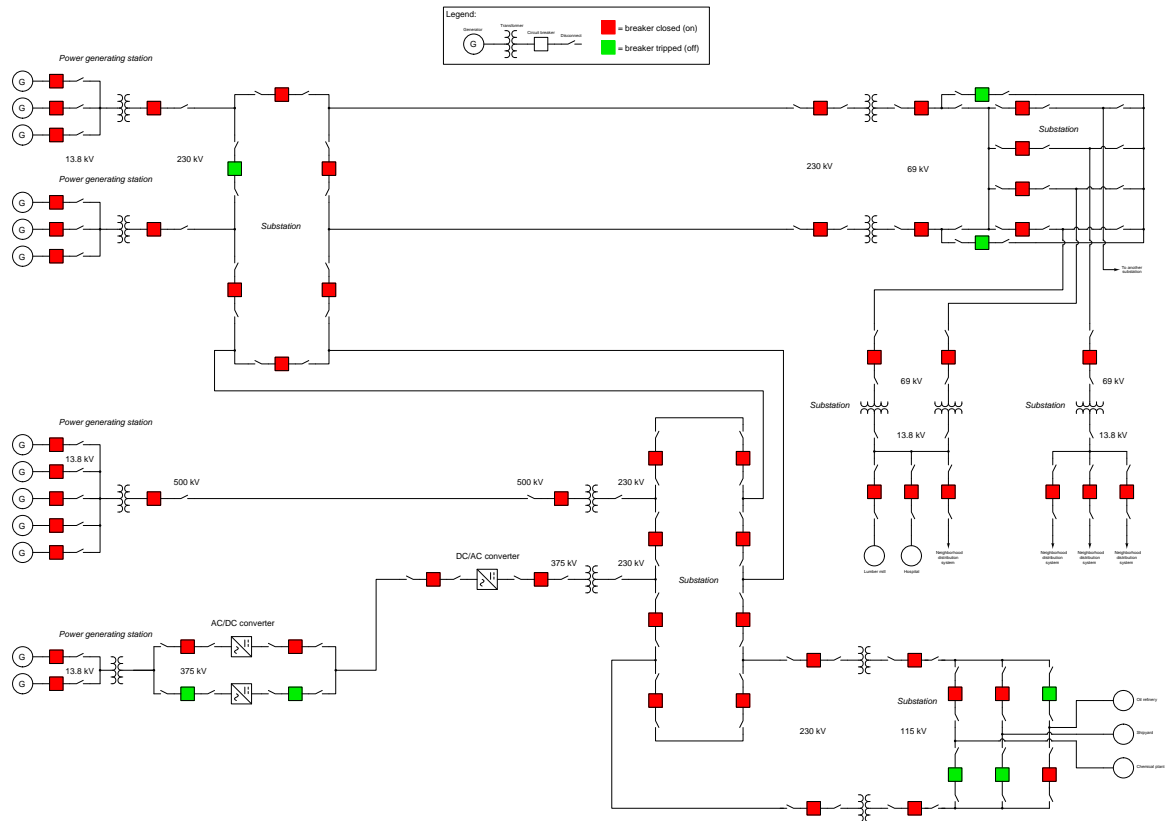
Note how much simpler and “cleaner” the single-line diagram is compared to the schematic diagram of the same power system: each three-conductor set of power wires is shown as a single line, each transformer appears as a single primary winding and single secondary winding (rather than three of each), each motor and generator is a simple circle rather than a complete set of windings, motor starter contacts aren’t triplicated, current transformers and ammeters appear as single units instead of triplets, and each three-pole circuit breaker appears as either a single square or a single breaker symbol. For those familiar with industrial instrumentation and control system diagrams, the distinction between schematic diagrams and single-line diagrams is analogous to the distinction between loop diagrams and P&IDs: the former shows a much greater degree of detail than the latter. As with loop diagrams and P&IDs for instrument technicians, schematic and single-line diagrams serve different purposes for professionals analyzing power systems. There are circumstances when the intricate conductor-by-conductor detail of a schematic is necessary, but for quick analysis of operations and faults in large systems it is hard to compete with the elegance of a single-line diagram.

A set of commonly-used single-line diagram symbols appears on the next two pages.





An example of a single-line diagram showing multiple generating stations, substations, transmission lines, and distribution lines appears here. Note the coloring used to illustrate circuit breaker states (green = off and red = on) which is how single-line diagrams typically appear on computer-based SCADA system displays:



It should be abundantly clear from this example that the single-line diagram format greatly simplifies what would otherwise be a cluttered schematic diagram, in illustrating a system containing over a dozen generators and nearly as many loads. As such, single-line diagrams are indispensable for electrical power system operators and other personnel who must make quick decisions in oversight of a power grid.

25.5 Circuit breakers and disconnects

Circuit breakers are the “final control elements” of the electric power industry, akin to control valves in the process industries. They are strictly on/off devices, used to make and break connections under load in power systems. Circuit breakers automatically open when dangerous circuit conditions are detected. Some low-voltage circuit breakers are strictly local-controlled devices, but larger circuit breakers (especially medium- and high-voltage units) may also be operated remotely by electrical signals.

Disconnects are switches designed to isolate sections of a power system in case of damage or to allow for routine maintenance. These may be manually-operated devices, or operated remotely by electric motor, and are typically not intended to make or break load current. *Circuit breakers*, by contrast, are designed to interrupt very high levels of electric current so they may safely cut off power in the event of a short-circuit fault.

This next photograph shows a set of three disconnect switches used on the line (input) side of a three-phase 500 kV power transformer bank in a large substation. As you can see, each disconnect is a simple “knife” switch design, bearing a striking similarity to its electrical schematic symbol. The long metal arm hinges on the left-hand side (at the top of the double insulator stands) and makes or breaks electrical contact on the right-hand side (where the arm is tipped with a sphere):



This particular high-voltage disconnect is motor-actuated, allowing all three disconnects to be operated in unison by remote control. When in the “open” state each metal arm points vertically toward the sky, clearly revealing its status to visual inspection. Lower-voltage disconnects are often built as manually-actuated devices, a crank handle or lever installed at ground level for a human lineman to actuate.

In series with these disconnect switches are a triad of 500 kV circuit breakers (with only two of them appearing in this next photograph). The circuit-breaking elements reside in the horizontal tube portion of each unit (called the “tank”), the tall ribbed structures being insulated conductors bringing the power down to the low-mounted breaker tank. Since disconnect switches are generally not rated for load current interruption, circuit breakers are necessary to “break” the current and safely extinguish the inevitable arc that forms when a live circuit is broken. In the background you can see a hinged disconnect switch in series with the furthest circuit breaker, a set of which serve to isolate power from the three circuit breakers (and every other component “downstream” of the disconnects) to permit maintenance on those components:



All three circuit breakers are remotely controlled by 125 VDC signals energizing solenoid coils within the breaker units. One solenoid coil called the “close” coil causes the breaker mechanism to move to the closed (conducting) position when momentarily energized. Another solenoid coil called the “trip” coil causes the breaker mechanism to move to the open (non-conducting) position when momentarily energized. The breakers also contain status contacts signaling the breaker’s change of state. It is these solenoid coils and status contacts which permit the circuit breaker to be a part of an automatic control system rather than function merely as a manual switching device.

25.5.1 Low-voltage circuit breakers

This photograph shows a typical low-voltage²¹ (480 volt) circuit breakers in a “Motor Control Center” (MCC) panel, for 480 volt 3-phase industrial power circuits:




Note how each circuit breaker has its own on/off handle, for manual operation. These circuit breakers, like most low-voltage breakers, are capable of turning off (“tripping”) on their own when

²¹In the United States, the term “low voltage” with reference to power circuits usually refers to circuits of 600 volt or less potential.

high current is detected, but must be turned on (“closed”) manually. That is to say, they lack “close” and “trip” solenoid coils present in larger circuit breaker units which would permit remote operation.

Some low-voltage circuit breakers utilize thermal elements to detect overcurrent conditions, much like the circuit breakers traditionally found in residential applications. When this thermal element inside the circuit breaker becomes too warm with current, it mechanically forces the mechanism to trip and open the contacts. Other low-voltage circuit breakers are magnetically operated, tripping when the magnetic field caused by conductor current becomes excessive. In either case, the trip mechanism for a low-voltage circuit breaker is typically contained within the circuit breaker itself.

A close-up photograph shows one of these breaker panels, containing two separate three-phase circuit breakers inside:



Note how the “Sump Pump” circuit breaker has been placed in the “off” position, its handle locked there by a padlock, a danger tag attached to notify any personnel of the reason for the breaker’s lock-out.

This next photograph shows a different brand of MCC (manufactured by Gould) where each unit contains not only a circuit breaker, but also an entire motor starter assembly (contactor, overload heaters, and associated switch contacts) for controlling a three-phase electric motor. One of these motor control “buckets” has been removed, revealing the line and load bus connections in the rear:




Industrial circuit breakers such as this are typically designed to be unplugged for ease of maintenance and replacement. If the “bucket” units are heavy, a lifting hoist is provided on the MCC to facilitate their removal and replacement.

25.5.2 Medium-voltage circuit breakers

Circuit breaker design and construction becomes more complicated at higher voltages, such as the voltage range extending from 2.4 kV to 35 kV commonly classified as “medium-voltage” in the electrical power distribution industry. Aside from being physically larger than low-voltage circuit breakers, medium-voltage circuit breakers are generally not self-tripping as low-voltage circuit breakers are. Rather, medium-voltage circuit breakers are electrically commanded to trip (and to close) by external devices called *protective relays* monitoring dangerous electrical conditions. Internally, these circuit breakers are equipped with “trip” and “close” electromagnet solenoids allowing the mechanism to be triggered by remote electrical signals.


Medium-voltage circuit breakers are designed to be unplugged from the breaker panel for maintenance and replacement. This is referred to in the electrical power industry as *racking out* a circuit breaker. Some circuit breakers “rack” by moving horizontally, sliding in and out of the panel on guide rails. Other circuit breakers “rack” by moving vertically, sliding up into and down out of the stationary panel terminals, as shown in the following illustration:



The primary purpose of being able to “rack” a medium voltage breaker into and out of its place in a breaker panel is to facilitate regular maintenance on the circuit breaker mechanism. Unlike the circuit breakers you find in your home, these units may be frequently cycled and will suffer definite wear with each actuation. After a certain number of closing/tripping cycles, the breaker must be removed from service for inspection and testing.

Racking out a circuit breaker also provides another advantage, and that is an extra measure of safety when securing a power circuit in a zero-energy state. When a circuit breaker has been locked into its “racked out” position, the load conductors serviced by this breaker absolutely cannot become energized even if the circuit breaker contacts were made to close. This is analogous to unplugging an electrical appliance from a wall receptacle: it cannot be powered up even if the switch is turned on!

An example of a vertically-racking circuit breaker is the General Electric “Magneblast” unit shown below, designed for use in power systems operating up to 15 kV. The particular unit shown rests on a wooden pallet in a storage area. Normally, it would be installed in a metal-clad breaker panel, its components hidden from direct view:




The six “stabs” seen on the top of this breaker unit engage with six sockets connected to the six bus-bar conductors inside the breaker panel (three for the three phases of the line supply, plus three more for the three phases of the load conductors). When this circuit breaker is “racked out,” it is dropped down so that these six stabs disengage from the bus-bar connections, making it impossible to energize the load conductors even if the breaker contacts close.

A detail not seen in this photograph is the hoisting mechanism necessary to lift this breaker into its “racked-in” position. Medium-voltage circuit breakers such as the General Electric Magneblast are quite heavy, requiring special “lift truck” frames to hoist into and out of their engaged positions in the circuit breaker panel.

Not only does “racking out” a circuit breaker add an extra measure of safety for personnel working on the load circuit, but it also allows the breaker to be tested in-place without energizing the load. The electrical connections commanding the breaker to open (trip) and close may still be connected to the control circuitry even in the racked-out state, permitting such tests. The following illustration shows how such a test may be performed in the “racked-out” condition:

*Breaker "racked out" (removed from service)
but control plug still engaged for operational testing*



At medium-voltage and greater levels of potential, a significant design problem is how to rapidly extinguish the arc formed when contacts separate under load. Low-voltage circuit breakers simply rely on a wide and rapid enough separation of contact points to ensure the electric arc formed when the breaker trips cannot continue more than a fraction of a second. In medium-voltage circuits, both the heat output and the potential length of the electric arc formed by separating contacts is huge and therefore the arc must be extinguished as quickly as possible, both for personnel safety and for extending the operating life of the circuit breaker.


The original General Electric Magneblast circuit breaker design used a series of *arc chutes*, electromagnet coils, and pneumatic jets to direct the arc away from the separating contacts and thereby extinguish it rapidly. Other medium-voltage circuit breaker designs submerge the electrical contacts in an *oil bath* to keep them completely isolated from air so that an arc could never form. This oil has a high dielectric value (i.e. it is an excellent electrical insulator with a high breakdown rating), but needs to be tested on a regular basis to ensure good integrity.

A modern approach to the problem of extinguishing the arc drawn by opening circuit breaker contacts is to encapsulate the contacts inside of an air-tight *vacuum chamber*. This rear view of this GE Magneblast circuit breaker shows it retrofitted with vacuum contacts (the three white-colored components seen inside the breaker frame), replacing the old open-air contacts and arc chutes:



By removing all air from the vicinity of the contacts, there are no gas molecules to ionize when the contacts separate. Not only does this completely eliminate the problem of contact arcing, but it also permits the circuit breaker mechanism to perform its job with a shorter “throw” (less contact motion), since less gap distance is necessary to prevent current in a vacuum than in air. The only real challenge now is ensuring the integrity of the vacuum inside these chambers. This requires periodic testing of the contacts’ dielectric rating by maintenance personnel using high-voltage testing equipment.

An interesting feature of the GE Magneblast and other medium-voltage circuit breakers is the mechanism for actuation. These circuit breaker contacts must be moved swiftly and with significant force in order to ensure quick and repeatable make/break times. In order to achieve this rapidity of motion, the breaker is designed to actuate by the stored energy of large mechanical springs. A side-view of a Magneblast circuit breaker shows a pair of large coil springs used to trip and close the circuit breaker contacts:




Much like the spring on the hammer of a firearm, the springs inside this Magneblast circuit breaker provide the mechanical driving force for opening and closing the breaker's three electrical power contacts. The act of opening or closing this circuit breaker is analogous to pulling the trigger of a firearm: a small mechanical movement unleashes the stored energy of these springs to do the actual work of rapidly opening and closing the contacts.

These springs are tensed (“charged”) by an electric motor in the times following an actuation cycle, so they will be ready for the next actuation. Typically these charging motors are powered by 125 VDC supplied by the substation’s “station power” battery bank, so they may operate even in the event of a total black-out condition where the substation loses AC line power from its incoming transmission lines. Indicator flags on the front of the circuit breaker reveal the breaker’s contact status as well as its spring charge status:



Green-colored flags seen on the front panel of this breaker show the contact status as “open” and the spring status as “discharged”. This circuit breaker is incapable of any action until its spring is charged. Once the spring has been charged, pushing the button labeled “Manual trip” will cause the breaker contacts to open, and pushing the button labeled “Manual close” will cause the breaker contacts to close.

A photograph of the front panel of a Westinghouse vacuum circuit breaker reveals the same basic indicators and manual controls seen on the (older) General Electric circuit breaker:



In this particular example, the actuating spring is charged, which means the breaker is in a state of readiness to switch from its present status (open, or tripped) to its opposite status (closed).

Both of these medium-voltage circuit breakers share another feature of interest: a mechanical counter tracking the number of close/trip cycles the breaker has experienced. The act of making and breaking high-power electric circuits takes a toll on the components of a circuit breaker – especially the contacts – and therefore this count value is a useful parameter for maintenance purposes. The breaker should be serviced at manufacturer-specified intervals of close/trip cycles, just like an automobile should be serviced at manufacturer-specified intervals of distance traveled.

25.5.3 High-voltage circuit breakers

At voltages 46 kV and above (classified as “high voltage” in the electrical power industry), the challenge of extinguishing the electric arc formed by separating breaker contacts becomes severe. Two popular strategies for mitigating contact arc in modern high voltage circuit breakers are *oil immersion* and *gas quenching*.

A set of three oil-bath circuit breakers (OCB's) rated for 230 kV service is shown here, retired from service:



Each of the three circuit breakers (one for each line of the three-phase circuit) is mechanically linked by a common shaft at the top of the breaker tanks, so they all trip and close as one unit.

The fast and reliable actuation of such a bulky mechanism requires a large amount of stored energy, and in the case of the oil circuit breaker shown above the energy storage medium is compressed air. An on-board electric air compressor powered by “station power” maintains air pressure inside a pressure vessel, and this compressed air is directed to a piston actuator through

solenoid valves to provide the actuation force necessary to move the breaker contact assemblies open and closed.

A view inside the enclosure on the far side of this oil circuit breaker reveals the air compressor (upper-right), compressed air storage tank (right) and actuation cylinder (middle):



The man shown in this photograph is pointing to a solenoid valve designed to pass compressed air to and from the piston actuator. A large-diameter black hose runs from this solenoid through the bottom of the enclosure, allowing compressed air from the cylinder to vent to atmosphere.

A more modern breaker design for 230 kV service is this gas-quenched circuit breaker unit, a mere fraction²² of the physical size of the oil circuit breaker shown previously:




The ribbed porcelain structures are the high-voltage terminals for this circuit breaker: three for the three-phase lines coming in, and three for the three-phase load terminals exiting. The actual contact assemblies reside in the gas-filled horizontal metal tubes (“tanks”). It is striking to note that the same current interruption and isolation functions performed by the gigantic oil-filled tanks of the retired circuit breaker previously shown are performed by this relatively tiny gas-quenched breaker.

The gas inside the breaker’s tanks is *sulfur hexafluoride*, a very dense gas (about 5 times denser than air) with excellent electrical insulating and arc-extinguishing properties. SF_6 gas is contained in these breaker contact chambers under pressure to maximize its dielectric breakdown strength (its ability to withstand high voltage without ionizing and passing current across the gap between the circuit breaker’s open contacts). SF_6 gas is non-toxic²³ and safe to handle.

²²For an equitable size comparison between the two different types of circuit breaker, consider the fact that the insulators on this gas-quenched circuit breaker are approximately the same physical height as the insulators on the previously-shown oil-tank circuit breaker.

²³While pure SF_6 gas is benign, it should be noted that one of the potential chemical byproducts of arcing in an SF_6 -quenched circuit breaker is hydrofluoric acid (HF) which is extremely toxic. HF is formed when SF_6 gas arcs in the presence of water vapor (H_2O), the latter being nearly impossible to completely eliminate from the interior chambers of the circuit breaker. This means any maintenance work on an SF_6 -quenched circuit breaker must take this chemical hazard into consideration.


Like the large oil-filled circuit breakers seen previously, this SF₆ circuit breaker has an enclosure on one side where the actuation and control components are located. Inside this enclosure we see a large stack of Belleville spring washers (the dark-colored discs located in the center of the enclosure), which are used as the mechanical energy-storage medium instead of compressed air. This stack of spring-steel washers is compressed by an electric motor and gear mechanism, then the spring tension is released through another mechanism to close and trip the breaker's contacts on demand. As usual this charging motor receives its power from the uninterruptible "station power" supply of the substation, allowing the breaker to actuate even in the event of a total "blackout" condition:



Inside this enclosure we also see a small pair of pushbuttons (one red, one green) just below and to the right of the Belleville washer stack for manually closing and tripping the breaker, respectively. The common color coding used in the United States for electric power switchgear is *red for energized, green for de-energized*. This may seem backward to most people familiar with red and green traffic lights, where red means "stop" and green means "go," but the concept here is one of safety: red means "dangerous" (power on) while green means "safe" (power off).

It should be noted that most of the time these high-voltage circuit breakers are triggered remotely, rather than manually by someone standing near them. These command signals may come from a manual switch located in the control room, or from some automatic circuit such as a *protective relay* instructing the breaker to open due to an abnormal system condition.

The nameplate photographed on a similar SF₆ circuit breaker reveals some interesting features:

		TYPE:	DT1-123 F1	CUSTOMER P.O.:	00055131
		SERIAL NUMBER:	11385-DT12	ORDER NUMBER:	C03303
		MANUFACTURE DATE:	2 / 2006	INSTRUCTION BOOK No.:	DT12-IM
RATED MAXIMUM VOLTAGE	123 kV	RATED INTERRUPTING TIME	3.0 CYCLES	PARTS LIST No.:	DT12A0748
RATED VOLTAGE RANGE FACTOR (K)	1.0	RATED SHORT CIRCUIT CURRENT	40 kA	WIRING DIAGRAM:	DT12C0748
RATED FREQUENCY	60 Hz	-TERMINAL FAULT	40 kA	SF ₆ PRESSURE (@ 20° C / 68° F)	RELATIVE PRESSURE
DUTY CYCLE	0-0.3e-C0-15e-C0	-90% SLF WITH 0mFL_g CAPACITANCE (WITHIN 100m)	40 kA	-NORMAL OPERATING PRESSURE	98.6 psig
RATED CONTINUOUS CURRENT	2000 A	TEMPERATURE RANGE	-30° C to +40° C	-LOW PRESSURE ALARM	85 psig
CAPACITANCE CURRENT SWITCHING		SHORT TIME CURRENT DURATION	3e	-MINIMUM PRESSURE	80 psig
-OVERVOLTAGE FACTOR	2.0	% DC COMPONENT		WEIGHT OF GAS	86 lb
-LINE CHARGING	160 A	FULL WAVE IMPULSE WITHSTAND	550 kV	TOTAL WEIGHT	6577 lb
-ISOLATED BANK	315 A	SWITCHING IMPULSE WITHSTAND		MECHANISM TYPE	CRRS-4E
-BACK-TO-BACK	315 A	-TERMINAL-TO-GROUND	N/A	CONTROL VOLTAGE	SPRING
-INRUSH PEAK	16 kA	-TERMINAL-TO-TERMINAL	N/A	-CLOSING	125VDC, 3.1A
-INRUSH FREQUENCY	4250 Hz			-TRIPPING (EACH)	125VDC, 1.8A
OUT-OF-PHASE SWITCHING	10 kA			-MOTOR	125VDC, 850W
				-SPACE HEATER	120VAC, 190W
AREVA T&D INC., HIGH VOLTAGE SWITCHGEAR ONE POWER LANE - CHARLEROI, PA 15022					
MADE IN THE USA					

According to this nameplate, the normal operating pressure for the SF₆ gas is 98.6 PSI. A low-pressure alarm triggers if the SF₆ gas pressure happens to fall below 85 PSI. When opening (tripping), the circuit breaker only takes *3 cycles*' worth of time at 60 Hz to completely interrupt the current. One solenoid coil closes the breaker, and that coil requires a signal of 125 volts DC at just over 3 amps of current. The breaker's contacts may be tripped by energizing one or more

redundant solenoid coils at 125 VDC and 1.8 amps (each). In either direction, the breaker's actuation is powered by a pre-charged spring, much like the 230 kV breaker seen previously. This particular breaker is rated for 123 kV at 2000 amps full-load.

Sulfur hexafluoride gas circuit breaker technology is popular for higher voltage applications as well, such as these 500 kV circuit breakers seen here:



In this application, where three separate circuit breaker units independently interrupt current for the three-phase power lines, there is no mechanical link to synchronize the motion of the three contact sets. Instead, each single-phase circuit breaker actuates independently.

So far all of the high-voltage circuit breakers shown in previous photographs are of the *dead tank* type, where the structure housing the interrupting contact(s) is maintained at ground potential (i.e. the outside surface of the circuit breaker mechanism is electrically “dead”). Some high-voltage circuit breakers are built such that their interrupting assemblies are at line potential, the entire breaker suspended above ground from insulators. This type of circuit breaker is called a *live tank*, because the “tank” containing the contact(s) operates at a high voltage with respect to earth ground. A photograph of a 500 kV, single-pole²⁴, SF₆-quenched, live tank breaker appears next:



The actuating mechanism for this live-tank breaker is housed in the “can” assembly seen at the base, where the vertical insulator meets the steel supporting tower.

²⁴This particular circuit breaker, like most live-tank circuit breakers, interrupts just one phase (i.e. one “pole”) of a three-phase bus. Portions of the second and third live-tank SF₆ breakers comprising the full three-phase breaker array for this bus may be seen near the left-hand edge of the photograph.

Another 500 kV, single-pole, live-tank circuit breaker assembly appears in the next photograph, this particular breaker being an older unit using *compressed air* as the interrupting medium rather than sulfur hexafluoride gas:



Air nozzles powered by hundreds of PSI of compressed air are used to “blow out” the arc formed when the circuit breaker’s contacts separate. These air nozzles are not visible in the photograph, being internal to the circuit breaker’s construction. An interesting feature of this style of circuit breaker is the loud report generated when it trips: the sound of the compressed air jets extinguishing the arc across the separating contact poles within the breaker is not unlike that of a firearm discharging.


Multiple series-connected contact assemblies comprise this circuit breaker, distributing the energy of the arc across multiple points in the breaker assembly rather than across a single contact. This is evident in the photograph as multiple clusters of “tanks” on the top of the left-hand assembly, as well as the second live-tank assembly connected in series to the right. Such arrangements are necessary because air is a less effective medium for extinguishing an electric arc than either oil or SF₆ gas.

25.5.4 Reclosers

A special type of medium-voltage circuit breaker used to quickly interrupt and re-establish power in distribution lines is called a *recloser*. Reclosers are designed to trip if ever a distribution line suffers a “transient” (momentary) fault due to some natural event such as a lightning strike causing an insulator to “flash over” to ground or a tree branch touching one or more line conductors, then automatically re-close moments later to test whether or not the fault still persists. If the fault clears on its own – a common occurrence with tree branches, as the branch may break off or burn away following the initial arc – then the recloser remains closed and continues providing power to customers. By some estimations transient faults account for 70% to 90% of all faults occurring on overhead power lines. If non-reclosing fault protection were applied to all distribution and transmission lines, extended interruptions of electric power service would be far more common than they are now.

If you have ever experienced momentary cessations of electrical power service to your home or business where the power “blinks” off and on in rapid succession, you have experienced a recloser at work. The recloser opens as soon as an overcurrent condition is detected, then recloses briefly to “test” for continued presence of the fault. If the fault persists, the recloser trips again and then recloses once more to “re-test” for the fault. If the fault has cleared by then, the recloser remains closed and restores normal power service to customers. Only if the fault persists after multiple “shots” does the recloser remain in the tripped (open) state and waits for line crews to repair the fault.

Reclosers are typically located some distance “downstream” of the substation, to isolate certain remote portions of the distribution network. Circuit breakers at the substation provide protection for each distribution line as a whole:



A typical recloser resembles a “circuit breaker on a stick,” located near the top of a distribution power pole near the line conductors. Modern reclosers use SF₆ gas quenching to extinguish arcing resulting from the interruption of high-magnitude fault currents. Legacy reclosers typically employed oil quenching. This photograph shows a modern (SF₆-quenched) recloser, with its three phase current-interrupting contact units clearly visible:



The grey enclosure located near ground level on this pole contains the *protective relay* responsible for issuing the “trip” and “close” signals to the recloser’s coils. Current transformers located within the recloser provide isolated sensing of line current to the reclosing relay, necessary for detecting any overcurrent conditions that might result from a transient fault. Each attempt by the reclosing relay to re-close the breaker is called a *shot*. If the reclosing relay fails to clear the fault by a certain number of shots, it enters the *lockout* state whereby the recloser remains open and must be re-closed by human intervention.

Distribution line reclosers, unlike circuit breakers located in substations, cannot rely on an auxiliary “station power” energy source for opening and closing its line-interrupting contacts. Therefore these small units utilize AC line voltage as the actuating power for the contacts. Low-voltage “trip” and “close” circuits still exist for control purposes, but the actual energy source for rapid tripping/reclosing cycles comes from the AC line itself.

The principle of automatic reclosing may be applied to transmission lines as well as distribution

lines, but new challenges exist at this level of a power grid. When transmission lines serve to interconnect distributed generating stations, interruption of that line for any significant time invites generator de-synchronization. Recall that AC generators, once synchronized with each other and connected in parallel on a common grid circuit, tend to remain synchronized with each other as though their mechanical shafts had become coupled. If a circuit breaker opens along a transmission line system and de-couples generators from each other, those generators will be free to fall out of synchronization. Reclosing that circuit breaker when those generators are out of sync with each other can be disastrous. Automatic reclosing at the distribution line level of a power grid, therefore, must either be fast enough that generators will not have enough time to fall out of sync with each other, or blocked by other protective relay logic to prevent reclosure in an out-of-sync situation.

25.6 Electrical sensors

The two “process variables” we rely on most heavily in the field of electrical measurement and control are *voltage* and *current*. From these primary variables we may determine impedance, reactance, resistance, as well as the reciprocals of those quantities (admittance, susceptance, and conductance).


Other sensors more common to general process measurement such as temperature, pressure, level, and flow are also used in electric power systems, but their coverage in other chapters of this book is sufficient to avoid repetition in this chapter.

Two common types of electrical sensors used in the power industry are *potential transformers* (PTs) and *current transformers* (CT). These are precision-ratio electromagnetic transformers used to step high voltages and high currents down to more reasonable levels for the benefit of panel-mounted instruments to receive, display, and/or process.

25.6.1 Potential transformers

Electrical power systems typically operate at dangerously high voltage. It would be both impractical and unsafe to connect panel-mounted instruments directly to the conductors of a power system if the voltage of that power system exceeds several hundred volts. For this reason, we must use a special type of step-down transformer referred to as a *potential transformer* to reduce and isolate the high line voltage of a power system to levels safe for panel-mounted instruments to input.

Shown here is a simple diagram illustrating how the high phase and line voltages of a three-phase AC power system may be sensed by low-voltage voltmeters through the use of step-down potential transformers:



Potential transformers are commonly referred to as “PT” units in the electrical power industry. It should be noted that the term “voltage transformer” and its associated abbreviation VT is becoming popular as a replacement for “potential transformer” and PT.

When driving a voltmeter – which is essentially an open-circuit (very high resistance) – the PT behaves as a voltage source to the receiving instrument, sending a voltage signal to that instrument proportionately representing the power system’s voltage.

The following photograph shows a potential transformer sensing the phase-to-ground voltage on a three-phase power distribution system. The normal phase voltage in this system is 7.2 kV (12.5 kV three-phase line voltage), and the PT's normal secondary voltage is 120 volts, necessitating a ratio of 60:1 (as shown on the transformer's side):



Any voltage output by this PT will be $\frac{1}{60}$ of the actual phase voltage, allowing panel-mounted instruments to read a precisely scaled proportion of the 7.2 kV (typical) phase voltage safely and effectively. A panel-mounted voltmeter, for example, would have a scale registering 7200 volts when its actual input terminal voltage was only 120 volts. This is analogous to a 4-20 mA indicating meter bearing a scale labeled in units of “PSI” or “Degrees Celsius” because the 4-20 mA analog signal merely represents some other physical variable sensed by a process transmitter. Here, the physical variable being sensed by the potential transformer is still voltage, just at a 60:1 ratio greater than what the panel-mounted instrument receives. Like the 4-20 mA DC analog signal standard so common in the process industries, 115 or 120 volts is the standard potential transformer output voltage used in the electrical industry to represent normal power system voltage.

This next photograph shows a set of three PTs used to measure voltage on a 13.8 kV substation bus. Note how each of these PTs is equipped with *two* high-voltage insulated terminals to facilitate phase-to-phase (line voltage) measurements as well as phase-to-ground:




Another photograph of potential transformers appears here, showing three large PTs used to precisely step the phase-to-ground voltages for each phase of a 230 kV system (230 kV line voltage, 133 kV phase voltage) all the way down to 120 volts for the panel-mounted instruments to monitor:



A loose-hanging wire joins one side of each PT's primary winding to the respective phase conductor of the 230 kV bus. The other terminal of each PT's primary winding connects to a common neutral point, forming a Wye-connected PT transformer array. The secondary terminals of these PTs connect to two-wire shielded cables conveying the 120 volt signals back to the control room where they terminate at various instruments. These shielded cables run through underground conduit for protection from weather.

Just as with the previous PT, the standard output voltage of these large PTs is 120 volts, equating to a transformer turns ratio of about 1100:1. This standardized output voltage of 120 volts allows PTs of any manufacture to be used with receiving instruments of any manufacture, just as the 4-20 mA standard for analog industrial instruments allows "interoperability" between different manufacturers' brands and models.


A special form of instrument transformer used on very high-voltage systems is the *capacitively-coupled voltage transformer*, or CCVT. These sensing devices employ a series-connected set of capacitors dividing the power line voltage down to a lesser quantity before it gets stepped down further by an electromagnetic transformer. A simplified diagram of a CCVT appears here, along with a photograph of three CCVTs located in a substation:



25.6.2 Current transformers


For the same reasons necessitating the use of potential (voltage) instrument transformers, we also see the use of *current transformers* to reduce high current values and isolate high voltage values between the electrical power system conductors and panel-mounted instruments.

Shown here is a simple diagram illustrating how the line current of a three-phase AC power system may be sensed by a low-current ammeter through the use of a current transformer:




When driving an ammeter – which is essentially a short-circuit (very low resistance) – the CT behaves as a current source to the receiving instrument, sending a current signal to that instrument proportionately representing the power system's line current.

In typical practice a CT consists of an iron toroid²⁵ functioning as the transformer core. This type of CT does not have a primary “winding” in the conventional sense of the word, but rather uses the line conductor itself as the primary winding. The line conductor passing once through the center of the toroid functions as a primary transformer winding with exactly 1 “turn”. The secondary winding consists of multiple turns of wire wrapped around the toroidal magnetic core:



A view of a current transformer’s construction shows the wrapping of the secondary turns around the toroidal magnetic core in such a way that the secondary conductor remains parallel to the primary (power) conductor for good magnetic coupling:



With the power conductor serving as a single-turn²⁶ winding, the multiple turns of secondary

²⁵A “toroid” is shaped like a donut: a circular object with a hole through the center.

²⁶This raises an interesting possibility: if the power conductor were to be wrapped around the toroidal core of the CT so that it passes through the center *twice* instead of *once*, the current step-down ratio will be cut in half. For example, a 100:5 CT with the power conductor wrapped around so it passes through the center twice will exhibit an actual current ratio of only 50:5. If wrapped so that it passed through the CT’s center three times, the ratio would be reduced to 33.33:5. This useful “trick” may be used in applications where a lesser CT ratio cannot be found, and one must make do with whatever CT happens to be available. If you choose to do this, however, beware that the current-measuring capacity of the CT will be correspondingly reduced. Each extra turn of the power conductor

wire around the toroidal core of a CT makes it function as a step-up transformer with regard to voltage, and as a *step-down* transformer with regard to current. The turns ratio of a CT is typically specified as a ratio of full line conductor current to 5 amps, which is a standard output current for power CTs. Therefore, a 100:5 ratio CT outputs 5 amps when the power conductor carries 100 amps.

The turns ratio of a current transformer suggests a danger worthy of note: if the secondary winding of an energized CT is ever open-circuited, it may develop an extremely high voltage as it attempts to force current through the air gap of that open circuit. An energized CT secondary winding acts like a current source, and like all current sources it will develop as great a potential (voltage) as it can when presented with an open circuit. Given the high voltage capability of the power system being monitored by the CT, and the CT turns ratio with more turns in the secondary than in the primary, the ability for a CT to function as a *voltage step-up* transformer poses a significant hazard.

Like any other current source, there is no harm in short-circuiting the output of a CT. Only an open circuit poses risk of damage. For this reason, CT circuits are often equipped with *shorting bars* and/or *shorting switches* to allow technicians to place a short-circuit across the CT secondary winding before disconnecting any other wires in the circuit. Later subsections will elaborate on this topic in greater detail.

adds to the magnetic flux experienced by the CT's core for any given amount of line current, making it possible to magnetically saturate the core if the line current exceeds the reduced value (e.g. 50 amps for the home-made 50:5 CT where the line passes twice through the center of a 100:5 CT).

Current transformers are manufactured in a wide range of sizes, to accommodate different applications. Here is a photograph of a current transformer showing the “nameplate” label with all relevant specifications. This nameplate specifies the current ratio as “100/5” which means this CT will output 5 amps of current when there is 100 amps flowing through a power conductor passed through the center of the toroid:



The black and white wire pair exiting this CT carries the 0 to 5 amp AC current signal to any monitoring instrument scaled to that range. That instrument will see $\frac{1}{20}$ (i.e. $\frac{5}{100}$) of the current flowing through the power conductor.


The following photographs contrast two different styles of current transformer, one with a “window” through which any conductor may be passed, and another with a dedicated busbar fixed through the center to which conductors attach at either end. Both styles are commonly found in the electrical power industry, and they operate identically:



Here is a photograph of some much larger CTs intended for installation inside the “bushings²⁷” of a large circuit breaker, stored on a wooden pallet:



The installed CTs appear as cylindrical bulges at the base of each insulator on the high-voltage circuit breaker. This particular photograph shows flexible conduit running to each bushing CT, carrying the low-current CT secondary signals to a terminal strip inside a panel on the right-hand end of the breaker:




Signals from the bushing CTs on a circuit breaker may be connected to *protective relay* devices to trip the breaker in the event of any abnormal condition. If unused, a CT’s secondary terminals are simply short-circuited at the panel.

²⁷High-voltage devices situate their connection terminals at the ends of long insulators, to provide a large air gap between the conductors and the grounded metal chassis of the device. The point at which the long insulator (with a conductor inside of it) penetrates the housing of the device is called the *bushing*.

Shown here is a set of three very large CTs, intended for installation at the bushings of a high-voltage power transformer. Each one has a current step-down ratio of 600-to-5:



In this next photograph we see a tiny CT designed for low current measurements, clipped over a wire carrying only a few amps of current. This particular current transformer is constructed in such a way that it may be clipped around an existing wire for temporary test purposes, rather than being a solid toroid where the conductor must be threaded through it for a more permanent installation:



This CT's ratio of 3000:1 would step down a 5 amp AC signal to 1.667 milliamps AC.

This last photograph shows a current transformer used to measure line current in a 500 kV substation switchyard. The actual CT coil is located inside the red-colored housing at the top of the insulator, where the power conductor passes through. The tall insulator stack provides necessary separation between the conductor and the earth below to prevent high voltage from “jumping” to ground through the air:




25.6.3 Transformer polarity


An important characteristic to identify for transformers in power systems – both power transformers and instrument transformers – is *polarity*. At first it may seem inappropriate to speak of “polarity” when we know we are dealing with *alternating* voltages and currents, but what is really meant by this word is *phasing*. When multiple power transformers are interconnected in order to share load, or to form a three-phase transformer array from three single-phase transformer units, it is critical that the phase relationships between the transformer windings be known and clearly marked. Also, we need to know the phase relationship between the primary and secondary windings (coils) of an instrument transformer in order to properly connect it to a receiving instrument such as a protective relay. For some instruments such as simple indicating meters, polarity (phasing) is unimportant. For other instruments comparing the phase relationships of two or more received signals from instrument transformers, proper polarity (phasing) is critical.

Polarity markings for any transformer may be symbolized several different ways:

Transformer polarity symbols




The marks should be interpreted in terms of *voltage polarity*, not current. To illustrate using a “test circuit²⁸” feeding a momentary pulse of DC to a transformer from a small battery:




Note how the secondary winding of the transformer develops the same polarity of voltage drop as is impressed across the primary winding by the DC pulse: for both the primary and secondary windings, the sides with the dots share the same positive potential.

²⁸The battery-and-switch test circuit shown here is not just hypothetical, but may actually be used to test the polarity of an unmarked transformer. Simply connect a DC voltmeter to the secondary winding while pressing and releasing the pushbutton switch: the voltmeter’s polarity indicated while the button is pressed will indicate the relative phasing of the two windings. Note that the voltmeter’s polarity will reverse when the pushbutton switch is released and the magnetic field collapses in the transformer coil, so be sure to pay attention to the voltmeter’s indication *only* during the time of switch closure! This is an application where an *analog* voltmeter may actually be superior to a digital voltmeter, since the instantaneous movement of a mechanical needle (pointer) is easier to visually interpret than the sign of a digital number display.

If the battery were reversed and the test performed again, the side of each transformer winding with the dot would be negative:



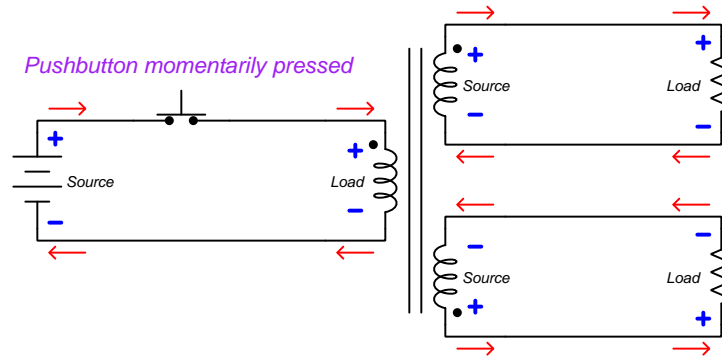
If we reverse the secondary winding's connection to the resistor and re-draw all voltages and currents, we see that the polarity dot always represents common voltage potential, regardless of source polarity:



It should be noted that this battery-and-switch method of testing should employ a fairly low-voltage battery in order to avoid leaving residual magnetism in the transformer's core²⁹. A single 9-volt dry-cell battery works well given a sensitive meter.



²⁹The amount of magnetic force H applied to the transformer's core is a direct function of winding current. If the DC test source is capable of pushing significant amounts of current through the transformer, it may leave the core in a partially magnetized state which will then affect its performance when powered by AC. A relatively "weak" source such as a 9 volt "transistor" battery helps ensure this will not happen as a result of the polarity test.

Transformers with multiple secondary windings act the same, with each secondary winding's polarity mark having the same polarity as every other winding:



To emphasize this important point again: *transformer polarity dots always refer to voltage, never current*. The polarity of voltage across a transformer winding will always match the polarity of every other winding on that same transformer in relation to the dots. The direction of current through a transformer winding, however, depends on whether the winding in question is functioning as a *source* or a *load*. This is why currents are seen to be in opposite directions (into the dot, out of the dot) from primary to secondary in all the previous examples shown while the voltage polarities all match the dots. A transformer's primary winding functions as a *load* (conventional-flow current drawn flowing into the positive terminal) while its secondary winding functions as a *source* (conventional-flow current flowing out of the positive terminal).


Transformer polarity is very important in the electric power industry, and so terms have been coined for different polarity orientations of transformer windings. If polarity dots for primary and secondary windings lie on the same physical side of the transformer it means the primary and secondary windings are wrapped the same direction around the core, and this is called a *subtractive* transformer. If polarity dots lie on opposite directions of the transformer it means the primary and secondary windings are wrapped in opposite directions, and this is called an *additive* transformer. The terms “additive” and “subtractive” have more meaning when we view the effects of each configuration in a grounded AC power system. The following examples show how voltages may either add or subtract depending on the phase relationships of primary and secondary transformer windings:




Transformers operating at high voltages are typically designed with subtractive winding³⁰ orientations, simply to minimize the dielectric stress placing on winding insulation from inter-winding voltages. Instrument transformers (PTs and CTs) by convention are *always* subtractive.

³⁰The IEEE standard C57.12.00-2010 (“IEEE Standard for General Requirements for Liquid-Immersed Distribution, Power, and Regulating Transformers”) states that single-phase transformers having power ratings of 200 kVA and below *and* high-voltage winding ratings of 8.66 kV and below must have *additive* polarity, and that all other types of power transformers must have *subtractive* polarity.

When three single-phase transformers are interconnected to form a three-phase transformer bank, the winding polarities *must* be properly oriented. Windings in a delta network must be connected such that the polarity marks of no two windings are common to each other. Curved arrows are drawn next to each winding to emphasize the phase relationships:




Windings in a wye network must be connected such that the polarity marks all face the same direction with respect to the center of the wye (typically, the polarity marks are all facing away from the center):



Failure to heed these phase relationships in a power transformer bank may result in catastrophic failure as soon as the transformers are energized!

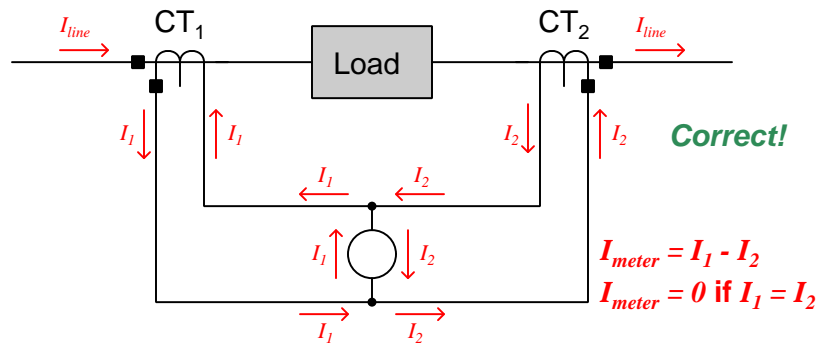
The following photograph shows the diagram for a large utility power transformer³¹ equipped with a number of current transformers permanently installed in the bushings (the points at which power conductors penetrate the steel casing of the power transformer unit). Note the solid black squares marking one side of each CT secondary winding as well as one side of each primary and secondary winding in this three-phase power transformer. Comparing placement of these black squares we can tell all CTs as well as the power transformer itself are wound as *subtractive* devices:



³¹This particular transformer happens to be a *tap-changing* unit, designed to provide a number of ratio increments useful for adjusting voltages in a power distribution system. Its typical primary voltage is 115 kV and its typical secondary voltage is 12.5 kV. If the secondary voltage happens to sag due to a heavy-load conditions, the transformer's tap setting may be manually adjusted to output a slightly greater secondary voltage (i.e. a lesser step-down ratio). This is how electric power distribution utilities manage to keep voltages to customers relatively stable despite ongoing changes in load conditions.

An example of the importance of polarity marks to the connection of instrument transformers may be seen here, where a pair of current transformers with equal turns ratios are connected in parallel to drive a common instrument which is supposed to measure the *difference* in current entering and exiting a load:

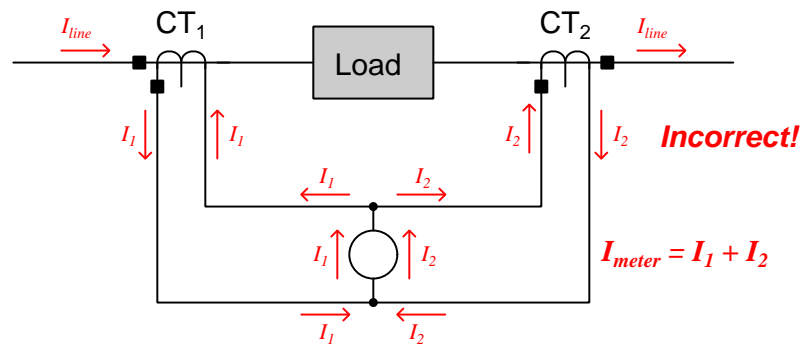
Differential current measurement using two CTs



Properly connected as shown above, the meter in the center of the circuit registers only the *difference* in current output by the two current transformers. If current into the load is precisely equal to current out of the load (which it should be), and the two CTs are precisely matched in their turns ratio, the meter will receive zero net current. If, however, a ground fault develops within the load causing more current to enter than to exit it, the imbalance in CT currents will be registered by the meter and thus indicate a fault condition in the load.

Let us suppose, though, that a technician mistakenly connected one of these CT units backwards. If we examine the resulting circuit, we see that the meter now senses the *sum* of the line currents rather than the *difference* as it should:


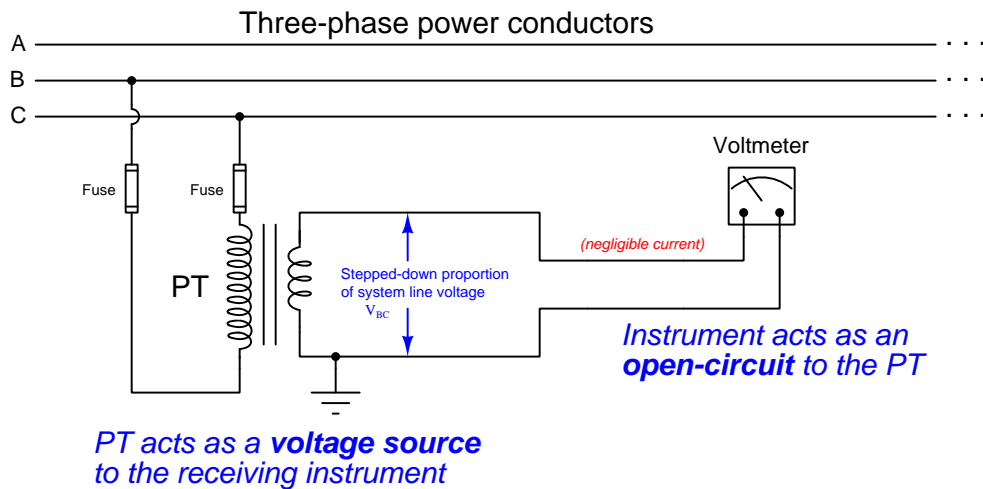
Differential current measurement using two CTs



This will cause the meter to falsely indicate a current imbalance in the load when none exists.

25.6.4 Instrument transformer safety

Potential transformers (PTs or VTs) tend to behave as *voltage sources* to the voltage-sensing instruments they drive: the signal output by a PT is supposed to be a proportional representation of the power system's voltage. Conversely, current transformers (CTs) tend to behave as *current sources* to the current-sensing instruments they drive: the signal output by a CT is supposed to be a proportional representation of the power system's current. The following schematic diagrams show how PTs and CTs should behave when sourcing their respective instruments:



In keeping with this principle of PTs as voltage sources and CTs as current sources, *a PT's secondary winding should never be short-circuited and a CT's secondary winding should never be open-circuited!* Short-circuiting a PT's secondary winding may result in a dangerous amount of current developing in the circuit because the PT will attempt to maintain a substantial voltage across a very low resistance. Open-circuiting a CT's secondary winding may result in a dangerous amount of voltage³² developing between the secondary terminals because the CT will attempt to drive a substantial current through a very high resistance.

This is why you will never see fuses in the secondary circuit of a current transformer. Such a fuse, when blown open, would pose a greater hazard to life and property than a closed circuit with any amount of current the CT could muster.

While the recommendation to never short-circuit the output of a PT makes perfect sense to any student of electricity or electronics who has been drilled never to short-circuit a battery or a laboratory power supply, the recommendation to never *open-circuit* a powered CT often requires some explanation. Since CTs transform current, their output current value is naturally limited to a fixed ratio of the power conductor's line current. That is to say, short-circuiting the secondary winding of a CT will *not* result in more current output by that CT than what it would output to any normal current-sensing instrument! In fact, a CT encounters minimum "burden" when powering a short-circuit because it doesn't have to output any substantial voltage to maintain that amount of secondary current. It is only when a CT is forced to output current through a substantial impedance that it must "work hard" (i.e. output more power) by generating a substantial secondary voltage along with a secondary current.

The latent danger of a CT is underscored by an examination of its primary-to-secondary turns ratio. A single conductor passed through the aperture of a current transformer acts as a winding with one turn, while the multiple turns of wire wrapped around the toroidal core of a current transformer provides the ratio necessary to step down current from the power line to the receiving instrument. However, as every student of transformers knows, while a secondary winding possessing more turns of wire than the primary *steps current down*, that same transformer conversely will *step voltage up*. This means an open-circuited CT behaves as a voltage step-up transformer. Given the fact that the power line being measured usually has a dangerously high voltage to begin with, the prospect of an instrument transformer stepping that voltage up even higher is sobering indeed. In fact, the only way to ensure a CT will not output high voltage when powered by line current is to keep its secondary winding loaded with a low impedance.

It is also imperative that all instrument transformer secondary windings be solidly *grounded* to prevent dangerously high voltages from developing at the instrument terminals via capacitive coupling with the power conductors. Grounding should be done at only one point in each instrument transformer circuit to prevent *ground loops* from forming and potentially causing measurement errors. The preferable location of this grounding is at the first point of use, i.e. the instrument or panel-mounted terminal block where the instrument transformer's secondary wires land. If any test switches exist between the instrument transformer and the receiving instrument, the ground connection must be made in such a way that opening the test switch does not leave the transformer's secondary winding floating (ungrounded).

³²The hazards of an open-circuited CT can be spectacular. I have spoken with power electricians who have personally witnessed huge arcs develop across the opened terminals in a CT circuit! This safety tip is not one to be lightly regarded.

25.6.5 Instrument transformer test switches


Connections made between instrument transformers and receiving instruments such as panel-mounted meters and relays must be occasionally broken in order to perform tests and other maintenance functions. An accessory often seen in power instrument panels is a *test switch bank*, consisting of a series of knife switches. A photograph of a test switch bank manufactured by ABB is seen here:



Some of these knife switches serve to disconnect potential transformers (PTs) from receiving instruments mounted on this relay panel, while other knife switches in the same bank serve to disconnect current transformers (CTs) from receiving instruments mounted on the same panel.

For added security, covers may be installed on the switch bank to prevent accidental operation or electrical contact. Some test switch covers are even lock-able by padlock, for an added measure of access prevention.

Test switches used to disconnect potential transformers (PTs) from voltage-sensing instruments are nothing more than simple single-pole, single-throw (SPST) knife switches, as shown in this diagram:




There is no danger in open-circuiting a potential transformer circuit, and so nothing special is needed to disconnect a PT from a receiving instrument.

A series of photographs showing the operation of one of these knife switches appears here, from closed (in-service) on the left to open (disconnected) on the right:



Test switches used to disconnect current transformers (CTs) from current-sensing instruments, however, must be specially designed to avoid opening the CT circuit when disconnecting, due to the high-voltage danger posed by open-circuited CT secondary windings. Thus, CT test switches are designed to place a *short-circuit* across the CT's output before opening the connection to the current-measuring device. This is done through the use of a special *make-before-break* knife switch:




A series of photographs showing the operation of a make-before-break knife switch appears here, from closed (in-service) on the left to shorted (disconnected) on the right:




The shorting action takes place at a spring-steel leaf contacting the moving knife blade at a cam cut near the hinge. Note how the leaf is contacting the cam of the knife in the right-hand and middle photographs, but not in the left-hand photograph. This metal leaf joins with the base of the knife switch adjacent to the right (the other pole of the CT circuit), forming the short-circuit between CT terminals necessary to prevent arcing when the knife switch opens the circuit to the receiving instrument.

A step-by-step sequence of illustrations shows how this shorting spring works to prevent the CT circuit from opening when the first switch is opened:




It is typical that the non-shorting switch in a CT test switch pair be equipped with a “test jack” allowing the insertion of an additional ammeter in the circuit for measurement of the CT’s signal. This test jack consists of a pair of spring-steel leafs contacting each other in the middle of the knife switch’s span. When that knife switch is in the open position, the metal leafs continue to provide continuity past the open knife switch. However, when a special ammeter adapter plug is forced between the leafs, spreading them apart, the circuit breaks and the current must flow through the two prongs of the test plug (and to the test ammeter connected to that plug).

A step-by-step sequence of illustrations shows how a test jack maintains continuity across an opened knife switch, and then allows the insertion of a test probe and ammeter, without ever breaking the CT circuit:




When using a CT test probe like this, one must be sure to thoroughly test the electrical continuity of the ammeter and test leads before inserting the probe into the test jacks. If there happens to be an “open” fault anywhere in the ammeter/lead circuit, a dangerous arc will develop at the point of that “open” the moment the test probe forces the metal leafs of the test jack apart! Always remember that a live CT is dangerous when open-circuited, and so your personal safety depends on always maintaining electrical continuity in the CT circuit.

This close-up photograph shows a closed CT test switch equipped with a test jack, the jack's spring leafs visible as a pair of "hoop" shaped structures flanking the blade of the middle knife switch:

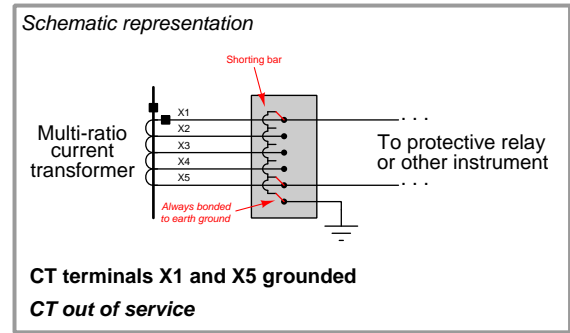
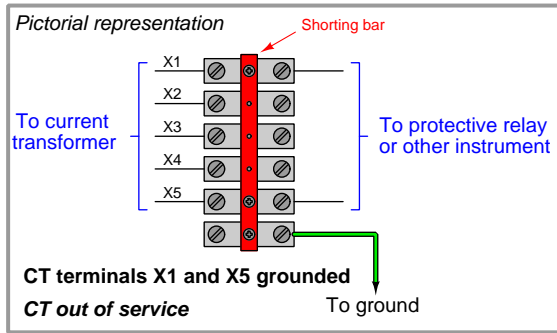


In addition to (or sometimes in lieu of) test switches, current transformer secondary wiring often passes through special “shorting” terminal blocks. These special terminal blocks have a metal “shorting bar” running down their center, through which screws may be inserted to engage with wired terminals below. Any terminals made common to this metal bar will necessarily be equipotential to each other. One screw is always inserted into the bar tapping into the earth ground terminal on the terminal block, thus making the entire bar grounded. Additional screws inserted into this bar force CT secondary wires to ground potential. A photo of such a shorting terminal block is shown here, with five conductors from a multi-ratio (multi-tap) current transformer labeled 7X1 through 7X5 connecting to the terminal block from below:

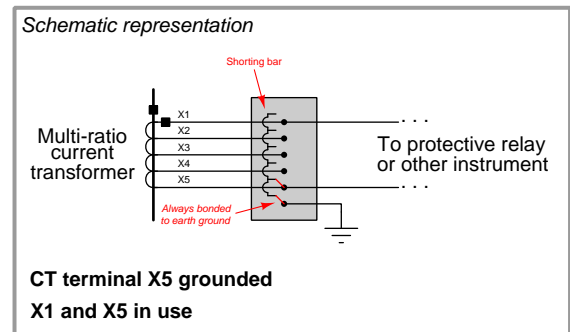



This shorting terminal block has three screws inserted into the shorting bar: one bonding the bar to the ground (“G”) terminal on the far-left side, another one connecting to the “7X5” CT wire, and the last one connecting to the “7X1” CT wire. While the first screw establishes earth ground potential along the shorting bar, the next two screws form a short circuit between the outer two conductors of the multi-ratio current transformer. Note the green “jumper” wires attached to the top side of this terminal block shorting 7X1 to 7X5 to ground, as an additional measure of safety for this particular CT which is currently unused and not connected to any measuring instrument.

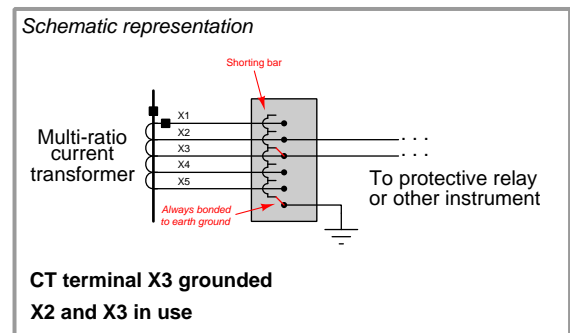
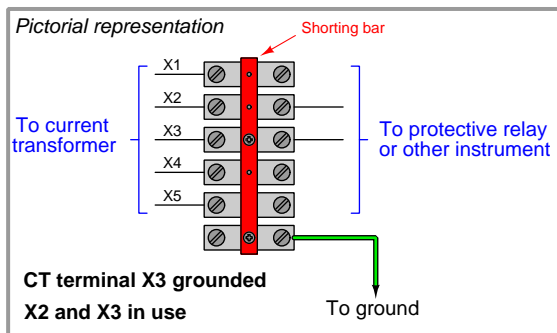
The following illustrations show combinations of screw terminal positions used to selectively ground different conductors on a multi-ratio current transformer. The first of these illustrations show the condition represented in the previous photograph, with the entire CT shorted and grounded:



This next illustration shows how the CT would be used in its full capacity, with X1 and X5 connecting to the panel instrument and (only) X5 grounded for safety:



This final illustration shows how the CT would be used in reduced capacity, with X2 and X3 connecting to the panel instrument and (only) X3 grounded for safety:



25.6.6 Instrument transformer burden and accuracy

In order for an instrument transformer to function as an accurate sensing device, it must not be unduly tasked with delivering power to a load. In order to minimize the power demand placed on instrument transformers, an ideal voltage-measuring instrument should draw zero current from its PT, while an ideal current-measuring instrument should drop zero voltage across its CT.

The goal of delivering zero power to any instrument is difficult to achieve in practice. Every voltmeter does indeed draw some current, however slight. Every ammeter does drop some voltage, however slight. The amount of apparent power drawn from any instrument transformer is appropriately called *burden*, and like all expressions of apparent power is measured in units of volt-amps. The greater this burden, the more the instrument transformer's signal will "sag" (decrease from loading). Therefore, minimizing burden is a matter of maximizing accuracy for power system measurement.

The burden value for any instrument transformer is a function of apparent power, impedance, and either voltage or current according to the familiar apparent power formulae $S = \frac{V^2}{Z}$ and $S = I^2 Z$:

$$\text{PT burden} = \frac{V_{\text{signal}}^2}{Z_{\text{instrument}}}$$

$$\text{CT burden} = (I_{\text{signal}}^2)(Z_{\text{instrument}})$$

Burden for any device or circuit connected to an instrument transformer may be expressed as an impedance value (Z) in ohms, or as an apparent power value (S) in volt-amps. Similarly, instrument transformers themselves are usually rated for the amount of burden they may source and still perform within a certain accuracy tolerance (e.g. $\pm 1\%$ at a burden of 2 VA).

Potential transformer burden and accuracy ratings

Potential transformers have maximum burden values specified in terms of apparent power (S , measured in volt-amps), standard burden values being classified by letter code:

Letter code	Maximum allowable burden at stated accuracy
W	12.5 volt-amps
X	25 volt-amps
M	35 volt-amps
Y	75 volt-amps
Z	200 volt-amps
ZZ	400 volt-amps

Standard accuracy classes for potential transformers include 0.3, 0.6, and 1.2, corresponding to uncertainties of $\pm 0.3\%$, $\pm 0.6\%$, and $\pm 1.2\%$ of the rated turns ratio, respectively. These accuracy class and burden ratings are typically combined into one label. A potential transformer rated "0.6M" therefore has an accuracy of $\pm 0.6\%$ (this percentage being understood as its *turns ratio* accuracy) while powering a burden of 35 volt-amps at its nominal (e.g. 120 volts) output.

Current transformer burden and accuracy ratings

Current transformer accuracies and burdens are more complicated than potential transformer ratings. The principal reason for this is the wider range of CT application. If a current transformer is to be used for *metering* purposes (i.e. driving wattmeters, ammeters, and other instruments used for regulatory control and/or revenue billing where high accuracy is required), it is assumed the transformer will operate within its standard rated current values. For example, a 600:5 ratio current transformer used for metering should rarely if ever see a primary current value exceeding 600 amps, or a secondary current exceeding 5 amps. If current values through the CT ever do exceed these maximum standard values, the effect on regulation or billing will be negligible because these should be transient events. However, protective relays are designed to interpret and act upon transient events in power systems. If a current transformer is to be used for *relaying* rather than metering, it must reliably perform under overload conditions typically created by power system faults. In other words, relay applications of CTs demand a much larger dynamic range of measurement than meter applications. Absolute accuracy is not as important for relays, but we must ensure the CT will give a reasonably accurate representation of line current during fault conditions in order for the protective relay(s) to function properly. PTs, even those used for protective relaying purposes, never see voltage transients as wide-ranging as the current transients seen by CTs.

Meter class CT ratings typically take the form of a percentage value followed by the letter “B” followed by the maximum burden expressed in ohms of impedance. Therefore, a CT with a metering classification of 0.3B1.8 exhibits an accuracy of $\pm 0.3\%$ of turns ratio when powering a 1.8 ohm meter impedance at 100% output current (typically 5 amps).

Relay class CT ratings typically take the form of a maximum *voltage* value dropped across the burden at 20 times rated current (i.e. 100 amps secondary current for a CT with a 5 amp nominal output rating) while maintaining an accuracy within $\pm 10\%$ of the rated turns ratio. Not coincidentally, this is how CT ratios are usually selected for power system protection: such that the maximum expected symmetrical fault current through the power conductor does not exceed 20 times the primary current rating of the CT³³. Therefore, a CT with a relay classification of C200 is able to output up to 200 volts while powering its maximum burden at $20\times$ rated current. Assuming a rated output current of 5 amps, 20 times this value would be 100 amps delivered to the relay. If the relay’s voltage drop at this current is allowed to be as high as 200 volts, it means the CT secondary circuit may have an impedance value of up to 2 ohms ($200\text{ V} \div 100\text{ A} = 2\ \Omega$). Therefore, a relaying CT rating of C200 is just another way of saying it can power as much as 2 ohms of burden.

The letter “C” in the “C200” rating example stands for *calculated*, which means the rating is based on theory. Some current transformers use the letter “T” instead, which stands for *tested*. These CTs have been actually tested at the specified voltage and current values to ensure their performance under real-world conditions.

³³For example, in an application where the maximum fault current is expected to be 40,000 amps, we would choose a CT with a ratio of at least 2000:5 to drive the protective relay, because 40,000 amps is twenty times this CT’s primary current rating of 2000 amps. We could also select a CT with a larger ratio such as 3000:5. The point is to have the CT be able to faithfully transform any reasonable fault current into a proportionately lower value for the protective relay(s) to sense.

Current transformer saturation

It is worthwhile to explore the concept of maximum CT burden in some detail. In an ideal world, a CT acts as a current source to the meter or relay it is powering, and as such it is quite content to drive current into a short circuit (0 ohms impedance). Problems arise if we demand the CT to supply more power than it is designed to, which means forcing the CT to drive current through an excessive amount of impedance. In the days of electromechanical meters and protective relays where the devices were entirely powered by instrument transformer signals, the amount of burden imposed by certain meters and relays could be quite substantial³⁴. Modern electronic meters and relays pose much less burden to instrument transformers, approaching the ideal conditions of zero impedance for current-sensing inputs.

The voltage developed by any inductance, including transformer windings, is described by Faraday's Law of Electromagnetic Induction:

$$V = N \frac{d\phi}{dt}$$


Where,

V = Induced voltage (volts)

N = Number of turns of wire

$\frac{d\phi}{dt}$ = Rate of change of magnetic flux (Webers per second)

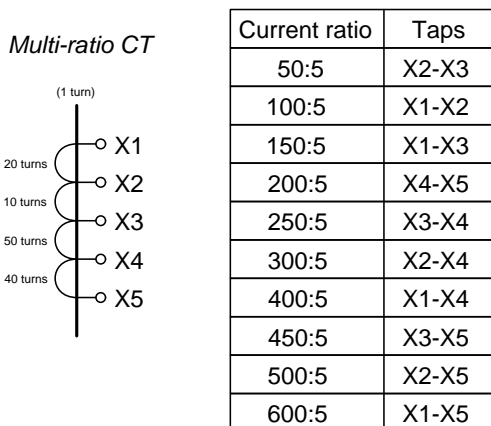
To generate a larger voltage, therefore, a current transformer must develop a faster-changing magnetic flux in its core. If the voltage in question is sinusoidal at a constant frequency, the magnetic flux also traces a sinusoidal function over time, the voltage peaks coinciding with the steepest points on the flux waveform, and the voltage "zero" points coinciding with the peaks on the flux waveform where the rate-of-change of magnetic flux over time is zero:



³⁴An illustrative example to consider is the venerable Westinghouse model CO-11 overcurrent relay, exhibiting a burden of 1.07 volt-amperes at a CT secondary current of 5 amperes with a 5-ampere tap setting. By contrast, an SEL-551 digital overcurrent relay exhibits only 0.16 volt-amperes of burden at the same CT current of 5 amperes: nearly *seven times less* burden than the electromechanical relay. The reason for this stark disparity in burden values is the design of each relay: the electromechanical relay demands power from the CT to spin an aluminum disk against the restraining forces of a spring and a drag magnet, while the electronic relay receives operating power from a separate source (station power) and only requires that the CT drive the input of an analog-to-digital converter (ADC) circuit.

Imposing a larger burden on a CT (i.e. more impedance the current must drive through) means the CT must develop a larger sinusoidal voltage for any given amount of measured line current. This equates to a flux waveform with a faster-changing rate of rise and fall, which in turn means a higher-peak flux waveform (assuming a sinusoidal shape). The problem with this at some point is that the required magnetic flux reaches such high peak values that the ferrous³⁵ core of the CT begins to saturate with magnetism, at which point the CT ceases to behave in a linear fashion and will no longer faithfully reproduce the shape and magnitude of the power line current waveform. In simple terms, if we place too much burden on a CT it will begin to output a distorted signal no longer faithfully representing line current.

The fact that a CT's maximum AC voltage output depends on the magnetic saturation limit of its ferrous core becomes particularly relevant to *multi-ratio* CTs where the secondary winding is provided with more than two “taps”. Multi-ratio current transformers are commonly found as the permanently mounted CTs in the bushings of power transformers, giving the end-user freedom in configuring their metering and protection circuits. Consider this distribution transformer bushing 600:5 CT with a C800 accuracy class rating:



This CT's “C800” classification is based on its ability to source a maximum of 800 volts to a burden *when all of its secondary turns are in use*. That is to say, its rating is “C800” only when connected to taps X1 and X5 for the full 600:5 ratio. If someone happens to connect to taps X1-X3 instead, using only 30 turns of wire in the CT's secondary instead of all 120 turns, this CT will be limited to sourcing 200 volts to a burden before saturating: the same magnitude of magnetic flux that could generate 800 volts across 120 turns of wire can only induce one-quarter as much voltage across one-quarter the number of turns, in accordance with Faraday's Law of Electromagnetic Induction ($V = N \frac{d\phi}{dt}$). Thus, the CT must be treated as a “C200” unit when wired for a 150:5 ratio.

³⁵Iron and iron alloys (“ferrous”) reach a point of maximum magnetization where all the magnetic “domains” in a sample are oriented in the same direction, leaving no more left to orient. Once a sample of ferrous material has thus “saturated”, it is of no further benefit to the establishment of a magnetic field. Increases in magnetic force will still produce additional lines of magnetic flux, but not at the rate experienced when the material was not saturated. In other words, a magnetically saturated inductor or transformer core essentially behaves like an air-core inductor or transformer for all additional current values beyond full saturation.

The presence of any direct current in AC power line conductors poses a problem for current transformers which may only be understood in terms of magnetic flux in the CT core. Any direct current (DC) in a power line passing through a CT biases the CT's magnetic field by some amount, making CT saturate more easily in one half-cycle of the AC than the other. Direct currents never sustain indefinitely in AC power systems, but are often present as transient pulses during certain fault conditions. Even so, transient DC currents will leave CT cores with some residual magnetic bias predisposing them to saturation in future fault conditions. The ability of a CT core to retain some magnetic flux over time is called *remanence*.


Remanence in a transformer core is an undesirable property. It may be mitigated by designing the core with a air gap (rather than making the core as an unbroken path of ferrous metal), but this compromises other desirable properties such as saturation limits (i.e. maximum output voltage). Some industry experts advise CTs be demagnetized by maintenance personnel as part of the repair work following a high-current fault, in order to ensure optimum performance when the system is returned to service. Demagnetization consists of passing a large AC current through the CT and then slowly reducing the magnitude of that AC current to zero amps. The gradual reduction of alternating magnetic field strength from full to zero tends to randomize the magnetic domains in the ferrous core, returning it to an unmagnetized state.

Whatever the cause, CT saturation can be a significant problem for protective relay circuits because these relays must reliably operate under all manner of transient overcurrent events. The more current through the primary of a CT, the more current it should output to the protective relay. For any given amount of relay burden (relay input impedance), a greater current signal translates into a greater voltage drop and therefore a greater demand for the CT to output a driving voltage. Thus, CT saturation is more likely to occur during overcurrent events when we most need the CT to function properly. Anyone tasked with selecting an appropriate current transformer for a protective relaying application must therefore carefully consider the maximum expected value of overcurrent for system faults, ensuring the CT(s) will do their job while driving the burdens imposed by the relays.

Current transformer testing

Current transformers may be bench-tested for turns ratio and saturation³⁶ by applying a variable AC voltage to the secondary winding while monitoring secondary current and primary voltage. For common “window” style CTs, the primary winding is a single wire threaded through its center hole. An ideal current transformer would present a constant impedance to the AC voltage source and a constant voltage ratio from input to output. A real current transformer will exhibit less and less impedance as voltage is increased past its saturation threshold:

Testing a CT for ratio and saturation



An ideal CT (with no saturation) would trace a straight line. The bent shape reveals the effects of magnetic saturation, where there is so much magnetism in the CT’s core that additional current only yields miniscule increases in magnetic flux (revealed by voltage drop).

Of course, a CT is never *powered* by its secondary winding when installed and operating. The purpose of powering a CT “backwards” as shown is to avoid having to drive very high currents through the primary of the CT. If high-current test equipment is available, however, such a *primary injection test* is actually the most realistic way to test a CT.

³⁶In the electric power industry this is commonly referred to as a “rat/sat” test.

The following table shows actual voltage and current values taken during a secondary excitation test on a C400 class relay CT with a 2000:5 ratio. The source voltage was increased from zero to approximately 600 volts AC at 60 Hz for the test while secondary voltage drop and primary voltage were measured. Around 575 volts a “buzzing” sound could be heard coming from the CT – an audible effect of magnetic saturation. Calculated values of secondary winding impedance and turns ratio are also shown in this table:

I_S	V_S	V_P	$Z_S = V_S \div I_S$	Ratio = $V_S \div V_P$
0.0308 A	75.14 V	0.1788 V	2.44 k Ω	420.2
0.0322 A	100.03 V	0.2406 V	3.11 k Ω	415.8
0.0375 A	150.11 V	0.3661 V	4.00 k Ω	410.0
0.0492 A	301.5 V	0.7492 V	6.13 k Ω	402.4
0.0589 A	403.8 V	1.0086 V	6.86 k Ω	400.4
0.0720 A	500.7 V	1.2397 V	6.95 k Ω	403.9
0.0883 A	548.7 V	1.3619 V	6.21 k Ω	402.9
0.1134 A	575.2 V	1.4269 V	5.07 k Ω	403.1
0.1259 A	582.0 V	1.4449 V	4.62 k Ω	402.8
0.1596 A	591.3 V	1.4665 V	3.70 k Ω	403.2
0.2038 A	600.1 V	1.4911 V	2.94 k Ω	402.5

As you can see from this table, the calculated secondary winding impedance Z_S begins to drop dramatically as the secondary voltage exceeds 500 volts (near the “knee” point of the curve). The calculated turns ratio appears remarkably stable – close to the ideal value of 400 for a 2000:5 CT – but one must remember this ratio is calculated on the basis of *voltage* and not current. Since this test does not compare primary and secondary currents, we cannot see the effects saturation would have on this CT’s current-sensing ability. In other words, this test reveals when saturation begins to take place, but it does not necessarily reveal how the CT’s current ratio is affected by saturation.

What makes the difference between a 2000:5 ratio CT with a relay classification of C400 and a 2000:5 ratio CT with a relay classification of C800 is not the number of turns in the CT’s secondary winding (N)³⁷ but rather the amount of ferrous metal in the CT’s core. The C800 transformer, in order to develop upwards of 800 volts to satisfy relay burden, must be able to sustain twice as much magnetic flux in its core than the C400 transformer, and this requires a magnetic core in the C800 transformer with (at least) twice as much flux-carrying capacity. All other factors being equal, the higher the burden capacity of a CT, the larger and heavier it must be due to the girth of its magnetic core.

³⁷If you think carefully about this, you realize that the number of turns of wire in either CT must be identical, because there is only one “turn” of wire passing through the center of either CT. In order to achieve a 2000:5 ratio, you must have 400 turns of wire wrapped around the toroidal ferrous core per the 1 “turn” of wire passing through the center of that core.

Current transformer circuit wire resistance

The burden experienced by an operating current transformer is the total series impedance of the measuring circuit, consisting of the sum of the receiving instrument's input impedance, the total wire impedance, and the internal secondary winding impedance of the CT itself. Legacy electromechanical relays, with their "operate" coils driven by CT currents, posed significant burden. Since the burden imposed by an electromechanical relay stems from the operation of a wire coil, this burden impedance is a complex quantity having both real (resistive) and imaginary (reactive) components. Modern digital relays with analog-to-digital converters at their inputs generally pose purely resistive burdens to their CTs, and those burden values are generally much less than the burdens imposed by electromechanical relays.

A significant source of burden in any CT circuit is the resistance of the wire carrying the CT's output current to and from the receiving instrument. It is quite common for the total "loop" distance of a CT circuit to be several hundred feet or more if the CTs are located in remote areas of a facility and the protective relays are located in a central control room. For this reason an important aspect of protective relay system design is wire size (gauge), in order to ensure the total circuit resistance does not exceed the CT's burden rating.

Larger-gauge wire has less resistance per unit length than smaller-gauge wire, all other factors being equal. A useful formula for approximating the resistance of copper wire is shown here:

$$R_{1000ft} = e^{0.232G-2.32}$$

Where,

R_{1000ft} = Approximate wire resistance in ohms per 1000 feet of wire length

G = American Wire Gauge (AWG) number of the wire

AWG wire sizes, like most "gauge" scales, is inverse: a larger number represents a thinner wire. This is why the formula predicts a smaller R value for a larger G value. An easy example value to plug into this formula is the number 10 representing #10 AWG wire, a common conductor size for CT secondary circuits:

$$R_{1000ft} = e^{(0.232)(10)-2.32}$$

$$R_{1000ft} = e^{2.32-2.32}$$


$$R_{1000ft} = e^0 = 1 \Omega \text{ per 1000 feet}$$

Bear in mind that this result of 1 ohm³⁸ wire resistance per 1000 feet of length applies to the *total circuit length*, not the distance between the CT and the receiving instrument. A complete CT secondary electrical circuit of course requires *two* conductors, and so 1000 feet of wire will be needed to cover 500 feet of distance between the CT and the instrument. Some sources cite #12 AWG wire as the minimum gauge to use for CT secondary circuits regardless of wire length.

³⁸Calculations based on the specific resistance of copper at 20 °C place 10 AWG wire at 0.9989 ohms per 1000 feet.
 $R = \frac{\rho l}{A}$

Example: CT circuit wire sizing, simple


A practical example will help illustrate how wire resistance plays a role in CT circuit performance. Let us begin by considering a C400 accuracy class current transformer to be used in a protective relay circuit, the CT itself possessing a measured secondary winding resistance of 0.3Ω with a 600:5 turns ratio. By definition, a C400 current transformer is one capable of generating 400 volts at its terminals while supplying 20 times its rated current to a burden. This means the maximum burden value is 4 ohms, since that is the impedance which will drop 400 volts at a secondary current of 100 amps (20 times the CT's nominal output rating of 5 amps):

C400 current transformer at 20 times rated current, maximum burden value

Although the CT has a C400 class rating which means 400 volts (maximum) produced at its terminals, the winding must actually be able to produce more than 400 volts in order to overcome the voltage drop of its own internal winding resistance. In this case, with a winding resistance of 0.3 ohms carrying 100 amps of current (worst-case), the winding voltage must be 430 volts in order to deliver 400 volts at the terminals. This value of 430 volts, at 60 Hz with a sinusoidal current waveform, represents the maximum amount of magnetic flux this CT's core can handle while maintaining a current ratio within $\pm 10\%$ of its 600:5 rating. Thus, 430 volts (inside the CT) is our limiting factor for the CT's winding at *any* current value.

This step of calculating the CT's maximum internal winding voltage is not merely an illustration of how a CT's "C" class rating is defined. Rather, this is an essential step in any analysis of CT circuit burden because we must know the maximum winding potential the CT is limited to. One might be tempted to skip this step and simply use 400 volts as the maximum terminal voltage during a fault condition, but doing so will lead to minor errors in a simple case such as this, and much more significant errors in other cases where we must de-rate the CT's winding voltage for reasons described later in this section.

Suppose this CT will be used to supply current to a protective relay presenting a purely resistive burden of 0.2 ohms. A system study reveals maximum symmetrical fault current to be 10,000 amps, which is just below the 20× rated primary current for the CT. Here is what the circuit will look like during this fault condition with the CT producing its maximum (internal) voltage of 430 volts:



The CT's internal voltage limit of 430 volts still holds true, because this is a function of its core's magnetic flux capacity and not line current. With a power system fault current of 10,000 amps, this CT will only output 83.33 amps rather than the 100 amps used to define its C400 classification. The maximum total circuit resistance is easily predicted by Ohm's Law, with 430 volts (limited by the CT's magnetic core) pushing 83.33 amps (limited by the system fault current):

$$R_{total} = \frac{V_w}{I_{fault}} = \frac{430 \text{ V}}{83.33 \text{ A}} = 5.16 \Omega$$

Since we know the total resistance in this series circuit is the sum of CT winding resistance, wire resistance, and relay burden, we may easily calculate maximum wire resistance by subtraction:

$$R_{total} = R_{CT} + R_{wire} + R_{relay}$$

$$R_{wire} = R_{total} - (R_{CT} + R_{relay})$$

$$R_{wire} = 5.16 \Omega - (0.3 \Omega + 0.2 \Omega) = 4.66 \Omega$$

Thus, we are allowed to have up to 4.66 Ω of total wire resistance in this CT circuit while remaining within the CT's ratings. Assuming the use of 12 gauge copper wire:

$$R_{1000ft} = e^{(0.232)(12) - 2.32} = 1.59 \Omega \text{ per } 1000 \text{ feet}$$

$$\frac{4.66 \Omega}{1.59 \Omega/1000 \text{ ft}} = 2.93 \times 1000 \text{ ft} = 2930 \text{ ft}$$

Of course, this is *total* conductor length, which means for a two-conductor cable between the CT and the protective relay the maximum distance will be half as much: 1465 feet.

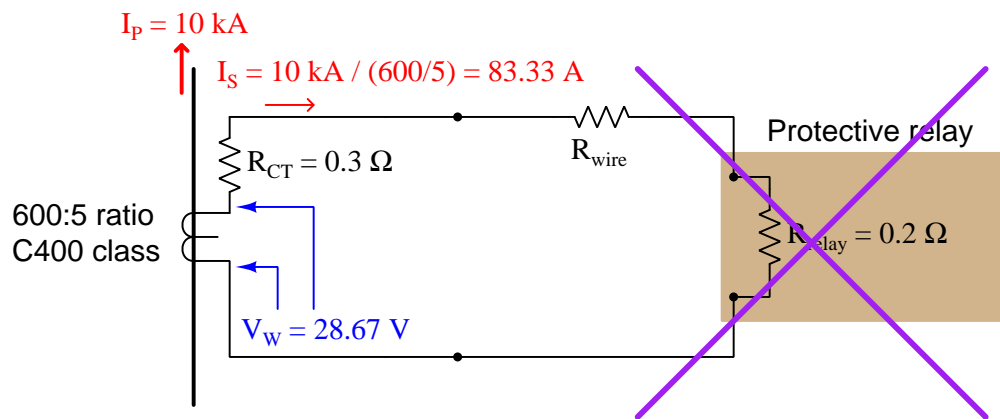
Example: CT circuit wire sizing, with DC considered

The previous scenario assumes purely AC fault current. Real faults may contain significant DC components for short periods of time, the duration of these DC transients being related to the $\frac{L}{R}$ time constant of the power circuit. As previously mentioned, direct current tends to magnetize the ferrous core of a CT, predisposing it to magnetic saturation. Thus, a CT under these conditions will not be able to generate the full AC voltage possible during a controlled bench test (e.g. a C400 current transformer under these conditions will not be able to live up to its 400-volt terminal rating). A simple way to compensate for this effect is to de-rate the CT's winding voltage by a factor equal to $1 + \frac{X}{R}$, the ratio $\frac{X}{R}$ being the reactance-to-resistance ratio of the power system at the point of measurement. De-rating the transformer provides a margin of safety for our calculations, anticipating that a fair amount of the CT's magnetic core capacity may be consumed by DC magnetization during certain faults, leaving less magnetic "headroom" to generate an AC voltage.

Let's re-do our calculations assuming the power system being protected now has an $\frac{X}{R}$ ratio of 14. This means our C400 current transformer (with a maximum internal winding potential of 430 volts) must be "de-rated" to a maximum winding voltage of:

$$\frac{430 \text{ V}}{1 + \frac{X}{R}} = \frac{430 \text{ V}}{1 + 14} = 28.67 \text{ V}$$


If we apply this de-rated winding voltage to the same CT circuit, we find it is insufficient to drive 83.33 amps through the relay:



This cannot work! 28.67 volts is not enough to drive 83.33 amps through the circuit even with no wire resistance at all!!

With 0.5Ω of combined CT and relay resistance (and no wire resistance), a winding voltage of 28.67 volts could only drive 57.33 amps which is far less than we need. Clearly this CT will not be able to perform under fault conditions where DC transients push it closer to magnetic saturation.

Upgrading the CT to a different model having a higher accuracy class (C800) and a larger current step-down ratio (1200:5) will improve matters. Assuming an internal winding resistance of 0.7 ohms for this new CT, we may calculate its maximum internal winding voltage as follows: if this CT is rated to supply 800 volts at its terminals at 100 amps secondary current through 0.7 ohms of internal resistance, it must mean the CT's secondary winding internally generates 70 volts more than the 800 volts at its terminals, or 870 volts under purely AC conditions. With our power system's $\frac{X}{R}$ ratio of 14 to account for DC transients, this means we must de-rate the CT's internal winding voltage from 870 volts to 15 times less, or 58 volts. Applying this new CT to the previous fault scenario:



Calculating the allowable total circuit resistance given the new CT's improved voltage:

$$R_{total} = \frac{V_w}{I_{fault}} = \frac{58 \text{ V}}{41.67 \text{ A}} = 1.392 \Omega$$

Once again, we may calculate maximum wire resistance by subtracting all other resistances from the maximum total circuit resistance:

$$R_{wire} = R_{total} - (R_{CT} + R_{relay})$$

$$R_{wire} = 1.392 \Omega - (0.7 \Omega + 0.2 \Omega) = 0.492 \Omega$$

Thus, we are allowed to have up to 0.492 Ω of wire resistance in this circuit while remaining within the CT's ratings. Using 10 AWG copper wire (exhibiting 1 ohm per 1000 feet), this allows us a total conductor length of 492 feet, which is 246 feet of distance between the CT terminals and the relay terminals.

25.7 Introduction to protective relaying

Circuit breakers used in residential, commercial, and light industrial service are self-tripping devices: they internally sense the amount of electric current going through them, automatically opening when that current exceeds a pre-determined level. Circuit breakers used in medium-voltage (2.4 kV to 35 kV) and higher-voltage applications, however, must be *triggered* to trip by external devices. This remote-control philosophy not only eliminates the technical problem of incorporating accurate and rugged current-sensing devices inside the body of a large circuit breaker, but it also opens the possibility of having these circuit breakers trip and close based on practically any condition imaginable, not just overcurrent.

An electrical device designed to detect some specified condition in a power system, and then command a circuit breaker either to trip or to close in order to protect the integrity of the power system, is called a *protective relay*. As we will see in this chapter, there is a wide variety of protective relay types and functions: overcurrent is just one of a great many power system conditions monitored and guarded against by protective relays.


The following photograph shows a pair of protective relays installed in the control panel for a medium-voltage power distribution circuit breaker. The relay on the left (just above the manual trip/close control switch) is a “time overcurrent” unit, designed to automatically trip the circuit breaker based on the product of current and time³⁹. The relay on the right (just above the “Reclose Cutout” switch) is a *reclosing* relay, designed to automatically trip the circuit breaker in the event of an instantaneous overcurrent event (e.g. a tree branch short-circuiting a power line) and then automatically re-close the breaker to test whether or not the fault has cleared. If the fault clears on its own, the breaker remains closed; if the fault remains, the reclosing relay will trip the breaker again.



If ever you have experienced the electric power service to your home “blink” a few times and then resume as normal, you have been the beneficiary of a reclosing relay. If it were not for the reclosing relay’s programmed strategy of multiple attempts to restore power, your electric service would be shut off for extended periods of time following *any* momentary power line fault.

³⁹What this means is that the relay will permit the circuit breaker to remain in its closed state indefinitely so long as the current is at or below 100% of its rated value. If the current ever exceeds the 100% limit, the protective relay begins to measure the length of time for the overcurrent event, commanding the circuit breaker to trip open after a certain amount of time inversely proportional to the degree of overcurrent. A 300% overcurrent condition, for example, will cause the circuit breaker to trip in a shorter amount of time than a 200% overcurrent condition.

An illustrative diagram shows how a simple protective relay monitors and interrupts power. The protective relay senses load current via the three line current transformers (CTs), closing a “trip” contact to trip the circuit breaker if ever the line current exceeds any limits pre-programmed into the relay:



Inside most remote-tripping circuit breakers is an auxiliary contact (sometimes designated “52a”) connected in series with the trip coil. This auxiliary contact is actuated by the same mechanism actuating the three large power contacts inside the circuit breaker, and thus the auxiliary contact will be closed when the breaker is closed and open when the breaker is tripped. The purpose of this normally-open auxiliary contact is to interrupt power to the trip coil once the circuit breaker has reached the trip position, so that the trip coil does not overheat (nor will the station battery be needlessly discharged) if the protective relay happens⁴⁰ to continuously output a trip command signal.

⁴⁰In many legacy electromechanical protective relays, the trip contact is designed to latch in the closed position even after the event prompting the breaker trip has passed. A special “seal-in” circuit with its own coil and contact provides this latching action, the purpose of which is to ensure the relay will continuously command the breaker to trip for as long as it takes for the breaker to reach the tripped condition. Only the 52a auxiliary contact inside the circuit breaker can interrupt a latched trip circuit, and that will only happen when the breaker achieves a tripped state.

Note the use of a 125 volt *DC* “station battery” supply for the circuit breaker’s “trip” circuit. A battery provides uninterruptible DC power, so that breakers may be tripped and closed even in the event of a total AC power failure in the facility. A photograph of the station battery for a large substation is shown here:




Protective relay circuits have been powered by station batteries for many decades, because a large battery bank is the simplest form of *uninterruptible power supply* (UPS) in existence. A trickle-charging AC-to-DC power supply keeps the station battery in a constant state of full charge while AC power is available. In the event of an AC power interruption, all protective relays and other critical instrumentation in the facility will continue to operate normally. Even the most modern digital protective relays operate on the traditional 125 VDC supply voltage⁴¹ rather than 120 VAC as is common with other types of industrial controls.

⁴¹It should be noted that some microprocessor-based protective relays may operate on DC *or* AC power, as well as at power supply voltages other than 125 volts, in addition to the standard of 125 VDC.

Protective relays have seen widespread use in industrialized power systems since the early twentieth century, with continued technological development. The earliest protective relay technologies were electromagnetic in design, a great many of them based on the “induction disk” design whereby out-of-phase AC magnetic fields caused an aluminum disk to torque like the rotor of an induction electric motor. The induction disk technology became popular as the basis for rotating-disk watt-hour meters used in residential and commercial electric power service as well.

An example of an induction disk protective relay typical of the genre is seen here, a General Electric model 121AC overcurrent relay:




This relay uses an aluminum disk approximately 4 inches in diameter to sense and time overcurrent conditions, the disk slowly rotated by the torque generated from a set of electromagnet coils energized by current received from a current transformer (CT). In order for the disk to rotate at all, the coils’ induced torque must exceed the restraining torque applied on the disk’s shaft by a spiral spring. This amount of current necessary to overcome the spring’s torque is called the *pickup* current value for the induction relay. Current in excess of the pickup value causes the disk

to slowly rotate, the speed of rotation being a function of current magnitude (more current = faster rotation). If the disk rotates all the way to its end-point, it closes an electrical contact to signal a “time-overcurrent” trip, causing the system’s circuit breaker to trip open and interrupt current.

Protective relays such as this General Electric model were built to be “drawn out” of their sockets for easy maintenance and replacement. The relay shown in the above photograph has already been drawn out of its case for inspection.

Later protective relay designs used electronic circuits rather than electromagnetic mechanisms to detect and time overcurrent conditions. This Basler model BE1-79M “reclosing” relay is illustrative of early solid-state generation protective relays:



Like the earlier generation induction-disk relays, this electronic relay is also a “draw-out” style, permitting convenient maintenance and replacement.

A more modern example of a protective relay is this Schweitzer Engineering Laboratories model 551 overcurrent/reclosing relay:



The accuracy, stability, and reliability of modern microprocessor-based protective relays is such that there is no longer a need to regularly remove them for service and replacement. This is why the traditional “draw-out” design has been replaced by a more permanent rack-mount design.

Another advantage of a microprocessor-based relay design is the ability to communicate digitally with other microprocessor-based systems. This permits remote querying of the relay’s status and parameter settings. Additionally, the digital memory capabilities of a microprocessor-based relay allows for power instrument data (voltage, current, phase shift, timestamps, etc.) to be stored, so that personnel may determine the sequence of events leading up to a breaker trip.

An interesting footnote to modern protective relays is their persistent use of anachronistic terms. Even in the most modern protective relays such as the Schweitzer model 551 shown previously, you will find parameters inside the relay designated *torque control*, *time dial*, *pickup*, and *dropout*: all terms designed to describe moving components inside an electromagnetic relay mechanism such as the old General Electric model 121AC induction-disk unit. Protective relay controls were developed and perfected for so many years using electromagnetic relay technology that the nomenclature remains in common use even though the mechanisms inspiring these terms are obsolete. It is for this reason that electromechanical relay technology will be presented in this book’s discussions of protective relay functions: to orient the reader to the origins of the terms so that they will make more sense when encountered in modern protective relays.

As you can see, the strategy of using independent “relay” devices to command a large power circuit breaker to trip is a much more sophisticated way of ensuring power system protection and reliability than constructing each circuit breaker with its own internal overcurrent mechanism. This is truly an “instrumentation” approach to electric power control: intentionally locating the intelligence of the system to a set of dedicated control devices which may be upgraded and reconfigured on demand to meet continually changing needs.

25.8 ANSI/IEEE function number codes

In the United States, the ANSI and IEEE organizations have standardized a set of numerical codes referring to different types of power system devices and functions (IEEE C 37.2). Some of these codes refer to specific pieces of equipment (e.g. circuit breakers) while other codes refer to abstract functions (e.g. overcurrent protection). Two partial listings of these ANSI/IEEE code numbers show some of the devices and functions covered by the ANSI/IEEE standard:

ANSI/IEEE code	Device
33	Position switch
41	Field circuit breaker
52	AC power circuit breaker
57	Shorting/grounding switch
63	Pressure switch
70	Rheostat
71	Liquid level switch
72	DC power circuit breaker
80	Flow switch
84	Operating mechanism (generic)
88	Auxiliary motor or motor/generator
89	Line switch (power disconnect)

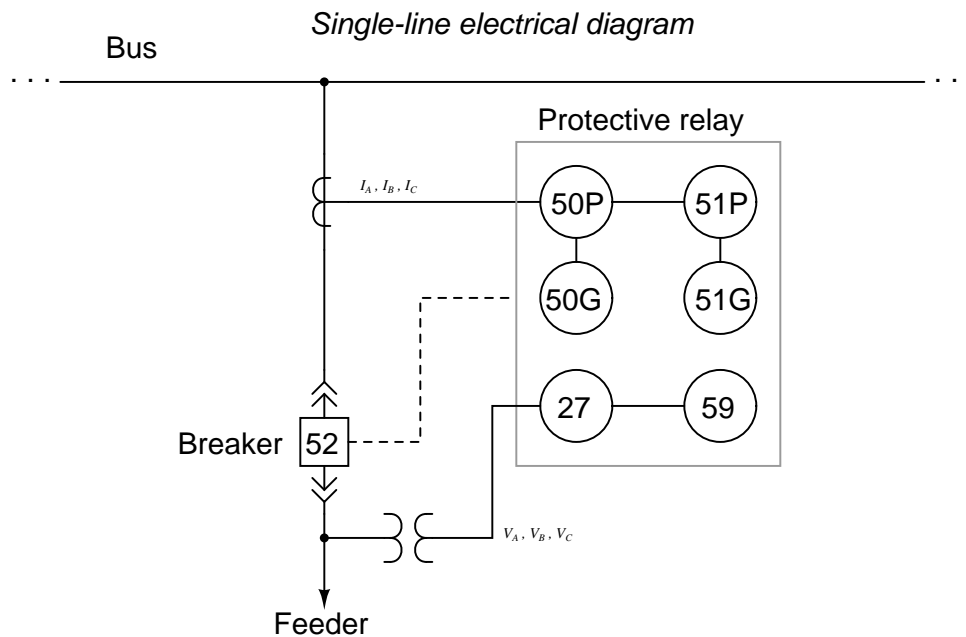
ANSI/IEEE code	Function
12	Over-speed
14	Under-speed
19	Reduced voltage start
21	Distance
23	Temperature control
24	V/Hz (overfluxing)
25	Synchronism check
27	Undervoltage
28	Flame safety detection
30	Annunciator
32	Directional (reverse) power
37	Undercurrent/underpower
38	Bearing overtemperature
40	Loss of excitation
43	Manual transfer/selecter
46	Current unbalance
46R	Broken conductor
47	Phase reversal
48	(Motor) stall
49	Thermal overload
50	Instantaneous overcurrent
50G	Instantaneous overcurrent (on ground conductor)
50ARC	Arc fault
51	Time overcurrent
51G	Time overcurrent (on ground conductor)
55	Power factor
58	Rectifier failure
59	Overvoltage
64	Ground fault
65	Speed governing
66	Starts per hour / time between starts
67	Directional overcurrent
68	Blocking
74	Alarm
78	Phase angle / out-of-step
79	Automatic reclose
81H/81L	Overfrequency/Underfrequency
81R	Rate of frequency change
86	Lockout or Auxiliary
87	Differential

It is typical to find multiple *functions* performed by a single *device* in an electrical power system. A common example of this is an instantaneous/time overcurrent relay: a single device monitoring the

signals coming from a set of current transformers (CTs), commanding a circuit breaker to trip if the current exceeds a pre-determined limit for any length of time (instantaneous overcurrent protection, ANSI/IEEE code 50) or if the time-current product exceeds a pre-determined limit (time overcurrent protection, ANSI/IEEE code 51). Both the 50 and 51 functions are usually implemented by the same protective relay. Modern digital electronic protective relays may provide a multitude of protective functions in one unit.


These code designations have become so common within industry parlance that it is typical to hear technicians and engineers alike refer to relays by number rather than name (e.g. “The 50/51 relays need to be calibrated next month”).

Protective relay functions are typically represented in single-line electrical diagrams as circles, with the ANSI/IEEE number code specifying each function. This is analogous to ISA-standard loop diagrams and P&IDs where instruments and control functions are represented as circles with ISA tagnames written inside the circles. Here is an example of a protective relay system for a circuit breaker sending power from a bus to a feeder:



In this system, a single protective relay device performs multiple functions: instantaneous overcurrent on the phase conductors (50P) and ground (50G), time overcurrent on the phase conductors (51P) and ground (51G), undervoltage (27), and overvoltage (59). Note how letters immediately following the number code qualify the purpose of the function, such as “G” for “ground” or “P” for “phase”. If the signals received from the CTs and/or PT suggest any of these abnormal conditions, the protective relay will send a “trip” command signal to the circuit breaker to open it. The circuit breaker itself is designated by the number code 52, as shown in the box symbol on the diagram.

ANSI/IEEE function codes also find application in relay trip circuit diagrams. Consider the following example of an electromechanical time-overcurrent (function 51) relay set, monitoring current through three power conductors and tripping the circuit breaker (device 52) if the current in any line exceeds safe levels. This format of diagram is typical for electromechanical protective relays, showing power circuitry on the left and trip circuitry on the right:




Note the labeling conventions used in the trip circuit diagram: each relay or breaker component bears a label beginning with its ANSI/IEEE device or function number. *52* refers to the power circuit breaker, *51* refers to the time-overcurrent function, dashed numbers specify which relay out of the three-relay set (one electromechanical overcurrent relay assembly per phase), and letters found below the horizontal line identify elements of the component's function (e.g. *TC* stands for *Trip Coil*, *a* refers to a form-A “normally open” contact inside a device). This labeling is used to advantage in eliminating duplicated lines and components in the trip diagram for relays 2 and 3 of the three-relay set (i.e. not having to show the seal-in coil, trip contact, or seal-in contact for the other two relays because their form is identical to those elements inside the first relay). As with ladder-style electrical diagrams, associations between components such as relay coils and relay contacts are done by *name* and not by physical proximity or dashed connecting lines as is the case with electronic schematics. For example, we can tell the left-hand current transformer monitoring current in line 1 activates relay number 1 because that is the label on the left-hand coil (51-1) connected to that CT. We can tell which coil activates relay 1's seal-in contact because the seal-in coil bears the same label (51-1/SI).

25.9 Instantaneous and time-overcurrent (50/51) protection


Perhaps the most basic and necessary protective relay function is *overcurrent*: commanding a circuit breaker to trip when the line current becomes excessive. The purpose of overcurrent protection is to guard against power distribution equipment damage, due to the fact that excessive current in a power system dissipates excessive heat in the metal conductors comprising that system. Overcurrent protection is also applied to machines such as motors and generators for the exact same reason: electric current dissipates heat in the windings' resistance ($P = I^2R$), and excessive heat will damage those winding conductors.

Instantaneous overcurrent protection is where a protective relay initiates a breaker trip based on current exceeding a pre-programmed "pickup" value for *any* length of time. This is the simplest form of overcurrent protection, both in concept and in implementation (relay design). In small, self-tripping circuit breakers, this type of protection is best modeled by "magnetic" breakers where the tripping mechanism is actuated by the magnetic field strength of the line conductors: any amount of current greater than the tripping threshold will cause the mechanism to unlatch and open the breaker. In protective relay-based systems, the instantaneous overcurrent protection function is designated by the ANSI/IEEE number code *50*.

Time overcurrent protection is where a protective relay initiates a breaker trip based on the combination of overcurrent magnitude and overcurrent duration, the relay tripping sooner with greater current magnitude. This is a more sophisticated form of overcurrent protection than instantaneous, expressed as a "time curve" relating overcurrent magnitude to trip time. In small, self-tripping circuit breakers, this type of protection is best modeled by "thermal" breakers where the tripping mechanism is actuated by the force of a bimetallic strip heated by line current: excessive current heats the metal strip, which then forces the mechanism to unlatch and open the breaker. In protective relay-based systems, the time overcurrent protection function is designated by the ANSI/IEEE number code *51*. Time overcurrent protection allows for significant overcurrent magnitudes, so long as these overcurrent events are brief enough that the power equipment avoids heat damage.



Electromechanical 50 (instantaneous overcurrent) relays are models of simplicity, consisting of nothing more than a coil⁴², armature, and contact assembly (a “relay” in the general electrical/electronic sense of the word). Spring tension holds the trip contacts open, but if the magnetic field developed by the CT secondary current becomes strong enough to overcome the spring’s tension, the contacts close, commanding the circuit breaker to trip:



The protective relay circuit in the above diagram is for one phase of the three-phase power system only. In practice, three different protective relay circuits (three CTs, and three 50 relays with their trip contacts wired in parallel) would be connected together to the circuit breaker’s trip coil, so that the breaker will trip if *any* of the 50 relays detect an instantaneous overcurrent condition. The monitoring of all three line currents is necessary because power line faults are usually unbalanced: one line will see a much greater share of the fault current than the other lines. A single 50 relay sensing current on a single line would not provide adequate instantaneous overcurrent protection for all three lines.

The amount of CT secondary current necessary to activate the 50 relay is called the *pickup current*. Its value may be varied by adjusting a movable magnetic pole inside the core of the relay. Calibration of an instantaneous overcurrent (50) relay consists simply of verifying that the unit “picks up” within a reasonably short amount of time if ever the current magnitude exceeds the prescribed pickup value.

⁴²In protective relay circuit diagrams, it is conventional to show relay coils as “zig-zag” symbols rather than as actual coils of wire as is customary in electronic schematics. Those familiar with “ladder” style electrical wiring diagrams may recognize this as the symbol for a *solenoid* coil. Once again, we see here the context-dependence of symbols and diagram types: a component type may have multiple symbols depending on which type of diagram it’s represented in, while a common symbol may have different meanings in different diagrams.


Electromechanical 51 (time overcurrent) relays are more complicated in design, using a rotating metal “induction disk” to physically time the overcurrent event, and trip the circuit breaker only if the overcurrent condition persists long enough. A photograph of a General Electric time-overcurrent induction-disk relay appears here:



The round disk you see in the photograph receives a torque from an electromagnet coil assembly acting like the stator coils of an induction motor: alternating current passing through these coils cause alternating magnetic fields to develop through the rear section of the disk, inducing currents in the aluminum disk, generating a “motor” torque on the disk to rotate it clockwise (as seen from the vantage point of the camera in the above photo). A spiral spring applies a counter-clockwise restraining torque to the disk’s shaft. The pickup value for the induction disk (i.e. the minimum amount of CT current necessary to overcome the spring’s torque and begin to rotate the disk) is established by the spring tension and the stator coil field strength. If the CT current exceeds the pickup value for a long enough time, the disk rotates until it closes a normally-open contact to send 125 VDC power to the circuit breaker’s trip coil.


A silver-colored permanent magnet assembly at the front of the disk provides a consistent “drag” force opposing disk rotation. As the aluminum disk rotates through the permanent magnet’s field, eddy currents induced in the disk set up their own magnetic poles to oppose the disk’s motion (Lenz’s Law). The effect is akin to having the disk rotate through a viscous liquid, and it is this dynamic retarding force that provides a repeatable, inverse time delay.

A set of three photographs show the motion of a peg mounted on the induction disk as it approaches the stationary trip contact. From left to right we see the disk in the resting position, partially rotated, and fully rotated:



The mechanical force actuating the time-overcurrent contact is not nearly as strong as the force actuating the instantaneous overcurrent contact. The peg may only lightly touch the stationary contact when it reaches its final position, failing to provide a secure and lasting electrical contact when needed. For this reason, a *seal-in relay* actuated by current in the 125 VDC trip circuit is provided to maintain firm electrical contact closure in parallel with the rotating peg contact. This “seal-in” contact ensures a reliable circuit breaker trip even if the peg momentarily brushes or bounces against the stationary contact. The parallel seal-in contact also helps reduce arcing at the peg’s contact by carrying most of the trip coil current.


A simplified diagram of an induction disk time-overcurrent relay is shown in the following diagram, for one phase of the three-phase power system only. In practice, three different protective relay circuits (three CTs, and three 51 relays with their trip contacts wired in parallel) would be connected together to the circuit breaker's trip coil, so that the breaker will trip if *any* of the 51 relays detect a timed overcurrent condition:



The seal-in unit is shown as an electromechanical relay connected with its contact in parallel with the induction disk contact, but with its actuating coil connected in series to sense the current in the 125 VDC trip circuit. Once the induction disk contact closes to initiate current in the DC trip circuit, even momentarily, the seal-in coil will energize which closes the seal-in contact and ensures the continuation of DC trip current to the circuit breaker's trip coil. The relay's seal-in function will subsequently maintain the trip command until some external contact opens to break the trip circuit, usually an auxiliary contact within the circuit breaker itself.

Calibration of a time overcurrent (51) relay consists first of verifying that the unit “picks up” (begins to time) if ever the current magnitude exceeds the prescribed pickup value. In electromagnetic relays such as the General Electric model showcased here, this setting may be coarsely adjusted by connecting a movable wire to one of several *taps* on a transformer coil inside the relay, varying the ratio of CT current sent to the induction disk stator coils. Each tap is labeled with the number of whole amperes (AC) delivered by the secondary winding of the CT required for relay pick-up⁴³ (e.g. a tap value of “5” means that approximately 5 amps of CT secondary current is required for induction disk pickup). A fine adjustment is provided in the form of a variable resistor in series with the stator coils.


A photograph of the tap wire setting (coarse pickup adjustment) and resistor (fine pickup adjustment) are shown here. The tap in this first photograph happens to be set at the 4 amp position:



Proper setting of the pickup tap value is determined by the maximum continuous current rating of the system being protected and the ratio of the current transformer (CT) used to sense that current.

⁴³Note that this General Electric relay provides pickup tap settings well in excess of 5 amps, which is the nominal full-load rating of most current transformers. CTs rated for protective relay applications are fully capable of exceeding their normal full-load capacity for short time periods, which is a necessary feature due to the extreme nature of fault current conditions. It is not uncommon for fault currents in a power system to exceed full-load current conditions by a factor of 20!

After the proper pickup value has been set, the time value is established by rotating a small wheel called the *time dial* located above the induction disk. This wheel functions as an adjustable stop for the induction disk's motion, positioning the disk closer to or farther away from the trip contact in its resting condition:



The amount of disk rotation necessary to close the trip contact may be set by adjusting the position of this time dial: a low number on the time dial (e.g. 1) means the disk need only rotate a small amount to close the contact; a high number on the time dial (e.g. 10) sets the resting position farther away from contact, so that the disk must rotate farther to trip. These time dial values are linear multipliers: a time dial setting of 10, for example, exhibits twice the time to trip than a setting of 5, for any given overload condition.

Calibration of the time-overcurrent protective function must be performed at multiple values of current exceeding the pickup value, in order to ensure the relay trips within the right amount of time for those current values. Like process instruments which are often calibrated at five points along their measurement range, time-overcurrent relays must also be checked at multiple points⁴⁴ along their prescribed “curve” in order to ensure the relay is performing the way it should.

Time overcurrent relays exhibit different “curves” relating trip time to multiples of pickup current. All 51 relays are *inverse* in that the amount of time to trip varies inversely with overcurrent magnitude: the greater the sensed current, the less time to trip. However, the function of trip time versus overcurrent magnitude is a curve, and several different curve shapes are available for United States applications:

- Moderately inverse
- Inverse
- Very inverse
- Extremely inverse
- Short-time inverse

⁴⁴Geometrically, at least three points are required to define the shape of any *curve*, just as two points are the minimum for defining a line. However, since the curvature of a relay's timing function is fixed by the construction of its components and therefore not liable to drift over time, it is common within the protective relay field to check the curve at just two points to ensure the adjustments are correct. The drag magnet is the principal adjustment for the timing of an electromechanical 51 relay.

Time curves standardized by the Swiss standards agency IEC (International Electrotechnical Commission) include:

- Standard inverse
- Very inverse
- Extremely inverse
- Long-time inverse
- Short-time inverse

The purpose for having different curves in time-overcurrent relays is related to a concept called *coordination*, where the 51 relay is just one of multiple overcurrent protection devices in a power system. Other overcurrent protection devices include fuses and additional 51 relays at different locations along the same line. Ideally, only the device closest to the fault will trip, allowing power to be maintained at all “upstream” locations. This means we want overcurrent protection devices at the remote end(s) of a power system to be more sensitive and to trip faster than devices closer to the source, where a trip would mean an interruption of power to a greater number of loads.

Legacy electromechanical time-overcurrent (51) relays implemented these different inverse curve functions by using induction disks with different “cam” shapes⁴⁵. Modern microprocessor-based 51 relays contain multiple curve functions as mathematical formulae stored within read-only memory (ROM), and as such may be programmed to implement any curve desired. It is an amusing anachronism that even in digital 51 relays containing no electromagnets or induction disks, you will find parameters labeled “pickup” and “time dial” in honor of legacy electromechanical relay behavior.

The trip time formulae programmed within a Schweitzer Engineering Laboratories model SEL-551 overcurrent relay for inverse, very inverse, and extremely inverse time functions are given here:

$$t = T \left(0.18 + \frac{5.95}{M^2 - 1} \right) \quad \text{Inverse curve}$$

$$t = T \left(0.0963 + \frac{3.88}{M^2 - 1} \right) \quad \text{Very inverse curve}$$

$$t = T \left(0.0352 + \frac{5.67}{M^2 - 1} \right) \quad \text{Extremely inverse curve}$$

Where,

t = Trip time (seconds)

T = Time Dial setting (typically 0.5 to 15)

M = Multiples of pickup current (e.g. if $I_{pickup} = 4.5$ amps, a 9.0 amp signal would be $M = 2$)


⁴⁵If you examine the induction disk from a 51 relay, you will note that the disk’s radius is not constant, and that there is actually a “step” along the circumference of the disk where its radius transitions from minimum to maximum. The amount of disk material exposed to the stator coil’s magnetic field to generate operating torque therefore changes with rotation angle, providing a nonlinear function altering the shape of the relay’s timing curve.

25.10 Differential (87) current protection

One of the fundamental laws of electric circuits is *Kirchhoff's Current Law*, which states the algebraic sum of all currents at a circuit node (junction) must be zero. A simpler way of stating this is to say “what goes in must come out.” We may exploit this principle to provide another form of protection against certain faults in electric circuits, by measuring the amount of current entering and exiting a circuit component, then tripping a circuit breaker if those two currents ever fail to match.

An important advantage of differential protection as compared to either instantaneous- or time-overcurrent protection is that it is far more sensitive and faster-acting. Unlike either form of overcurrent protection, which picks up only if current exceeds the maximum rating of the conductors, differential protection is able to pick up at far lower levels of current because Kirchhoff's Current Law predicts that *any* amount of current imbalance, for *any* length of time, is abnormal. Lower pick-up thresholds along with no time delay means that differential protection is able to take action sooner than any form of overcurrent protection can, thereby limiting equipment damage by clearing the fault in a shorter amount of time.

Suppose we were to measure the amount of current at both ends of every phase winding in a three-phase generator, shown in the following diagram:




Like most large power generators, this unit brings both terminals of each phase winding to external points so that they may be connected in either a Wye or a Delta configuration as desired. In this particular case, the generator's windings are Wye-connected. So long as we measure current going in and out of each winding individually, it matters little whether those generator windings are Wye- or Delta-connected.


If the circuit is exactly as drawn above, the amount of current entering and exiting each phase winding must be the same in accordance with Kirchhoff's Current Law. That is to say:

$$I_{A1} = I_{A2} \quad I_{B1} = I_{B2} \quad I_{C1} = I_{C2}$$

Suppose now that one of the turns within the “C” phase winding were to accidentally contact the generator’s metal frame, such as what might happen as a result of insulation damage. This *ground fault* will cause a third path for current in the faulted winding. I_{C1} and I_{C2} will now be imbalanced by an amount equal to the fault current I_F :




Another fault detectable by Kirchhoff’s Current Law is a phase-to-phase winding fault, where current flows from one winding to another. In this example, a fault between B and C phases in the generator upsets the balance of incoming and outgoing currents for both phases:



It should be noted that the magnitude of a ground fault or a winding-to-winding fault current might not be large enough to pose an overcurrent threat to the generator, yet the very existence of a current imbalance in any phase proves the winding is damaged. In other words, this is a type of system fault that would not necessarily be detected by an overcurrent (50/51) relay, and so must be detected by some other means.

The relay type designated for this task is called a *differential current* relay. The ANSI/IEEE number code for differential protection is 87. Differential voltage relays also exist, with the same “87” ANSI/IEEE designation, making it necessary to specify whether the differential quantity in question is voltage or current whenever mentioning an “87” relay.

A simple form of differential current protection for this generator may be implemented by connecting CTs on either side of each winding to operating coils of an electromechanical relay like this. For the sake of simplicity, protection for only one phase winding (C) of the generator will be shown. A practical differential current protective relay system would monitor current through all six stator wires on the generator, comparing currents in and out of every phase:




If the CT primary currents I_{C1p} and I_{C2p} are equal and the CT ratios are equal, the CT secondary currents I_{C1s} and I_{C2s} will be equal as well. The result will be zero⁴⁶ current through the operating coil (OC) of the differential relay.

If, however, a fault to ground or to an adjacent winding were to develop anywhere within the generator's "C" stator winding, the primary currents of the two CTs will become unequal, causing unequal secondary currents, thereby causing a substantial amount of current to flow through the differential relay's operate coil (OC). If this current is sufficient to cause the differential relay to "pick up", the relay will send a signal commanding the generator's circuit breaker to trip.

⁴⁶In practice, perfect cancellation of currents is nearly impossible due to mismatched CTs and other imperfections, and so a small amount of current typically passes through the differential relay's operating coil even under normal circumstances. The pickup value of this relay must be set such that this small amount of current does not unnecessarily trip the relay.

Even with the relay's pickup value biased to avoid unnecessary tripping, it is still possible that a heavy phase current demanded from the generator may cause the differential relay to trip, due to the impossibility of a perfect match between the two "C" phase current transformers. Any mismatch between these two CTs will result in an inequality of secondary currents that will become larger as phase current grows in magnitude. Large, harmonic-rich *inrush currents*⁴⁷ occasionally experienced when a large power transformer is initially energized may also cause false trips in this simple form of differential protection. We do not wish this differential relay to trip for any condition but an internal generator fault in its phase winding, and so a modification is necessary to provide a different operating characteristic.

If we modify the relay to have three coils, one to move its mechanism in the trip direction, and two to help "restrain" its mechanism (working to hold the mechanism in its normal operating position), we may connect these coils in such a way that the two restraint coils⁴⁸ (RC) are energized by the two CT secondary currents, while the operating coil only sees the difference between the two CT secondary currents. We refer to this scheme as a *restrained differential relay*, and the former (simpler) design as an *unrestrained differential relay*:



The general characteristic of a restrained differential relay is to trip on the basis of the differential current exceeding a set *percentage* of phase current.


⁴⁷Transformers exhibit inrush current for reasons different than capacitors (reactance) or motors (counter-EMF). Residual magnetism in a transformer core from the last time it was energized biases that core toward saturation in one direction. If the applied power happens to match that direction, and have sufficient magnitude, the transformer core will saturate on power-up which results in abnormally high current for multiple cycles until the core's magnetic domains normalize.

⁴⁸Restraint coils are sometimes labeled as "RC" and other times labeled as "R". It should be noted that the principle of a "restraining element" within a protective relay is not unique to differential (87) relays. Other relay types, notably distance (21) relays, also employ restraint coils or other mechanisms to prevent the relay from tripping under specific circumstances.

This photograph shows three differential relays used to protect the windings of a three-phase generator at a gas turbine power plant. Note how one differential current relay is required to protect each of the generator's three phases:




Modern digital differential relays typically sense CT signals from all three phases, allowing protection in a single panel-mount unit. Digital protective relays offer much more sophisticated approaches to the problem of false tripping based on mismatches between current transformer pairs and/or harmonic currents. The following graph shows the characteristic for a General Electric model 745 transformer protective relay providing differential current protection:



Not only may the pickup value be adjusted by the user, but also the slope of each line segment on the graph, the height of the “kneepoint” step, etc. Note how the term “restraint” is still used in digital relay configuration, even though it originated in electromechanical relay designs.


It is noteworthy that a form of differential current protection also finds application in American households, where electrical codes require the installation of *Ground Fault Current Interruptor* (GFCI) protected circuits in areas where contact between electrical appliances and water is likely (e.g. bathrooms, kitchens). GFCI receptacles function by sensing any difference in current between the “hot” and “neutral” conductors carrying current to and from any load plugged into the receptacle:



A single current transformer (CT) within the GFCI unit senses any differential current by sensing the *net* magnetic field around both current-carrying conductors. If the “hot” and “neutral” currents are equal, their opposite directions will produce opposing magnetic fields, with zero net magnetic field sensed by the CT. If, however, a ground fault exists at the load plugged into this receptacle, these two currents will be unequal and the CT will detect a net magnetic field. These protective devices are extraordinarily sensitive, tripping the contacts with differential current values in the *milliamp* range. This is important, as a ground fault existing in an electrical appliance may very well pass through the body of a person or an animal, in which case mere milliamps may prove harmful or even fatal.

If a GFCI receptacle trips, it may be reset by pressing a “reset” button on its face. GFCI units may also be manually tested by pressing a “test” button also mounted on the front face.

A very important concept in the field of protective relaying is that of *protection zones*, which is easily explained in the context of differential current relays. Simply defined, a relay’s “protection zone” is the physical range wherein the specified electrical fault may be detected, and thereby any components and connections within the zone may be protected through proper relay action. Overcurrent (50/51) relays discussed in an earlier section of this book do not exhibit well-defined zones of protection, since overcurrent relays pick up on a certain minimum fault current *value*, not necessarily on any certain fault *location*. Differential current relays, however, exhibit very clear and unambiguous zones of protection: *the area lying between the current-sensing CT pair*:




Only a fault within the relay’s protection zone (i.e. an “internal” fault) is capable of forcing the two CTs currents to become unequal. Thanks to Kirchhoff’s Current Law, no fault outside the protection zone (i.e. an “external” fault), no matter how severe, can make the CT primary currents become unequal⁴⁹.

The concept of protection zones is a very important one in protective relaying, and finds application well beyond differential current (87) systems. It is closely related to the concept of *selectivity*, which means the ability of a protective relay to discriminate between a fault within its own protection zone and one lying outside of its zone. A relay with high selectivity is one capable of ignoring external faults, while a relay with poor selectivity may falsely trip when faced with external faults.


⁴⁹It should be mentioned that an external fault generating currents high enough to saturate one or more of the CTs used in the differential protection system may cause the differential current system to falsely trip, due to saturation causing the affected CT(s) to no longer faithfully represent line current to the relay.

Ground Fault Current-Interrupting (GFCI) household power receptacles also exhibit well-defined zones of protection. In the case of a GFCI the zone of protection is anything plugged in to the receptacle (i.e. to the right of the CT in the diagram):



A common residential wiring practice in the United States is to “daisy-chain” regular receptacles to a GFCI receptacle where water hazards exist, such that all receptacles powered through the GFCI become part of the GFCI’s protection zone. A bathroom wired this way, for example, provides the exact same degree of ground fault protection at all receptacles in the room. If someone were to plug an electric hair dryer into one of those “daisy-chained” receptacles and then accidentally drop that appliance into a bathtub full of water, the GFCI would trip and cut power to all of the receptacles just as surely as it would trip if the hair dryer had been plugged directly into the GFCI receptacle itself.

Differential current protection is most practical to implement over short physical distances, such as over the phase windings in a generator or some other power system component, but the fundamental concept is applicable over longer distances as well because Kirchhoff's Current Law knows no bounds. Consider for example a transmission line spanning miles of distance between two busses, shown in this single-line diagram:



Here, two differential relays control the tripping of circuit breakers (ANSI/IEEE function 52) at each end of the transmission line. The current at each end of the line is monitored by current transformers connected to local 87 relays, which makes the differential current protection zone cover the entire length of the transmission line. In order for this protection scheme to work, the two local 87 relays must somehow communicate with one another to continuously compare measured current values at both ends of the line. This is accomplished via a communication route between the two relays called a *pilot channel*. The term “pilot” is a general term in the field of protective relaying, referring to any form of data communication. If a significant difference in line current is detected (i.e. resulting from a fault anywhere along the length of the transmission line), both relays trip their respective circuit breakers and thereby de-energize the transmission line.

Pilot systems may take the form of an analog current or voltage “loop” circuit, a microwave radio link, a power-line carrier (PLC)⁵⁰ link, a fiber-optic cable⁵¹ data link, or any other form of point-to-point data link allowing the relays to communicate data with each other. The details of pilot systems in protection schemes is complex and will not be treated in any detail here.

An interesting caveat when applying differential current protection to long lines is that the lines’ capacitive *charging current* may in some cases be substantial enough to trip an 87 relay that is configured too sensitively. One can visualize line-to-ground capacitance as a form of AC “ground fault” because any current taking that path to earth ground is current passing through one CT but not the other.

⁵⁰Power-line carrier, or PLC as it is known in the electric power industry, consists of data communications conveyed over the power line conductors themselves. This usually takes the form of a high-frequency AC signal (in the hundreds of kilohertz range) which is then modulated with the data of interest, similar to radio communication except that the RF signals travel along power lines rather than through empty space as electromagnetic waves. Power-line carrier systems are generally less reliable than fiber optic networks, because the presence of faults on the protected line may compromise the pilot communication.

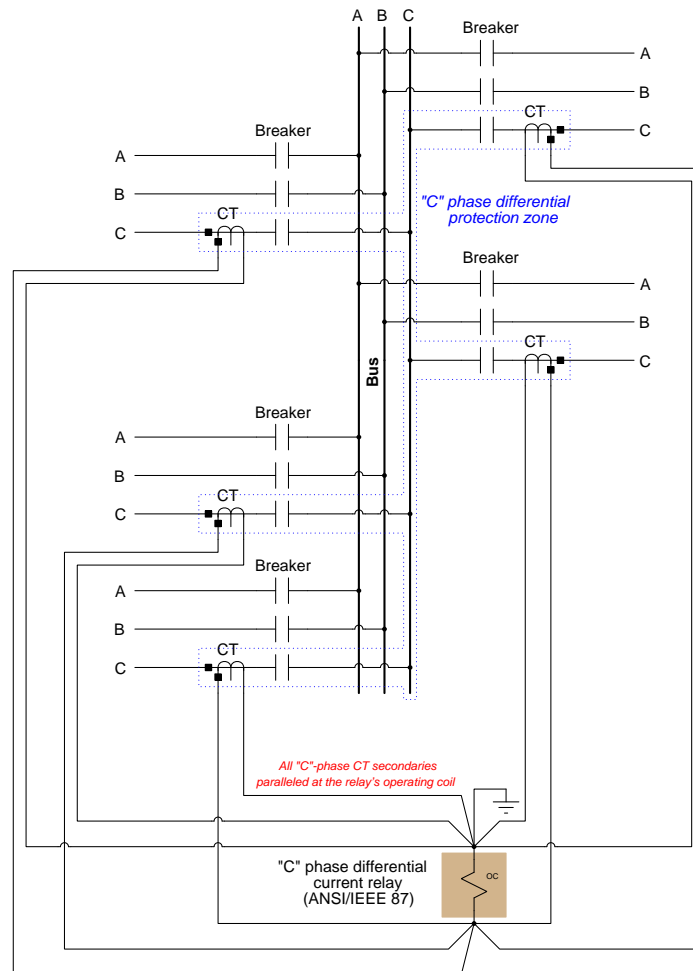
⁵¹Schweitzer Engineering Laboratories manufactures a differential current relay specifically designed for line protection called the model 387L. It is billed as a “zero settings” relay because there are no parameters to configure. Simply set up a pair of 387L’s (one at each end of the line), each one connected to matched CTs monitoring current through all three line conductors, and then link the relays together via a pair of fiber optic cables, and it’s ready to work.

Not only is Kirchoff's Current Law unbounded with regard to distance, it is also unlimited with regard to the number of lines entering or exiting a node. This fact permits us to apply differential current protection to *busses* where multiple power lines and/or devices interconnect. An example of a high-voltage bus photographed at Grand Coulee Dam in Washington state appears here, connecting multiple three-phase transformer banks (each one fed by a hydroelectric generator):




Busses are typically constructed from flexible cable or rigid tube, suspended from ground by insulators. Faults may arise in a bus if an insulator “flashes over” (i.e. develops an electric arc from a bus conductor to ground), or if anything conductive happens to bridge between bus lines. As such, busses may be protected by the differential current principle just like any other electrical component or power line. The algebraic sum of all currents entering and exiting each phase of a bus should equal zero, and if it doesn't it means the bus must be faulted.

A schematic diagram showing one bus with five different feeds reveals how differential current protection may be used to protect a bus with any number of lines. For simplicity's sake the CT and 87 relay wiring is shown only for one phase on this three-phase bus. In any realistic bus differential protection circuit all three phases would be equipped with CTs and there would be three separate 87 "operating coil" elements, one for each phase:




Kirchhoff's Current Law informs us that the algebraic sum of all currents at a node must equal zero. In this case the node in question is the sum of all conductors shown enclosed within the dotted blue protection zone outline. With all CTs possessing the same turns ratio and connected in parallel as shown, their combined secondary currents should all sum to a net value of zero amps through the 87 relay's operate coil during normal operation. However, if a ground fault or a phase-to-phase fault happens to develop anywhere within the protection zone, the CT secondary currents will *not* sum to zero, causing the differential relay to pick up.

Another important concept in protective relaying is *protection zone overlap*. The philosophy here is that each protection zone's size should be limited in order to avoid unnecessarily tripping any more sections of the power system than are necessary to isolate any fault, while leaving no component or conductor unprotected. The following single-line diagram shows how protection zones are configured to overlap each other at each circuit breaker where they connect:

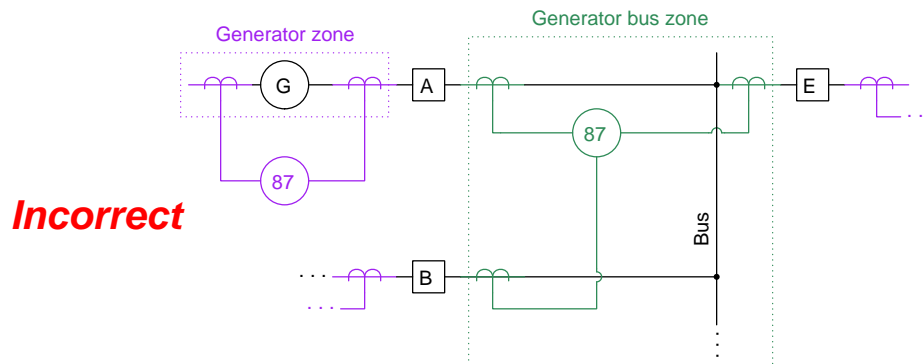


For example, a fault in the upper transmission line belongs to that protection zone only, and will therefore only trip circuit breakers F and G, leaving the other transmission line and associated components to carry power from the generating station to the substation. Note how each circuit breaker in the above system falls within *two* protection zones. If fault happened to develop within breaker F, it would trip the breaker E in the upper generating station transformer zone as well as breaker G in the upper transmission line zone, isolating the failed breaker.

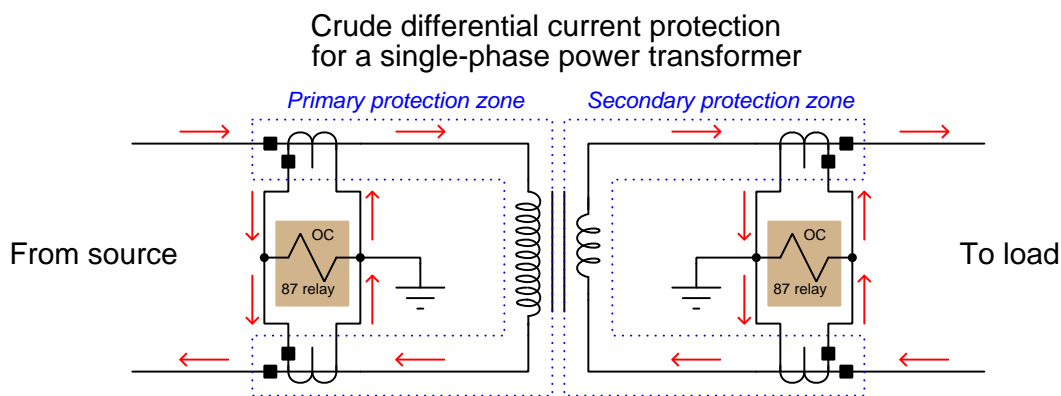
Differential protection zone overlap is accomplished by judicious placement of CTs on either side of a circuit breaker. Recall that the boundary of any differential current protection scheme is defined by the location of the CTs sensing current in and out of the node. Which CT a differential current relay connects to, therefore, defines how far the boundary of that relay's protection zone will reach. We will take a closer look at the single-line diagram in order to explore this concept further, focusing in on the upper-left corner of the generating station and omitting all transformers and all but one generator as well as breakers C, D, and F for simplicity:



Here we see how zone overlap is achieved by connecting each differential relay to the *far* CT on each circuit breaker. If we instead chose to connect each 87 relay to the *near* CT, the two protection zones would not overlap, leaving every circuit breaker unprotected:




Perhaps the most interesting and challenging application of differential current protection is the protection of power transformers, which suffer many of the same vulnerabilities as generators and motors (e.g. winding faults). At first we might be tempted to connect CTs to every conductor entering and exiting a transformer, with 87 relays installed to compare these currents and trip if ever an imbalance were detected, just like protecting the individual windings in a generator. A single-phase transformer suffices to illustrate this concept, again omitting the restraint coils (RC) inside each of the differential relays for simplicity:



So long as each pair of CTs for each differential current relay were matched (i.e. same turns ratio), this protective relay circuitry would detect ground faults and winding-to-winding faults within the power transformer. However, one common transformer fault which would go undetected is a turn-to-turn fault within one of the windings. Such a fault would skew the turns ratio of the power transformer, but it would *not* upset the balance of current going in and coming out of any given winding and therefore would go undetected by the differential relays as shown.

A very clever way to improve differential current protection for a transformer is to have a single 87 relay compare primary and secondary currents for that transformer, thereby extending the zone of protection across both windings with just one relay:



One necessary condition for this strategy to work is to employ CTs with the necessary turns ratios to complement the power transformer's turns ratio and give the 87 relay two equivalent currents to compare. For example, if our power transformer had a turns ratio of 20:1, our two CTs' ratios must differ from each other by the same factor (e.g. a 50:5 CT on the low-current primary winding and a 1000:5 CT on the high-current secondary winding).


This differential current protection scheme works to detect common transformer faults in the following ways:

- **Ground fault:** this kind of fault forces the currents entering and exiting the faulted winding to be unequal. Since the entire winding does not see the same current, it cannot induce the correct proportion of current in the other (non-faulted) winding. This incorrect difference in currents will be seen by the 87 relay.
- **Winding-to-winding fault:** in this kind of fault some of the current from one winding escapes and enters the other winding at a 1:1 ratio. This effectively skews the transformer's step-ratio, which imbalances the currents seen by the 87 relay.
- **Turn-to-turn fault:** this kind of fault directly skews the transformer's step-ratio, which imbalances the currents seen by the 87 relay.

An interesting caveat to using differential current protection on a transformer is the phenomenon of *inrush current* which often happens when a transformer is initially energized. Inrush current happens when the residual magnetism in a transformer's core from its last energized state happens to be substantial, and in the same polarity as the initial magnetization when first energized. The result is that the transformer core begins to magnetically saturate, the result being excess current in the primary winding that does *not* generate current in the secondary winding. Any differential current relay will naturally see this difference as a fault, and may trip power to the transformer unnecessarily.


A clever solution to the problem of false 87 relay tripping due to transformer inrush current is called *harmonic restraint* or *harmonic blocking*. Inrush currents tend to be asymmetrical when viewed on an oscilloscope, due to the bias of a pre-magnetized transformer core (i.e. the core's magnetic field attains stronger peaks in one polarity than the other). This asymmetry results in significant second-harmonic content (e.g. 120 Hz in a 60 Hz power system) in the primary current and is therefore an accurate indicator of inrush. If an 87 relay is designed to detect this harmonic frequency it may be configured to provide additional restraint or even completely inhibit ("block") its own tripping action until such time that the harmonics subside and the transformer stabilizes to normal operation.

Differential current protection of *three-phase* transformers and transformer banks is a more complicated matter, and not simply because there are three of everything. Power transformers are often wired with their primary and secondary sides in different configurations (e.g. Wye-Delta or Delta-Wye). Thus, the currents entering and exiting a power transformer may not be in-phase with each other, and in such cases cannot be compared directly against each other for differential current protection. Consider this example, where the primary winding is a Wye and the secondary winding is a Delta. For simplicity's sake we will consider a transformer with equal numbers of turns on every winding, such that each primary/secondary coil pair has a 1:1 turns ratio. Furthermore, we will label each of the primary phase currents as I_A , I_B , and I_C :



Since the secondary windings are Delta-connected, the secondary lines carry currents equal to $I_A - I_C$, $I_B - I_A$, and $I_C - I_B$, respectively as declared by Kirchhoff's Current Law at each of the Delta winding nodes. The result is that each secondary line current is $\sqrt{3}$ times larger and lags 30° behind each corresponding primary line current, as shown by the phasor diagrams. This 30° phase shift means we cannot simply connect CT pairs together to a common 87 relay as we could in the single-phase transformer example. In order to compensate for the 30° phase shift imparted by the power transformer, we must connect the CTs themselves in a complementary Delta-Wye configuration such that the 87 relays will be able to compare in-phase currents from primary and secondary sides of the power transformer.

In this schematic diagram we see how primary and secondary CTs need to be connected (CTs on the Wye side of the power transformer are Delta-connected, while CTs on the Delta side of the transformer are Wye-connected) to provide a matching 30° phase shift. The currents generated by each CT secondary winding are labeled with lower-case letters (i rather than I) in order to represent their smaller values:



Note how each current entering an 87 relay’s restraint coil (RC) exits out the other restraint coil with the same mathematical expression, indicating equal current values. This will be true so long as all CT ratios are correct and currents into and out of the power transformer correspond properly to each other.

If the power transformer windings happen to have 1:1 turns ratios as is the case in this demonstration circuit, the secondary line currents will be larger than the primary line currents by a factor of $\sqrt{3}$, owing to the fact that the primary windings are Wye-connected (winding currents same as line currents) while the secondary windings are Delta-connected (winding currents combine to make larger line currents). This means each of the secondary CTs will see a greater line current than each of the corresponding primary CTs. However, given the fact that the CTs on the primary side of the power transformer have their secondary windings Delta-connected, the actual amount of current they send to the 87 relay coils will be the same as the amount of current sent to the 87 relay by the other CTs, given equal CT ratios all around.

If the power transformer windings have turns ratios other than 1:1, the CTs installed on the primary and secondary lines will likely have differing ratios as well. It is unlikely that the CTs will exhibit precisely complementary ratios to the power transformer’s internal winding ratios, which means when these CTs are connected to 87 relays their output currents will *not* match in magnitude. Legacy electromechanical 87 relays were equipped with “taps” which could be set at different ratios to equalize the CT currents to within a few percent of agreement with each other. Modern digital 87 relays are able to do a much better job of matching primary-side and secondary-side CT outputs because they may be programmed with arbitrary correction factors. As you can see, care must be taken when connecting transformer CTs to differential current relays in order to ensure primary and secondary current values match in phase and magnitude.

It should be noted that the 30 degree phase shift between primary and secondary windings of the power transformer previously shown is actually a standard specified by the IEEE. The IEEE standard C57.12.00-2010 (“IEEE Standard for General Requirements for Liquid-Immersed Distribution, Power, and Regulating Transformers”) states that transformers having wye-wye or delta-delta winding configurations shall exhibit 0° phase shift from input to output, but transformers having wye-delta or delta-wye winding configurations shall exhibit 30° phase shift between primary and secondary sides with the lower-voltage side of the transformer lagging.

Modern digital 87 relays offer “CT compensation” which may be used in lieu of complementary connections to correct for the phase shift of a Wye-Delta power transformer, as well as correct for CT ratios that are not ideally matched. Rather than carefully connect the secondary windings of all CTs in such a manner that the primary- and secondary-side phase angles and current values match for all normal transformer operating conditions, we may connect the CTs as we see fit (typically in a Wye configuration on both⁵² sides, for simplicity) and let the relay mathematically match angles

⁵²There is a potential problem arising from CT secondaries in Wye when those CTs are measuring currents on the Wye-connected side of a power transformer, and that is the problem of *zero sequence* currents. A “zero sequence” set of currents is equivalent to in-phase currents flowing through all three lines of a three-phase power system, lacking the normal 120 degree phase shift from each other. The mathematical foundations of this concept are beyond the immediate scope of this section (for more information, refer to section 5.8.4 on “Symmetrical Components” beginning on page 457), but suffice to say zero-sequence currents are found in certain fault conditions as well as circuits containing “triplen” harmonics (i.e. harmonic frequencies that are some multiple of $3\times$ the fundamental, e.g. 180 Hz, 240 Hz, 540 Hz for a 60 Hz power system). Zero-sequence currents flow through the neutral conductor in a 4-wire Wye-connected system, but circle through the phase elements of a Delta-connected system. This means a Wye-Delta connected

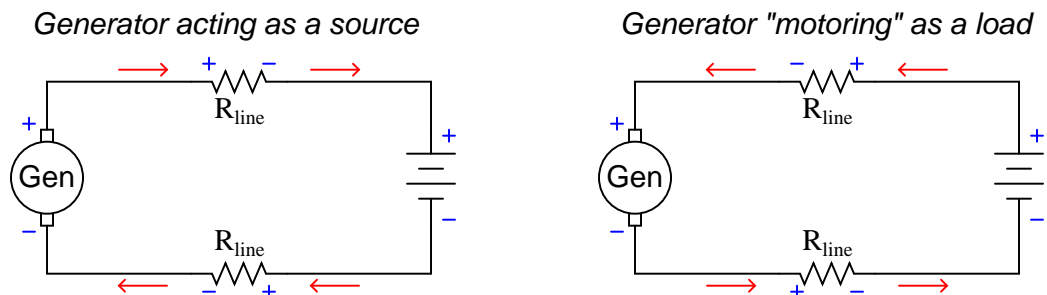
and magnitudes. This digital alternative, of course, requires careful attention to relay settings in order to work.

25.11 Directional overcurrent (67) protection

While 50 and 51 (instantaneous and time overcurrent) relay functions monitor line current magnitude and guard against excesses, there are applications where the *direction* of line current is just as relevant as the magnitude. In such cases, we need a protective relay function able to discriminate between current in one direction versus current in the other direction. The ANSI/IEEE number code designation for a directional current-sensing protection is 67.

One such application is *generator protection*, where an overcurrent relay monitors the amount of current at the point where an electrical power generator connects to a larger network of generators. The problem of directional current monitoring is easiest to understand in the context of a direct-current (DC) generator and battery circuit, which we will now explore as an introduction to the topic:


Consider a DC generator connected to a secondary-cell (i.e. rechargeable) battery. Here, the voltage polarity never changes, but the direction of current does change depending on whether the generator is acting as a power *source* (charging the battery) or “motoring” and acting as a power *load* (discharging the battery):



A generator acting as a source (in this case, to charge the battery) is fulfilling its intended function. A generator running as a motor, drawing energy from the battery as a load, is most definitely not fulfilling its intended function. Therefore, we would consider any current in the wrong (generator as load) direction to be excessive, while considerable current in the correct (generator as source) direction would be considered perfectly normal. If we were to install an overcurrent relay in this simple DC system, we would therefore prefer that it be more sensitive to current (i.e. pick up at a lower value) in the “reverse” direction than to current in the “forward” direction.

transformer where a fourth conductor attaches to the center of the Wye winding set may experience line currents on the Wye side that are not seen in the line conductors of the Delta side, and may therefore cause a differential current relay to operate. This is another reason why connecting CTs differently than the power transformer windings they sense (i.e. Delta-connected CTs on a power transformer’s Wye side) is a good idea: any zero-sequence currents within the power transformer’s Wye-connected winding will circulate harmlessly through the Delta-connected CT secondaries and never enter the 87 relay. For digitally compensated 87 relay installations where all CTs are Wye-connected, the relay must also be configured to mathematically cancel out any zero-sequence currents on the Wye-connected side of the power transformer.

Alternating-current (AC) power systems are not that different in this regard. In the forward-current direction, the generator acts as a power source, sending electrical power to any loads connected to the generator bus. In the reverse-current direction, the current experiences a phase-shift of 180 degrees from that of the forward-current direction, at which point the generator acts as a load to any other generator(s) on the network. This phase shift is evident if we compare the signal waveforms from PT and CT instrument transformers connected to the generator:



The fundamental problem we face in designing an AC directional current relay is how to detect this phase shift between forward and reverse current. In our DC generator circuit, a reverse flow of current could easily be detected by monitoring the polarity of voltage drop across a series resistance⁵³. In an AC circuit, however, the only way to tell if the line current is going the wrong way is if we compare the current waveform against another “reference” waveform (such as line voltage). The difference in phase shift between forward current and reverse current will be 180 degrees. Thus, an AC directional protective relay requires at least *two* signal inputs: one representing line current to be monitored, and another serving as a *polarizing* or *reference* quantity to be used for phase comparison.

This polarizing quantity may be line voltage, it may be a different current in the system, or it may even be a some combination where one signal provides backup in case the other polarizing signal becomes too weak. The challenge of finding a suitable polarizing signal in a power system for a directional relay stems from the fact that voltage and current signal strengths may vary wildly under fault conditions, which is precisely when we need the protective relay to do its job. Consider, for example, using generator line voltage as the polarizing signal to be compared with line current in the 67 relay. Imagine now if that generator suffers a major fault in its windings. Any other generators connected to the same bus will now send power into the faulted generator: a clear case of reverse power flow (into the generator) when we need the directional relay to trip. However, if the fault happens to significantly reduce the line voltage of the failed generator, the directional relay may receive too weak of a polarizing signal to properly operate, and thus may fail to trip the generator’s breaker connecting the failed generator to the bus.

Modern microprocessor-based directional relays have a definite advantage in this regard over legacy electromechanical relay designs, in being able to intelligently select the best polarizing quantity to use during fault conditions. Relays manufactured by Schweitzer Engineering Laboratories having directional protection elements for ground and neutral currents, for example, use a proprietary algorithm called “Best-Choice Ground Directional Element” logic to select from one

⁵³Note the reversal of polarity for the voltage drop across each line resistance in the DC example diagram. A shunt resistor intentionally placed in series with the generator current could fulfill that same directional-sensing role.

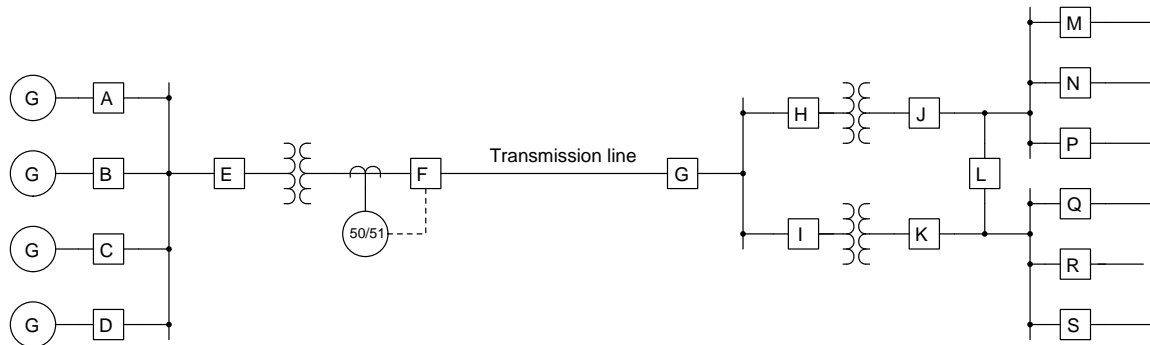
of several real and calculated polarizing quantities (e.g. neutral current, zero-sequence impedance, negative-sequence impedance). Microprocessor-based protective relays are able to calculate symmetrical component quantities (positive-sequence, negative-sequence, and zero-sequence) from live measurements, and then use those calculated quantities in protection logic just the same as raw voltage and current measurements. Given the fact that negative-sequence quantities manifest only during asymmetrical fault conditions, and that asymmetrical faults are more common than symmetrical faults, the ability to calculate these quantities during fault conditions and use them to make trip decisions is a powerful advantage indeed.

25.12 Distance (21) protection

A form of protection against faults on long-distance power lines is called *distance* relaying, so named because it is actually able to estimate the physical distance between the relay's sensing transformers (PTs and CTs) and the location of the fault. In this way, it is a more sophisticated form of fault detection than simple overcurrent (e.g. 50 or 51 relay). The ANSI/IEEE number code designation for distance relaying is *21*.

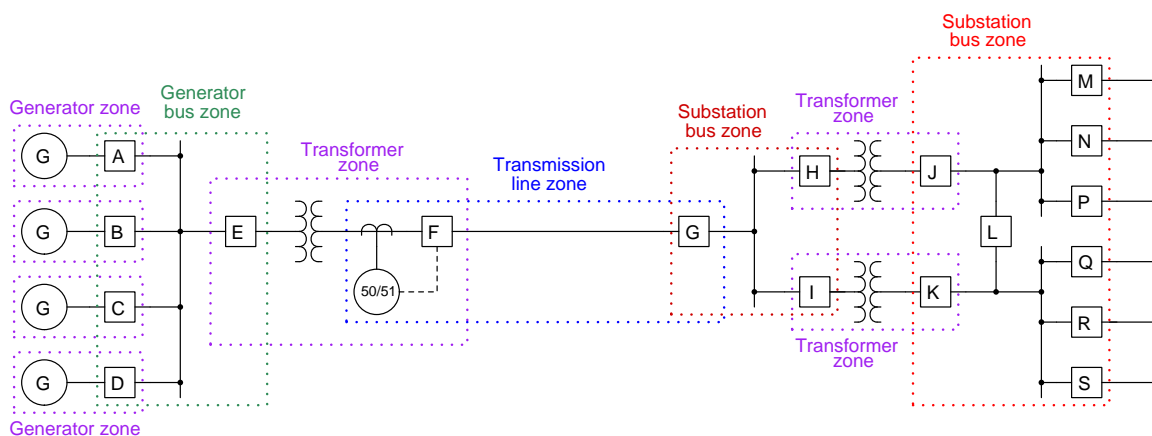
25.12.1 Zone overreach and underreach

In order to understand the rationale for distance relaying on transmission lines, it is useful to recognize the limitations of simple overcurrent (50/51) protection. Consider this single-line diagram of a transmission line bringing power from a set of bus-connected generators to a substation at some remote distance. For simplicity's sake, only one protective relay is shown in this diagram, and that is for breaker "F" feeding the transmission line from the generator bus:



The purpose of the overcurrent relay tripping breaker "F" is to protect the transmission line and associated equipment from damage due to overcurrent in the event of a fault along that line, and so the relay must be set appropriately for the task. The amount of fault current this relay will see depends on several factors, one of them being the location of the fault along the transmission line. If we imagine a fault occurring on the line near breaker "F," the fault current will be relatively high because it is close to the generator bus and therefore experiences little transmission line impedance to limit current. Conversely, if we imagine a fault farther out on the transmission line (closer to breaker "G"), the amount of current caused by the fault will be less, even for the exact same type of fault, simply due to the added series impedance of the transmission line's length (if $I = \frac{V}{Z}$ and Z increases while line voltage V remains the same, I must decrease). Any similar fault further downstream of the generators – such as a fault in one of the transformers in the substation – will draw even less current through breaker "F" than a similar fault on the transmission line for the same reason of greater series impedance.

An important concept in protective relaying is that of *protection zones*. Protective relays exist to protect the power system from damage due to faults, and they do so by tripping circuit breakers to interrupt⁵⁴ the flow of power to a fault. However, in the interest⁵⁵ of maintaining power to customers it is best for protective relays to *only* trip those breakers necessary to clear the fault, and no more. Thus, protective relays are designed to trip specific breakers to protect limited “zones” within the system. In this next single-line diagram, we show the same system with rectangular zones overlaid on the system components:



This zone diagram makes it clear that breaker “F” and its associated overcurrent relay should only act to protect the transmission line from fault-induced damage. If a fault happens to occur within one of the transformer zones within the substation, we would prefer that fault be cleared by the protective relays and breakers for that transformer alone (i.e. either breakers “H” and “J” or breakers “I” and “K” depending on which transformer faults), in order that power be maintained in the rest of the system. This means the overcurrent relay controlling breaker “F” needs to be sensitive to faults within the transmission line zone, but insensitive to faults lying outside of the transmission line zone. If the 50/51 relay happened to *overreach* its zone and trip breaker “F” because of current sensed from a fault in one of the substation transformers, it would unnecessarily cut power to the entire substation, with a loss of power to all loads fed by that substation.

At first, the problem of overreaching may seem simple to solve: just calculate the maximum fault current in the transmission line due to any worst-case fault outside of that zone, and be sure to set the overcurrent relay so that it will only trip at some current *greater* than that amount, or set it so it will trip after a longer time delay than the substation relay(s) will trip, to give the substation relays a chance to clear the fault first.

The weakness of this approach is that fault location is not the only factor influencing fault current magnitude. Another important variable is the number of generators in service at the time of the fault. If one or more of the generators happens to go off-line, it reduces the generator bus’s ability to

⁵⁴In the electric power industry, the probability that protective relays and associated equipment will reliably interrupt power in the event of a fault is called *dependability*.

⁵⁵In the electric power industry, the probability that protective relays and associated equipment will not interrupt power unnecessarily is called *security*. As one might guess, dependability and security are two competing interests in the design of any protection scheme, the challenge being how to strike a reasonable balance between the two.

supply current to a fault. *Another way of saying this is that the power source's impedance changes with the number of generators on-line.* This means any given fault downstream of breaker "F" will cause less fault current than it would if all generators were on-line.

This causes a problem for the "reach" of the overcurrent relay controlling breaker "F." With reduced current capacity from the generator bus, the same relay setting that worked well to protect the transmission line zone will now be too high for faults lying toward the far end of that line. In other words, the overcurrent relay may *underreach* and fail to trip breaker "F" because the amount of fault current for a transmission line fault is now less than what the relay has been set to protect against, and all because we happen to have fewer generators on-line to supply power. The impedance of the transmission line and fault may be precisely the same as before, but the overcurrent relay will not trip because the circuit's *total* impedance has changed due to fewer generators being on-line.


We see that the location of a fault within a long-distance power distribution system cannot be reliably detected by sensing current alone. In order to provide more consistent and reliable zone protection for the transmission line, we need a form of protection better able to discriminate fault location. One such method is to measure the impedance of the protected zone, based on current *and* voltage measurements at the entry point of power into that zone. This is the fundamental concept of *distance* protection: calculating the impedance of just the protected zone, and acting to trip breakers feeding power to that zone if the impedance suggests a fault within the boundaries of that zone.

25.12.2 Line impedance characteristics

Capacitance, inductance, and resistance are all naturally present along miles of power line conductors: capacitance due to electric fields existing within the separation of the lines from one another and from earth ground by the dielectric of porcelain insulators and air; inductance due to the magnetic fields surrounding the lines as they carry current; and resistance from the metal conductors' length.

The capacitive nature of a power line is evident when that line is open-circuited (i.e. no load connected). For the next few schematic diagrams, only a single phase (one "hot" conductor and one "neutral" conductor) will be represented for the sake of simplicity:

Load disconnected -- line capacitance dominates circuit impedance




Here, an oscilloscope shows the relative magnitudes and phase shifts of the voltage and current waveforms, allowing us to make determinations of total circuit impedance ($Z = \frac{V}{I}$).

Under typical load conditions, the resistance of the load draws a much greater amount of current than an open-circuited line draws due to its own capacitance. More importantly, this current is nearly in-phase with the voltage because the load resistance dominates circuit impedance, being substantially greater than the series reactance caused by line inductance while being substantially less than the parallel capacitive reactance:

Load connected -- load resistance dominates circuit impedance

$$X_{C(\text{line})} \gg R_{\text{load}} \gg X_{L(\text{line})}$$




Current and voltage nearly in-phase

Z_{total} = moderate value at nearly 0° phase angle

A significant fault behaves like a very low resistance connected in parallel. This not only decreases total circuit impedance but also shifts the phase angle closer toward $+90^\circ$ because now the line inductive reactance is substantial compared to the resistance of the fault. Real transmission lines tend to exhibit shorted impedance phase angles nearer 70 degrees rather than 90 degrees, owing to the effects of line resistance. The exact line impedance phase angle depends on conductor size and separation:

Line fault -- line inductive reactance dominates circuit impedance

$$R_{\text{fault}} \ll X_{L(\text{line})}$$




Current lags voltage by approx. 70°

$$Z_{\text{total}} = \text{low value at approx. } 70^\circ \text{ phase angle}$$

Since line inductance is a fairly linear function of line distance (a longer power line means more inductance, given a fixed inductance-per-mile value), and this inductive reactance is the dominant factor limiting fault current, the magnitude of the fault current becomes an approximate indication of distance between the instrument transformers and the fault.

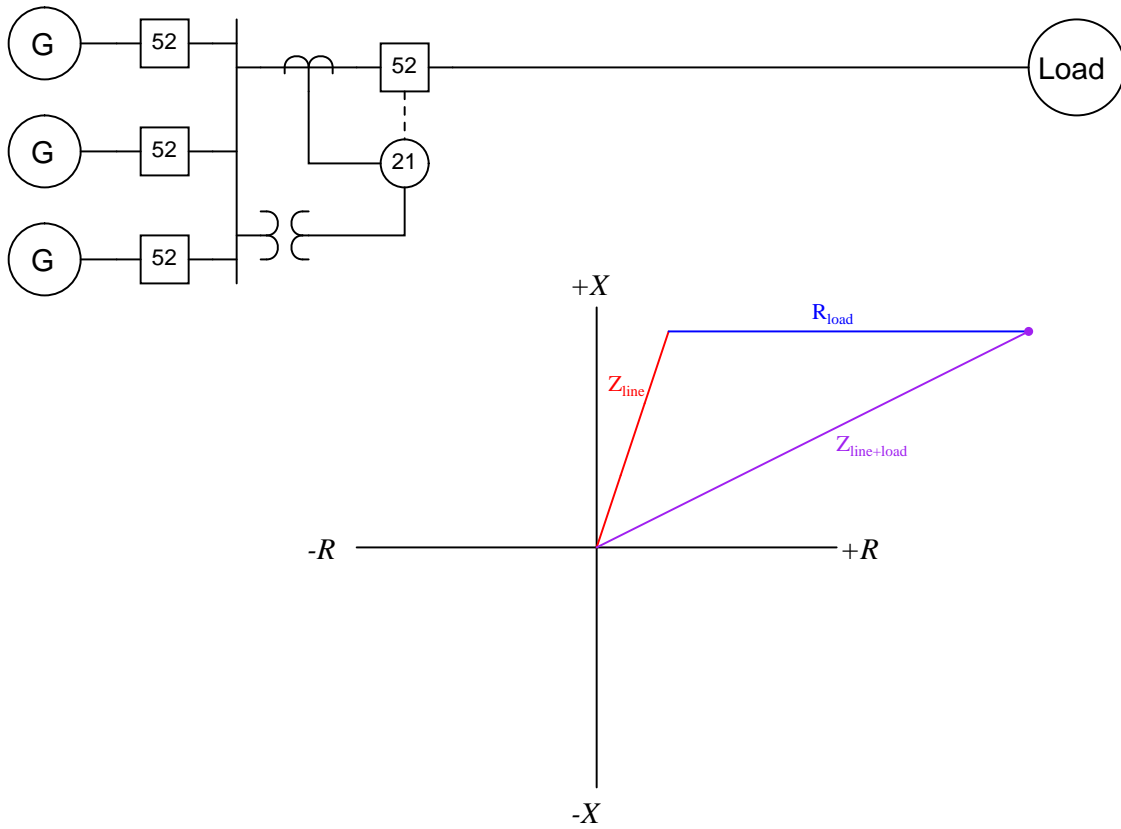
25.12.3 Using impedance diagrams to characterize faults

Oscilloscope displays showing the raw voltage and current waveforms are clumsy representations of line impedance. Better visual representations for impedance exist, the most popular being a *phasor diagram* for line impedance with resistance (R) on the horizontal axis and reactance (X) on the vertical axis, commonly referred to as an *R-X diagram*. The three line examples shown in the previous section using the oscilloscope are shown in phasor format here:




Keep in mind that these phasors represent *impedance*, and as such a short-circuited (faulted) condition is shown as a short phasor, while an unloaded condition is shown as a long phasor. It should also be noted that these impedances, while calculated from measurements of voltage and current, do not change unless the line, load, or fault characteristics change. If the system voltage were to sag due to a generator problem, for example, the impedance phasor representing the combined effects of line and load impedance would not be altered. Any protective relay operating on impedance would therefore ignore such changes, and trip only if the line's characteristics were to change. This is precisely the behavior we need from a "distance" relay, enabling it to discriminate line faults better than a simple overcurrent relay ever could.

For a normal load condition, the impedance phasor will be significantly longer than that of the line's full length (i.e. much higher impedance) with an angle significantly less than that of the line impedance alone:

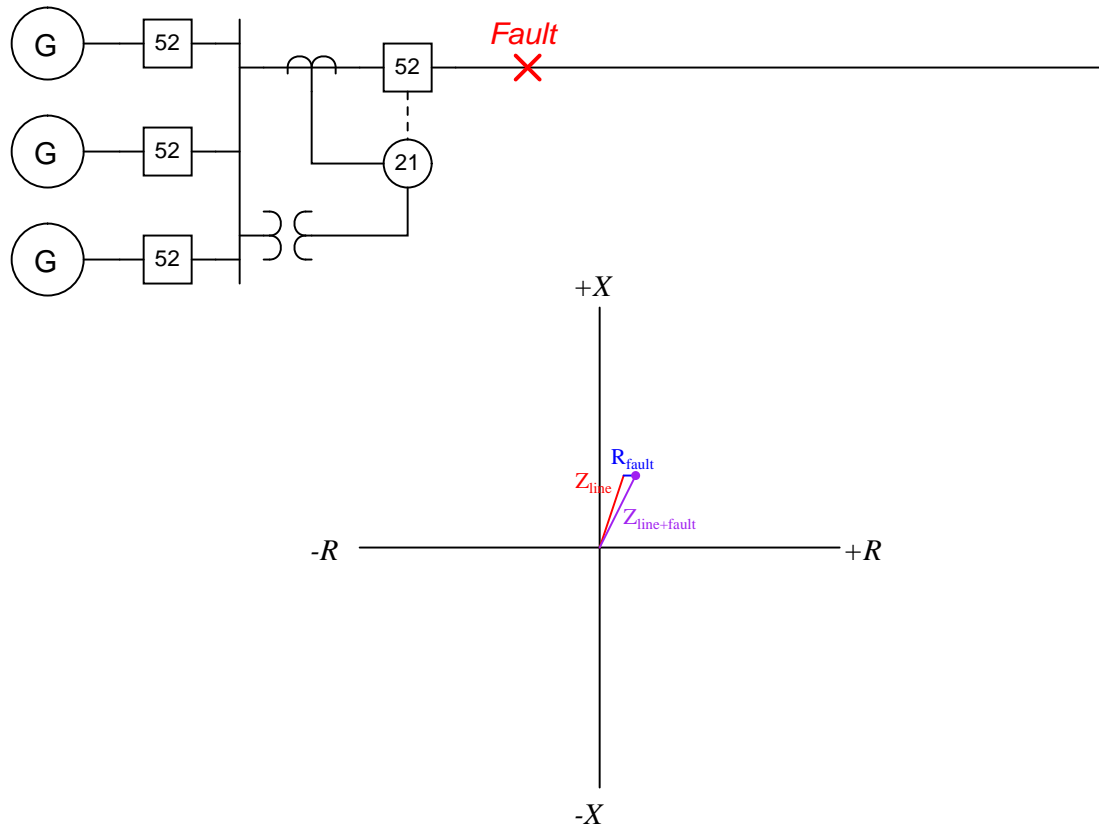


Short-circuit faults at various locations along a transmission line will cause the impedance phasor to vary primarily in magnitude and angle. Recall that during fault conditions, the resistance and reactance of the power line itself is the dominant impedance limiting fault current. The actual fault is predominantly resistive, with a very small impedance value.

For a fault far removed from the relay, the impedance phasor will be long (i.e. relatively high impedance) with angle nearly equal to that of the line impedance alone:




For a fault closer to the relay, the impedance phasor will be short (i.e. low impedance) with angle slightly less than that of the line impedance alone:



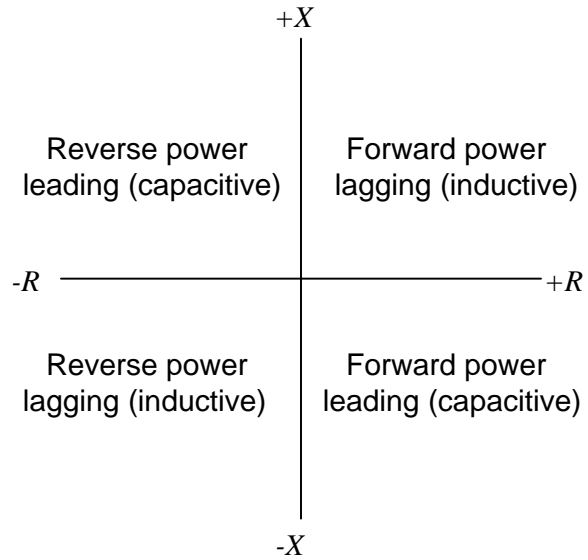
The goal of a distance relay (ANSI/IEEE code 21) is to trip its circuit breaker(s) if a fault occurs within its programmed “reach” and to ignore both normal operating loads and faults lying outside its reach.

If additional sources of electrical power are connected to the far end of the transmission line, it is possible for the distance relay to sense *reverse* power flow. Consider a case where a short-circuit fault occurs on the generator bus shown in this single-line diagram:




A fault to the left of the distance relay manifests as high current and low voltage just like a fault on the transmission line, but since the current waveform is inverted (180° phase shift) due to the opposite direction of fault current, the impedance phasor ends up in an entirely different quadrant of the R-X diagram. If the goal of the distance relay is to protect the transmission line, we need it to ignore such faults, because to operate on such a fault would be an example of overreach, the distance relay “reaching into” the generator bus zone where it should be concerned with the transmission line zone.

Each of the R-X diagram's quadrants may be labeled in terms of power direction and power factor, either "lagging" (predominantly inductive) or "leading" (predominantly capacitive):



25.12.4 Distance relay characteristics


A primitive electromechanical impedance relay design for detecting faults along long-distance transmission lines uses a simple balance-beam mechanism to sense when the ratio of line current to line voltage ($\frac{I}{V}$) becomes excessive. It will trip if ever the impedance became too small (i.e. too much I and too little V):



This relay's operating coil (OC) is energized by a current transformer signal, while the restraint coil (RC) is energized by a potential transformer signal. During normal and unloaded operation, when voltage is high and current is moderate-to-low, the attractive force exerted on the beam by the restraint coil's magnetic field exceeds the attractive force exerted on the beam by the operating coil's magnetic field, and the trip contact remains open. If, however, a fault condition occurs on the transmission line, current will dramatically increase while voltage decreases. This combination of changes causes the operating coil's magnetic attraction to exceed that of the restraint coil, causing the beam to tilt in a clockwise rotation, closing the trip contact to send a DC tripping signal to the circuit breaker's trip coil.

An important characteristic of this crude impedance relay design is that it is insensitive to phase shift between voltage and current. In other words, it does not discriminate between line impedance values having different phase angles, but will trip based solely on a minimum impedance *magnitude*.


The “reach” of this impedance relay may be plotted on an R-X phasor diagram as a circle centered around the origin of the graph:



Any line condition placing the impedance phasor tip within this circle will cause the relay to trip. Any line condition placing the impedance phasor tip outside this circle will cause the relay to be restrained (i.e. it will not trip). Thus, the reach of this relay is represented by the radius of the circle sketched on the R-X diagram.

As you can see, this design of distance relay will trip for reverse-power conditions just as easily as it will trip for forward-power conditions. Since we generally wish to de-sensitize distance relays from “reaching backward” into a reverse protection zone, we must find some way to limit the impedance relay’s tripping sensitivity in the reverse direction.

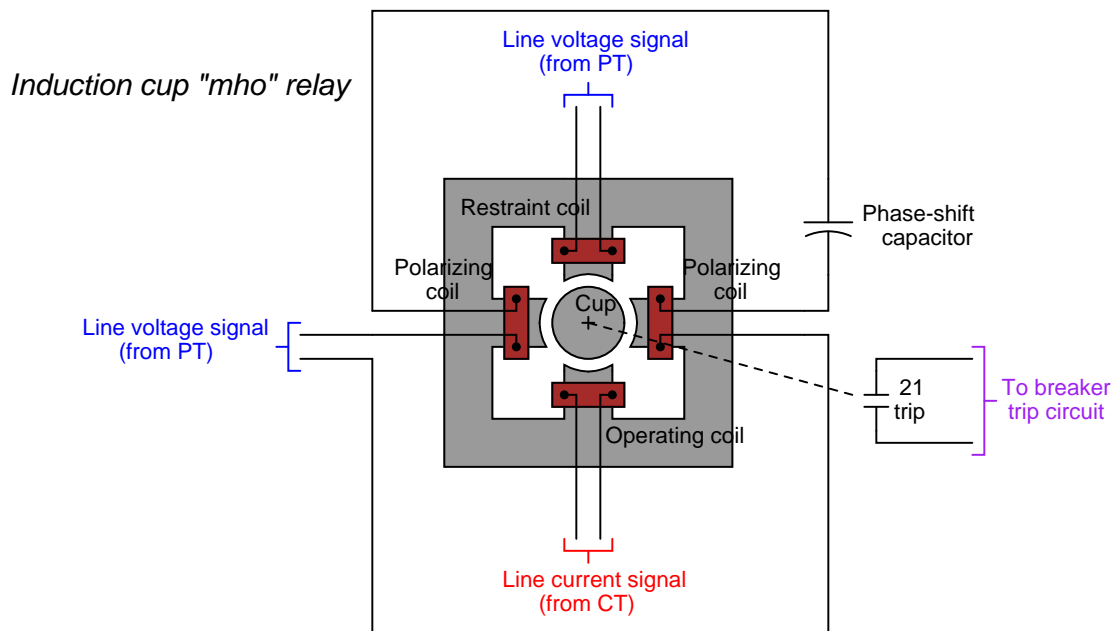
For the primitive balance-beam design, one solution to the problem of reverse-power sensitivity is to use a directional relay in conjunction with the distance relay to *block* the distance relay's action during reverse-power conditions. The trip contact of a directional relay could be wired in series with the trip contact of the impedance relay, such that the only way to trip the breaker is if both the impedance relay and the directional relay agree. We may represent this blocking behavior by drawing a line called a *blinder* on the R-X diagram showing a threshold beyond which the impedance relay cannot operate:



Phase-shifting components inside the directional relay mechanism tilt its blinder characteristic slightly clockwise from its natural (horizontal) characteristic. As you can see, the blinder prevents all operation in the lower-left quadrant, restricting operation of the impedance relay primarily to the upper-right and upper-left quadrants, with only a small portion of the lower-right quadrant active.

Blocking the impedance relay's action using a directional relay is a crude solution for a crude relay design. Much better distance relay characteristics have been developed since.

A major breakthrough in distance relay design came with the advent of the *induction cup* mechanism. This is similar in design to the *induction disk* mechanism explained in the section on time-overcurrent relays, but designed to operate very quickly rather than very slowly. An induction cup mechanism closely resembles a two-phase induction motor, where a small cup-shaped metal rotor is surrounded by two sets of electromagnet poles. Maximum torque will be induced on the rotor when the stators' magnetic fields are 90 degrees phase-shifted from one another in time. When a positive torque is applied to the cup, it rotates on its axis to close a trip contact, sending DC power to the circuit breaker's trip coil:



The amount of torque induced on the rotor (cup) is described by the following formula:

$$\tau = VI \cos(\theta - \phi) - KV^2$$

Where,

- τ = Torque exerted on induction cup
- V = Line voltage
- I = Line current
- θ = Phase angle of voltage with respect to current
- ϕ = Phase angle of maximum rotor torque
- K = Restraint constant of relay

Algebraically solving for the relay constant K at a point of zero torque (the pick-up value for the relay) yields units of *mho*, or inverse ohms, which is why this mechanism is called a “mho element”:

$$0 = VI \cos(\theta - \phi) - KV^2$$


$$0 = \frac{I}{V} \cos(\theta - \phi) - K$$

$$K = \frac{I}{V} \cos(\theta - \phi)$$

$$K = \left[\frac{\text{Amps}}{\text{Volts}} \right] \text{ [unitless]}$$


$$K = [\text{Mhos}]$$

With no phase-shifting capacitor, this mechanism will be maximally sensitive to impedance values of $+90^\circ$, with a characteristic resembling a circle passing through the origin of an R-X diagram:




This relay’s reach is defined as any impedance falling within the circle, just as we defined the reach of the impedance relay. The difference here, of course, is that the “mho” distance relay is entirely insensitive to conditions within the lower quadrants of the R-X diagram.

With the addition of the phase-shifting capacitor to the induction cup polarizing coil circuit, the circular characteristic becomes tilted. Ideally, the angle of this tilt is set to match the impedance phasor angle of the transmission line so as to make the relay maximally sensitive to faults along the line.



With a tilted axis, the longest chord within the circle beginning at the origin of the R-X diagram is one matching the axis of tilt. Therefore, the highest impedance value capable of operating the relay and tripping the circuit breaker is one where the phase angle matches the tilt: indicative of a low-resistance fault at the end of a transmission line, assuming the circle's diameter is proportional to the length of that line. Measured impedances at any other angle must be lower (i.e. a "heavier" loading condition) in order to operate the distance relay and trip the breaker.


If we compare circle characteristics for the simple impedance relay versus the “mho” relay capable of tripping at the same end-of-line fault condition, we see a remarkable contrast:




The "mho" relay reaches farthest only along the transmission line, while the impedance relay reaches the same distance indiscriminately

Both relays have the exact same reach at the transmission line's impedance angle, but the impedance relay's reach extends omnidirectionally for all phase angles and power flow directions, while the mho relay's reach is optimized for the forward power direction and the line's impedance, making it far more selective to faults along that line.

In some applications it is desirable to have the distance relay sensitized to certain values of reverse impedance (i.e. the lower-left quadrant on the R-X diagram). The induction cup relay mechanism is capable of having its circular reach characteristic “offset” with additional components so that the circle covers part of every quadrant like this:



Another variation on the “mho” characteristic is to equip the distance relay with multiple elements, each one with a different reach. The purpose of this is to provide *backup* protection for other zones by allowing the distance relay to overreach its primary protection zone:




Since distance relays trip whenever the tip of the impedance phasor falls within the prescribed area on the R-X diagram, at first it may seem as though zones 1 and 2 are pointlessly redundant to zone 3, since any fault lying within one of the inner zones will certainly be within the reach of the furthest zone. Indeed, this would be the case if all three distance elements operated at the same speed. However, if the zone 2 reach is purposely delayed in its action to be slower than zone 1, and zone 3 purposely delayed to make it slower than zone 2, the distance relay will serve to provide *remote backup* protection for the substation bus and transformer zones in the event the protective relays and/or breakers for those zones fail to properly clear a fault.

The following photograph shows a pair of Westinghouse electromechanical distance relays mounted next to a pair of time-delay units, each timer providing a different amount of delay for each of the two zones of protection afforded by the distance relays:




In the age of microprocessor relays, this design philosophy of physically wiring time-delay relays to the outputs of electromechanical protective relays for multi-zone protection may seem archaic, but it represents standard practice in distance relaying for a number of decades. Of course, modern microprocessor-based distance relays are able to perform all the necessary zone timing functions along with distance-sensing fault detection within the same unit, and do so with a degree of precision unthinkable with electromechanical relays.

Modern microprocessor-based distance relays are not constrained to circular reach characteristics either, and thus may be programmed to implement a variety of interesting reach functions. Although the traditional “mho” characteristic is still available as an option within digital distance relays, another form of reach called the *quadrilateral* (or “quad” for short) is often provided, whereby the distance relay may be configured to trip for any impedance phasor lying within a four-sided boundary:



Another common option within digital distance relays is the *load encroachment* blocking characteristic, designed to prevent relay tripping on heavy (legitimate) loads that might otherwise fall within the reach of the distance relay's characteristic:



25.13 Auxiliary and lockout (86) relays

An important type of “accessory” relay, especially for legacy electromechanical protective relays, is the so-called *auxiliary* or *lockout* relay, designated by the ANSI/IEEE number code 86. The purpose of an 86 relay is to serve as an intermediary element between one or more protective relays and one or more control devices, both expanding the number of control elements actuated by any one protective function and also providing a “latching” function which must be intentionally reset in order to resume normal operation of the system.

Lockout relays are typically panel-mounted devices equipped with handles for manual resetting. This photograph shows a lockout relay actuated by a breaker fault (BF) function, requiring manual intervention before the system may be returned to service:



This lockout relay is shown in the “Reset” position, having been placed in that state by someone rotating the handle clockwise until it points vertically. When a remote signal energizes this lockout relay’s trip coil, it moves by spring action into the “Trip” position, the handle turning about 45 degrees counter-clockwise and an orange “target” appearing just above the handle shaft. The lockout relay will remain in this “tripped” state until returned to its “Reset” position by human intervention.

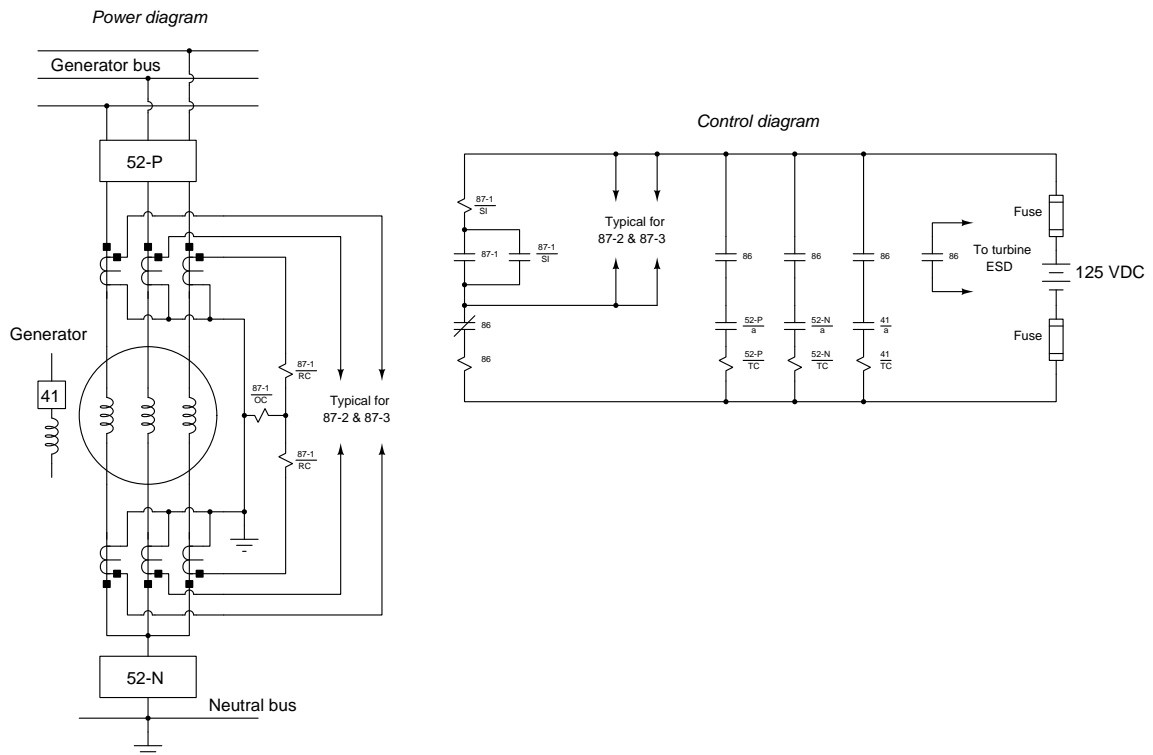
The following photographs show what a lockout relay looks like when tripped (left) and reset (right). The relay's trip coil may be seen in both photographs as a diagonally oriented metal component at the bottom:



Multiple sets of switch contacts operated by the common shaft inside the lockout relay may then be used for a variety of functions, including sending trip signals to circuit breakers, activating sirens, energizing indicator lamps, sending status signals to other protective relays, inhibiting (blocking) other relay functions, etc.

When interpreting switch states in a latching relay such as this, the “reset” state is considered the “normal” state (i.e. the state where the relay is “at rest”). The actuating stimulus for an 86 relay is energization of its trip coil, and so the “tripped” state of a lockout relay is its “actuated” state where all of its contacts transition to their non-resting states (e.g. NO contacts closing and NC contacts opening).

A typical application for an 86 lockout relay is in this generator protection system, where a set of three differential current relays (87) sense any imbalance in currents entering and exiting the generator's stator windings. This is a case where the detection of current imbalance needs to trip multiple circuit breakers and signal a steam turbine control system in order to protect the generator from possible damage:



If any of the differential (87) relays pick up, the lockout (86) relay's trip coil will become energized through one or more of the 87 contacts. This in turn causes the lockout relay to switch to its tripped state by spring tension, opening the NC contact to cut power to the 86 trip coil (and unlatch the 87 relay seal-in circuits) and closing all four NO 86 contacts. Three of these closed contacts then send DC station power to the trip coils of three different circuit breakers: the main phase breaker trip coil (52-P/TC), the neutral breaker trip coil (52-N/TC), and the trip coil of the breaker controlling DC power to the generator's field winding (41/TC). Another NO 86 contact sends a discrete shutdown signal to the steam turbine's emergency shutdown system (ESD) to halt the flow of steam to the turbine and thereby prevent over-speeding.

The latching nature of the 86 relay means none of these circuit breakers may be closed until someone turns the handle of the 86 relay back to the "reset" position. This latching action "locks out" the generator and prevents any remote attempt to place it back on-line until someone physically walks over to the control panel and has an opportunity to investigate the reason for the trip. This feature is particularly useful when a multitude of inputs are wired to trip the 86 relay. In the case

of a power generator, such inputs may include (but are not limited to) a loss-of-field-excitation relay (40), a thermal overload relay (49), and/or a directional power relay (32), the trip contacts of all these relays paralleled with the differential current relay (87) contacts shown in the example diagram such that closure of any one energizes the lockout relay's trip coil and takes the generator off-line. The requirement of manually resetting the 86 relay prompts the human operator to survey which of these relays was responsible for initiating the lockout trip.

25.14 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Transformer ratios:** the voltage or current step-ratio of a transformer is proportional to the turns ratio of its windings. Relevant to the design of potential and current transformers (PTs and CTs).
- **Voltage versus current sources:** voltage sources try to maintain constant voltage with variable current, while current sources try to maintain constant current with variable voltage. Relevant to the safe use of potential and current transformers (PTs and CTs): never short-circuit a PT and never open-circuit a CT!
- **Electrical sources versus loads:** electrical power sources output current (conventional flow) on their positive terminals and input current on their negative terminals (e.g. batteries and generators). Electrical loads do the opposite (e.g. resistors). Relevant to determining the phasing of voltages and currents for instrument transformers (PTs and CTs), as well as to the operation of a directional overcurrent (67) protective relay.
- **“Normal” switch status:** the “normal” status of a switch contact as defined by the manufacturer is its *resting* condition (minimum stimulus).
- **“Seal-in” circuit:** when an electrical relay uses one of its own switch contacts to continue its own coil energization after the initial triggering event has passed. Relevant to induction-disk protective relays, where the time-overcurrent trip contact moved by the disk requires a seal-in contact in parallel to maintain solid closure for a reliable breaker trip.
- **Lenz's Law:** any magnetic field arising from electromagnetic induction opposes the inducing field. Relevant to the use of “drag magnets” in induction-disk style overcurrent relays.
- **Impedance and phase shift:** impedance is a phasor combination of resistance and reactance. If current lags voltage, the impedance has an inductive characteristic. If current leads voltage, the impedance has a capacitive characteristic.

References

“Allowable CT Secondary Lead Lengths for Automatic Capacitor Bank and Harmonic Filter Bank Controllers”, Northeast Power Systems, Inc., Queensbury, NY, 1999-2012.

“Automatic Reclosing – Transmission Line Applications and Considerations”, Basler Electric Company, Highland, IL.

Blackburn, J. Lewis and Domin, Thomas J., *Protective Relaying Principles and Applications*, Third Edition, CRC Press, Taylor & Francis Group, Boca Raton, FL, 2007.

Boggs, Steven A., “Sulphur Hexafluoride: Introduction to the Material and Dielectric”, *IEEE Electrical Insulation Magazine*, Volume 5, Number 5, pages 18-21, September/October 1989.

Boylestad, Robert L., *Introductory Circuit Analysis*, 9th Edition, Prentice Hall, Upper Saddle River, NJ, 2000.

Choubey, Tribhuwan, “Communication Aided Protection Schemes”, Southern California Edison presentation.

“CT Application Guide for the 489 Generator Management Relay”, GE Publication number GET-8402, General Electric Power Managements, 2002.

Electrical Metermen’s Handbook, Seventh Edition, Edison Electric Institute, New York, NY, 1965.

Ganesan, Sethuraman, “Selection of Current Transformers & Wire Sizing in Substations”, ABB Inc., Allentown, PA.

Greene, S. Dana, “Distribution of the Electrical Energy from Niagara Falls”, *Cassier’s Magazine*, Volume 8, pp. 333-362, London, England, 1895.

IEEE C57.12.00-2010, IEEE Standard for General Requirements for Liquid-Immersed Distribution, Power, and Regulating Transformers, IEEE Power Engineering Society, New York, NY, 2010.

IEEE C57.13.3-2005, IEEE Guide for Grounding of Instrument Transformer Secondary Circuits and Cases, IEEE Power Engineering Society, New York, NY, 2005.

“Instrument Transformer Basic Technical Information and Application”, General Electric Company.

“Instrument Transformers Buyer’s Guide”, GE Meter, Somersworth, NH.

“Introduction to Synchronizing – Automatic Synchronizing Considerations and Applications”, Basler Electric Company, Highland, IL.

Jones, Kent, “Instrument Transformer Basics”, Line Power Manufacturing.

“Maintenance of Power Circuit Breakers”, Facilities Instructions, Standards and Techniques volume 3-16, United States Department of the Interior, Bureau of Reclamation, Denver, CO, December

1999.

Martin, David; Sharma, Pankaj; Sinclair, Amy; Finney, Dale; “Distance Protection in Distribution Systems: How It Assists With Integrating Distributed Resources”, Hydro One Networks Inc. and Schweitzer Engineering Laboratories, Inc., 2011.

Mason, C. Russell, *The Art and Science of Protective Relaying*, First Edition, John Wiley & Sons, 1956.

“Offset Mho Distance Relay Type CEB51B Instructions”, document GEK-26420B, GE Protection and Control, Malvern, PA, April 1993.

Protection and Control Reference Guide, Volume 22, GE Digital Energy, 2009.

Schweitzer, E.O. III, “A Review of Impedance-Based Fault Locating Experience”, paper presented before the fourteenth annual Iowa-Nebraska System Protection Seminar, Schweitzer Engineering Laboratories, Inc., Pullman, WA, 1990.

Schweitzer, E.O. III and Roberts, Jeff; “Distance Relay Element Design”, revision 2, Schweitzer Engineering Laboratories, Inc., 1993.

“SEL-551 Relay instruction manual”, PM551-01, Schweitzer Engineering Laboratories, Inc., 1998-2011.

Todd, Victor H., *Protective Relays – their theory, design, and practical operation*, First Edition, Fourth Impression, McGraw-Hill Book Company, Inc., New York, 1922.

Chapter 26

Signal characterization

Mathematics is full of complementary principles and symmetry. Perhaps nowhere is this more evident than with *inverse functions*: functions that “un-do” one another when put together. A few examples of inverse functions are shown in the following table:

$f(x)$	$f^{-1}(x)$
Addition	Subtraction
Multiplication	Division
Power	Root
Exponential	Logarithm
Derivative	Integral

Inverse functions are vital to master if one hopes to be able to manipulate algebraic (literal) expressions. For example, to solve for time (t) in this exponential formula, you must know that the natural logarithm function directly “un-does” the exponential e^x . This is the only way to “unravel” the equation and get t isolated by itself on one side of the equals sign:

$$V = 12e^{-t}$$

Divide both sides by 12

$$\frac{V}{12} = e^{-t}$$

Take the natural logarithm of both sides

$$\ln\left(\frac{V}{12}\right) = \ln(e^{-t})$$

The natural logarithm “cancels out” the exponential

$$\ln\left(\frac{V}{12}\right) = -t$$

Multiply both sides by negative one

$$-\ln\left(\frac{V}{12}\right) = t$$

In industry there exist a great many practical problems where inverse functions play a similar role. Just as inverse functions are useful for manipulating literal expressions in algebra, they are also useful in inferring measurements of things we cannot directly measure. Many continuous industrial measurements are *inferential* in nature, meaning that we actually measure some other variable in order to quantify the variable of interest. More often than not, the relationship between the primary variable and the inferred variable is nonlinear, necessitating some form of mathematical processing to complete the inferential measurement.

Practical examples of inferential measurement include:

- Liquid flow measurement through restrictions in pipes
- Liquid flow measurement in open channels (over weirs)
- Liquid level measurement in non-cylindrical vessels
- Temperature measurement by radiated energy
- Chemical composition measurement


The following sections will describe the mathematics behind each of these measurement applications.

26.1 Flow measurement from differential pressure


Take for instance the problem of measuring fluid flow through a pipe. To the layperson, this may seem to be a trivial problem. However there is no practical way to *directly* and continuously measure the flow rate of a fluid, especially when we cannot allow the fluid in question to become exposed to the atmosphere (e.g. when the liquid or gas in question is toxic, flammable, under high pressure, or any combination thereof).

One standard way to measure the flow rate of a fluid through a pipe is to intentionally place a restriction in the path of the fluid, and measure the pressure drop across that restriction. The most common form of intentional restriction used for this purpose is a thin plate of metal with a hole precisely machined in the center, called an *orifice plate*.

A side view of the orifice plate assembly and pressure-measuring instrument looks like this:



This approach should make intuitive sense: the faster the flow rate of the fluid, the greater the pressure difference developed across the orifice. The actual physics of this process has to do with energy exchanging between potential and kinetic forms, but that is incidental to this discussion. The mathematically interesting characteristic of this flow measurement technique is its nonlinearity. Pressure does not rise linearly with flow rate; rather, it increases with the *square* of the flow rate:




To write this as a proportionality, we relate flow rate (Q) to pressure (P) as follows (the constant k accounts for unit conversions and the geometries of the orifice plate and pipe):

$$P = kQ^2$$

This is a practical problem for us because our intent is to use pressure measurement (P) as an indirect (inferred) indication of flow rate (Q). If the two variables are not *directly* related to one another, we will not be able to regard one as being directly representative of the other. To make this problem more clear to see, imagine a pressure gauge connected across the restriction, with the face of the gauge labeled in percent:

Face of pressure gauge, calibrated to read in percent of full flow rate



Consider a pressure gauge such as the one shown above, registering 20 percent on a linear scale at some amount of flow through the pipe. What will happen if the flow rate through that pipe suddenly doubles? An operator or technician looking at the gauge *ought* to see a new reading of 40 percent, if indeed the gauge is supposed to indicate flow rate. However, this will not happen. Since the pressure dropped across the orifice in the pipe increases with the square of flow rate, a doubling of flow rate will actually cause the pressure gauge reading to *quadruple!* In other words, it will go from reading 20% to reading 80%, which is definitely not an accurate indication of the flow increase.

A couple of simple solutions exist for addressing this problem. One is to re-label the pressure gauge with a “square root” scale. Examine this photograph of a 3-15 PSI receiver gauge:




Now, a doubling of fluid flow rate still results in a quadrupling of needle motion, but due to the nonlinear (inner) scale on this gauge this needle motion translates into a simple doubling of indicated flow, which is precisely what we need for this to function as an accurate flow indicator.

If the differential pressure instrument outputs a 4-20 mA analog electronic signal instead of a 3-15 PSI pneumatic signal, we may apply the same “nonlinear scale” treatment to any current meter and achieve the same result:



Another simple solution is to use a *nonlinear manometer*, with a curved viewing tube¹:



The scale positioned alongside the curved viewing tube will be linear, with equal spacings between division marks along its entire length. The vertical height of the liquid column translates pressure into varying degrees of movement along the axis of the tube by the tube’s curvature. Literally, any


¹This solution works best for measuring the flow rate of gases, not liquids, since the manometer obviously must use a liquid of its own to indicate pressure, and mixing or other interference between the process liquid and the manometer liquid could be problematic.

inverse function desired may be “encoded” into this manometer by fashioning the viewing tube into the desired (custom) shape without any need to print a nonlinear scale.

Shown here is a photograph of an actual curved-tube manometer. This particular specimen does not have a scale reading in units of flow, but it certainly could if it had the correct curve for a square-root characterization:



A more sophisticated solution to the “square root problem” is to use a computer to manipulate the signal coming from the differential pressure instrument so the characterized signal becomes a direct, linear representation of flow. In other words, the computer *square-roots* the pressure sensor’s signal in order that the final signal becomes a direct representation of fluid flow rate:




Both solutions achieve their goal by mathematically “un-doing” the nonlinear (square) function intrinsic to the physics of the orifice plate with a complementary (inverse) function. This intentional compounding of inverse functions is sometimes called *linearization*, because it has the overall effect of making the output of the instrument system a direct proportion of the input:

$$\text{Output} = k(\text{Input})$$


26.2 Flow measurement in open channels

Measuring the flow rate of liquid through an open channel is not unlike measuring the flow rate of a liquid through a closed pipe: one of the more common methods for doing so is to place a restriction in the path of the liquid flow and then measure the “pressure” dropped across that restriction. The easiest way to do this is to install a low “dam” in the middle of the channel, then measure the height of the liquid upstream of the dam as a way to infer flow rate. This dam is technically referred to as a *weir*, and three styles of weir are commonly used:

Different styles of weirs for measuring open-channel liquid flow



Another type of open-channel restriction used to measure liquid flow is called a *flume*. An illustration of a *Parshall flume* is shown here:



Weirs and *flumes* may be thought of being somewhat like “orifice plates” and “venturi tubes,” respectively, for open-channel liquid flow. Like an orifice plate, a weir or a flume generates a differential pressure that varies with the flow rate through it. However, this is where the similarities end. Exposing the fluid stream to atmospheric pressure means the differential pressure caused by the flow rate manifests itself as a difference in liquid height at different points in the channel. Thus, weirs and flumes allow the indirect measurement of liquid flow by sensing liquid height. An interesting feature of weirs and flumes is that although they are nonlinear primary sensing elements, their nonlinearity is quite different from that of an orifice.

Note the following transfer functions for different weirs and flumes, relating the rate of liquid flow through the device (Q) to the level of liquid rise upstream of the device (called “head”, or H):

$$Q = 2.48 \left(\tan \frac{\theta}{2} \right) H^{\frac{5}{2}} \quad \text{V-notch weir}$$

$$Q = 3.367LH^{\frac{3}{2}} \quad \text{Cippoletti weir}$$

$$Q = 0.992H^{1.547} \quad \text{3-inch wide throat Parshall flume}$$

$$Q = 3.07H^{1.53} \quad \text{9-inch wide throat Parshall flume}$$

Where,

Q = Volumetric flow rate (cubic feet per second – CFS)

L = Width of notch crest or throat width (feet)

θ = V-notch angle (degrees)

H = Head (feet)

It is important to note these functions provide answers for flow rate (Q) with head (H) being the independent variable. In other words, they will tell us how much liquid is flowing given a certain head. In the course of calibrating the head-measuring instruments that infer flow rate, however, it is important to know the inverse transfer function: how much head there will be for any given value of flow. Here, algebraic manipulation becomes important to the technician. For example, here is the solution for H in the function for a Cippoletti weir:

$$Q = 3.367LH^{\frac{3}{2}}$$

Dividing both sides of the equation by 3.367 and L:

$$\frac{Q}{3.367L} = H^{\frac{3}{2}}$$

Taking the $\frac{3}{2}$ root of both sides:


$$\sqrt[3/2]{\frac{Q}{3.367L}} = H$$

This in itself may be problematic, as some hand calculators do not have an $\sqrt[y]{x}$ function. In cases such as this, it is helpful to remember that a root is nothing more than an inverse power. Therefore, we could re-write the final form of the equation using a $\frac{2}{3}$ power instead of a $\frac{3}{2}$ root:

$$\left(\frac{Q}{3.367L}\right)^{\frac{2}{3}} = H$$


26.3 Material volume measurement

A variety of technologies exist to measure the quantity of stored material in a vessel. For liquid applications, hydrostatic pressure, radar, ultrasonic, and tape-and-float are just a few of the more common technologies:



These measuring technologies share a common trait: they infer the quantity of material stored in the vessel by measuring *height*. If the vessel in question has a constant cross-sectional area throughout its working height (e.g. a vertical cylinder), then material height will directly correspond to stored material volume. However, if the vessel in question does not have a constant cross-sectional area throughout its height, the relationship between material height and material volume will not be linear.

For example, there is a world of difference between the height/volume functions for a vertical cylinder versus a horizontal cylinder:



The volume function for a vertical cylinder is a simple matter of geometry – height (h) multiplied by the cylinder's cross-sectional area (πr^2):


$$V = \pi r^2 h$$

Calculating the volume of material stored inside a horizontal cylinder as a function of material height (h) is a far more complicated matter, because the cross-sectional area is not constant. The solution to this problem is an exercise in integral calculus, the final result being this unwieldy function:

$$V = L \left[(h - r) \sqrt{2hr - h^2} + r^2 \sin^{-1} \frac{(h - r)}{r} + \frac{\pi r^2}{2} \right]$$

Any instrumentation system inferring stored material volume by measurement of material height in a horizontal cylinder must somehow apply this formula on a continuous basis. This is a prime example of how digital computer technology is essential to certain continuous measurement applications!

Spherical vessels, such as those used to store liquefied natural gas (LNG) and butane, present a similar challenge: the material height/volume function is more complicated than for a vertical cylinder because the cross-sectional area of a spherical vessel changes with height. Once again an exercise in integral calculus is necessary to derive a formula for calculating the stored volume of material inside a sphere based on material height:




Whereas the material volume function for a vertical cylinder is a simple linear formula ($V = \pi r^2 h$), the material volume formula for a sphere is a bit more complicated although not as complicated as it is for a horizontal cylinder:

$$V = \pi h^2 \left(r - \frac{h}{3} \right)$$


This function will “un-do” the inherent height/volume nonlinearity of a spherical vessel, allowing a height measurement to translate directly into a volume measurement. A “characterizing” function such as this is typically executed in a digital computer connected to the level sensor, or sometimes in a computer chip within the sensor device itself.

An interesting alternative to a formal equation for linearizing the level measurement signal is something called a *multi-segment characterizer* function, typically implemented as a function block inside a digital loop controller or distributed control system (DCS). This is an example of what mathematicians call a *piecewise function*: a function made up of line segments. Multi-segment characterizer functions may be programmed to emulate virtually any continuous function, with reasonable accuracy:

Continuous characterizing function




Piecewise characterizing function



The computer correlates the input signal (height measurement, h) to a point on this piecewise function, linearly interpolating between the nearest pair of programmed coordinate points. The number of points available for multi-point characterizers varies between ten and one hundred² depending on the desired accuracy and the available computing power.


²There is no theoretical limit to the number of points in a digital computer's characterizer function given sufficient processing power and memory. There is, however, a limit to the patience of the human programmer who must encode all the necessary x, y data points defining this function. Most of the piecewise characterizing functions I have seen available in digital instrumentation systems provide 10 to 20 (x, y) coordinate points to define the function. Fewer than 10 coordinate points risks excessive interpolation errors, and more than 20 would just be tedious to configure.

Although true fans of math might blanch at the idea of approximating an inverse function for level measurement using a piecewise approach rather than simply implementing the correct continuous function, the multi-point characterizer technique does have certain practical advantages. For one, it is readily adaptable to any shape of vessel, no matter how strange. Take for instance this vessel, made of separate cylindrical sections welded together:



Here, the vessel's very own height/volume function is fundamentally piecewise, and so *nothing but* a piecewise characterizing function could possibly linearize the level measurement into a volume measurement!

Consider also the case of a spherical vessel with odd-shaped objects welded to the vessel walls, and/or inserted into the vessel's interior:




The volumetric space occupied by these structures will introduce all kinds of discontinuities into the transfer function, and so once again we have a case where a continuous characterizing function cannot properly linearize the level signal into a volume measurement. Here, *only* a piecewise function will suffice.

To best generate the coordinate points for a proper multi-point characterizer function, one must collect data on the storage vessel in the form of a *strapping table*. This entails emptying the vessel completely, then filling it with measured quantities of material, one sample at a time, and taking level readings. The result of this exercise is a data table correlating known volumes of material inside the vessel with measured heights of that material:

Introduced material volume	Measured material level
150 gallons	2.46 feet
300 gallons	4.72 feet
450 gallons	5.8 feet
600 gallons	(etc., etc.)
750 gallons	(etc., etc.)

Each of these paired numbers would constitute the coordinates to be programmed into the characterizer function computer by the instrument technician or engineer:



With this programmed function, the computer is able to infer material volume quite accurately from the height measurement reported by a level transmitter.

Many “smart” level transmitter instruments possess enough computational power to perform the level-to-volume characterization directly, so as to transmit a signal corresponding directly to material volume rather than just material level. This eliminates the need for an external “level computer” to perform the necessary characterization. The following screenshot was taken from a personal computer running configuration software for a radar level transmitter³, showing the strapping table data point fields where a technician or engineer would program the vessel’s level-versus-volume piecewise function:

Configure/Setup of CTRL-01C04CH02 [3300 Rev. 2]

File Actions Help

Basic Setup Setup **Volume** LCD Analog Output Advanced Version

Tank Type: Ver Cylindr Tank Diameter: 39.37 in Tank Height: 196.85 in

Current Measurement: Level: 8.93 in Volume: 47 gal

Strapping Table

Entries used: 2 Max entries: 10

	Level	Volume		Level	Volume
1:	0.00 in	0 gal	6:	0.00 in	0 gal
2:	0.00 in	0 gal	7:	0.00 in	0 gal
3:	0.00 in	0 gal	8:	0.00 in	0 gal
4:	0.00 in	0 gal	9:	0.00 in	0 gal
5:	0.00 in	0 gal	10:	0.00 in	0 gal

Time: Current


OK Cancel Apply Help

Device Last Synchronized: 8/18/2008 2:35:12 PM

This configuration window actually shows more than just a strapping table. It also shows the option of calculating volume for different vessel shapes (*vertical cylinder* is the option selected here) including horizontal cylinder and sphere. In order to use the strapping table option, the user would have to select “Strapping Table” from the list of Tank Types. Otherwise, the level transmitter’s computer will attempt to calculate volume from an ideal tank shape.

³The configuration software is Emerson’s AMS, running on an engineering workstation in a DeltaV control system network. The radar level transmitter is a Rosemount model 3301 (guided-wave) unit.

Solids material volume measurement poses unique problems over liquid material volume measurement, because solids – unlike liquids – are not self-leveling. A common problem encountered in measuring the level of a solid (chip, granular, or powder) inside a vessel is the *angle of repose* created when that material forms a pile inside the vessel:



It should be clear from the above illustration that any level sensor detecting the level of the solid material at any one point inside the vessel will usually yield a misleading result because it cannot sense variations in material level across the whole surface. This problem exists regardless of level sensor type (e.g. ultrasonic, radar, tape) so long as that sensor only senses material level at one point. Liquid material applications do not pose this problem because liquids naturally self-level, achieving a relatively flat surface by nature.

A modern solution to this measurement problem is to use an array of echo-based sensors to generate a three-dimensional *map* of the solid material surface, and then compute material volume based on all height measurements spanning the vessel's cross-section. One such instrument designed for this purpose is Rosemount's model 5708 3D Solids Scanner.

26.4 Radiative temperature measurement

Temperature measurement devices may be classified into two broad types: *contact* and *non-contact*. Contact-type temperature sensors detect temperature by directly touching the material to be measured, and there are several varieties in this category. Non-contact temperature sensors work by detecting light emitted by hot objects.

Energy radiated in the form of electromagnetic waves (photons, or light) relates to object temperature by an equation known as the Stefan-Boltzmann equation, which tells us the rate of heat lost by radiant emission from a hot object is proportional to the *fourth power* of its absolute temperature:

$$P = e\sigma AT^4$$

Where,

P = Radiated energy power (watts)

e = Emissivity factor (unitless)

σ = Stefan-Boltzmann constant (5.67×10^{-8} W / m² · K⁴)

A = Surface area (square meters)

T = Absolute temperature (Kelvin)

Solving for temperature (T) involves the use of the fourth root, to “un-do” the fourth power function inherent to the original function:

$$T = \sqrt[4]{\frac{P}{e\sigma A}}$$

Any optical temperature sensor measuring the emitted power (P) must “characterize” the power measurement using the above equation to arrive at an inferred temperature. This characterization is typically performed inside the temperature sensor by a microcomputer.

26.5 Analytical measurements

A great many chemical composition measurements may be made indirectly by means of electricity, if those measurements are related to the concentration of ions (electrically charged molecules). Such measurements include:

- pH of an aqueous solution
- Oxygen concentration in air
- Ammonia concentration in air
- Lead concentration in water

The basic principle works like this: two different chemical samples are placed in close proximity to each other, separated only by an *ion-selective membrane* able to pass the ion of interest. As the ion activity attempts to reach equilibrium through the membrane, an electrical voltage is produced across that membrane. If we measure the voltage produced, we can infer the relative activity of the ions on either side of the membrane.

Not surprisingly, the function relating ion activity to the voltage generated is nonlinear. The standard equation describing the relationship between ionic activity on both sides of the membrane and the voltage produced is called the *Nernst equation*:

$$V = \frac{RT}{nF} \ln \left(\frac{a_1}{a_2} \right)$$

Where,

V = Electrical voltage produced across membrane due to ion exchange (volts)

R = Universal gas constant (8.315 J/mol·K)

T = Absolute temperature (Kelvin)

n = Number of electrons transferred per ion exchanged (unitless)


F = Faraday constant (96485 coulombs per mole)

a_1 = Activity of ion in measured sample

a_2 = Activity of ion in reference sample (on other side of membrane)

A practical application for this technology is in the measurement of oxygen concentration in the flue gas of a large industrial burner, such as what might be used to heat up water to generate steam. The measurement of oxygen concentration in the exhaust of a combustion heater (or boiler) is very important both for maximizing fuel efficiency and for minimizing pollution (specifically, the production of NO_x molecules). Ideally, a burner's exhaust gas will contain no oxygen, having consumed it all in the process of combustion with a perfect stoichiometric mix of fuel and air. In practice, the exhaust gas of an efficiently-controlled burner will be somewhere near 2%, as opposed to the normal 21% of ambient air.

One way to measure the oxygen content of hot exhaust is to use a *high-temperature zirconium oxide* detector. This detector is made of a “sandwich” of platinum electrodes on either side of a solid zirconium oxide electrolyte. One side of this electrochemical cell is exposed to the exhaust gas (process), while the other side is exposed to heated air which serves as a reference:



The electrical voltage generated by this “sandwich” of zirconium and platinum is sent to an electronic amplifier circuit, and then to a microcomputer which applies an inverse function to the measured voltage in order to arrive at an inferred measurement for oxygen concentration. This type of chemical analysis is called *potentiometric*, since it measures (“metric”) based on an electrical voltage (“potential”).

The Nernst equation is an interesting one to unravel, to solve for ion activity in the sample (a_1) given voltage (V):

$$V = \frac{RT}{nF} \ln \left(\frac{a_1}{a_2} \right)$$

Multiplying both sides by nF :

$$nFV = RT \ln \left(\frac{a_1}{a_2} \right)$$

Dividing both sides by RT :

$$\frac{nFV}{RT} = \ln \left(\frac{a_1}{a_2} \right)$$

Applying the rule that the difference of logs is equal to the log of the quotient:

$$\frac{nFV}{RT} = \ln a_1 - \ln a_2$$

Adding $\ln a_2$ to both sides:

$$\frac{nFV}{RT} + \ln a_2 = \ln a_1$$

Making both sides of the equation a power of e :

$$e^{\frac{nFV}{RT} + \ln a_2} = e^{\ln a_1}$$

Canceling the natural log and exponential functions on the right-hand side:

$$e^{\frac{nFV}{RT} + \ln a_2} = a_1$$

In most cases, the ionic activity of a_2 will be relatively constant, and so $\ln a_2$ will be relatively constant as well. With this in mind, we may simplify the equation further, using k as our constant value:

Substituting k for $\ln a_2$:

$$e^{\frac{nFV}{RT} + k} = a_1$$

Applying the rule that the sum of exponents is the product of powers:

$$e^k e^{\frac{nFV}{RT}} = a_1$$

If k is constant, then e^k will be constant as well (calling the new constant C):

$$C e^{\frac{nFV}{RT}} = a_1$$

Analytical instruments based on potentiometry must evaluate this inverse function to “undo” the Nernst equation to arrive at an inferred measurement of ion activity in the sample given the small voltage produced by the sensing membrane. These instruments typically have temperature sensors as well built in to the sensing membrane assembly, since it is apparent that temperature (T) also plays a role in the generation of this voltage. Once again, this mathematical function is typically evaluated in a microprocessor.

26.6 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Inverse mathematical functions:** an inverse function, when applied to the result of its counterpart function, “un-does” the operation and leaves you with the original quantity. Relevant to all manner of “linearization” techniques designed to make an inherently nonlinear measurement read linearly.
- **Bernoulli's equation:** $z_1\rho g + \frac{v_1^2\rho}{2} + P_1 = z_2\rho g + \frac{v_2^2\rho}{2} + P_2$, which is an application of the Law of Energy Conservation, stating that the sum of all forms of energy in a moving fluid stream (height, kinetic, and pressure) must remain the same. Relevant to calculations of pressure drop and pressure recovery across restrictions such as venturi tubes, orifice plates, etc.
- **Nernst equation:** $V = \frac{RT}{nF} \ln\left(\frac{a_1}{a_2}\right)$ predicting the amount of voltage developed across an ion-permeable membrane by the exchange of ions across that membrane. Relevant to all forms of potentiometric chemical analysis, where sensor voltage is proportional to the logarithm of concentration quotient across the sensor membrane.

References

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.

Stewart, James, *Calculus: Concepts and Contexts*, 2nd Edition, Brooks/Cole, Pacific Grove, CA, 2001.

Chapter 27

Control valves

One of the most common final control elements in industrial control systems is the *control valve*. A “control valve” works to restrict the flow of fluid through a pipe at the command of a remotely sourced signal, such as the signal from a loop controller or logic device (such as a PLC), or even a manual (“hand”) interface controlled by a human operator. Some control valve designs are intended for discrete (on/off) control of fluid flow, while others are designed to *throttle* fluid flow somewhere between fully open and fully closed (shut), inclusive. The electrical equivalent of an on/off valve is a switch, while the electrical equivalent of a throttling valve is a variable resistor.

Control valves are comprised of two major parts: the *valve body*, containing all the mechanical components necessary to influence fluid flow; and the *valve actuator*, providing the mechanical power necessary to move the valve body components. Often times, the major difference between an on/off control valve and a throttling control valve is the type of actuator applied to the valve¹: on/off actuators need only position a valve mechanism two one of two extreme positions (fully open or fully closed). Throttling actuators must be able to accurately position a valve mechanism anywhere between those extremes.


Within a control valve body, the specific components performing the work of throttling (or completely shutting off) of fluid flow are collectively referred to as the valve *trim*. For each major type of control valve, there are usually many variations of trim design. The choice of valve type, and of specific trim for any type of valve, is a decision dictated by the type of fluid being controlled, the nature of the control action (on/off versus throttling), the process conditions (expected flow rate, temperature, pressures, etc.), and economics.

An appendix of this book (Appendix C beginning on page 3167) photographically documents the complete disassembly of a typical control valve. The valve happens to be a Fisher E-body globe valve with a pneumatic diaphragm actuator.

¹To be honest, there are some valve body designs that work far better in on/off service (e.g. ball valves and plug valves) while other designs do a better job at throttling (e.g. double-ported globe valves). Many valve designs, however, may be pressed into either type of service merely by attaching the appropriate actuator.

27.1 Sliding-stem valves

A *sliding-stem* valve body is one where the moving parts slide with a linear motion. Some examples of sliding-stem valve body designs are shown here:



Most sliding-stem control valves are *direct acting*, which means the valve opens up wider as the stem is drawn out of the body. Conversely, a direct-acting valve shuts off (closes) when the stem is pushed into the body. Of course, a *reverse-acting* valve body would behave just the opposite: opening up as the stem is pushed in and closing off as the stem is drawn out.

27.1.1 Globe valves

Globe valves restrict the flow of fluid by altering the distance between a movable plug and a stationary seat (in some cases, a pair of plugs and matching seats). Fluid flows through a hole in the center of the seat, and is more or less restricted by the plug's proximity to that hole. The globe valve design is one of the most popular sliding-stem valve designs used in throttling service. A photograph of a small (2 inch) globe valve body appears here:



A set of three photographs showing a cut-away Masoneilan model 21000 globe valve body illustrates just how the moving plug and stationary seat work together to throttle flow in a direct-acting globe valve. The left-hand photo shows the valve body in the fully closed position, while the middle photo shows the valve half-open, and the right-hand photo shows the valve fully open:




As you can see from these photographs, the valve plug is guided by the stem to maintain alignment with the centerline of the seat. For this reason, this particular style of globe valve is called a *stem-guided* globe valve.

A variation on the stem-guided globe valve design is the *needle valve*, where the plug is extremely small in diameter and usually fits well into the seat hole rather than merely sitting on top of it. Needle valves are very common as manually-actuated valves used to control low flow rates of air or oil. A set of three photographs shows a needle valve in the fully-closed, mid-open, and fully-open positions (left-to-right):




Yet another variation on the globe valve design is the *port-guided* valve, where the plug has an unusual shape, projecting into the seat. Thus, the seat ring acts as a guide for the plug to keep the centerlines of the plug and seat always aligned, minimizing guiding stresses that would otherwise be placed on the stem. This means that the stem may be made smaller in diameter than if the valve trim were stem-guided, minimizing sliding friction and improving control behavior.



A photograph showing a small port-guided globe valve plug appears in the following photograph:



Some globe valves use a pair of plugs (on the same stem) and a matching pair of seats to throttle fluid flow. These are called *double-ported* globe valves. The purpose of a double-ported globe valve is to minimize the force applied to the stem by process fluid pressure across the plugs:



Differential pressure of the process fluid ($P_1 - P_2$) across a valve plug will generate a force parallel to the stem as described by the formula $F = PA$, with A being the plug's effective area presented for the pressure to act upon. In a single-ported globe valve, there will only be one force generated by the process pressure. In a double-ported globe valve, there will be *two opposed* force vectors, one generated at the upper plug and another generated at the lower plug. If the plug areas are approximately equal, then the forces will likewise be approximately equal and therefore nearly cancel. This makes for a control valve that is easier to actuate (i.e. the stem position is less affected by pressure drop across the valve).

The following photograph shows a disassembled Fisher “A-body” double-ported globe valve, with the double plug plainly visible on the right:



This particular double-ported globe valve happens to be stem-guided, with bushings guiding the upper stem and also a lower stem (on the bottom side of the valve body). Double-ported, port-guided control valves also exist, with two sets of port-guided plugs and seats throttling fluid flow.

While double-ported globe valves certainly enjoy the advantage of easier actuation compared to their single-ported cousins, they also suffer from a distinct disadvantage: the near impossibility of tight shut-off. With *two* plugs needing to come to simultaneous rest on *two* seats to achieve a fluid-tight seal, there is precious little room for error or dimensional instability. Even if a double-ported valve is prepared in a shop for the best shut-off possible², it may not completely shut off when installed due to dimensional changes caused by process fluid heating or cooling the valve stem and body. This is especially problematic when the stem is made of a different material than the body. Globe valve stems are commonly manufactured from stainless steel bar stock, while globe valve bodies are commonly made of cast steel. Cold-formed stainless steel has a different coefficient of thermal expansion than hot-cast steel, which means the plugs will no longer simultaneously seat once the valve warms or cools much from the temperature it was at when it seated tightly.


²The standard preparatory technique is called *lapping*. To “lap” a valve plug and seat assembly, an abrasive paste known as *lapping compound* is applied to the valve plug(s) and seat(s) at the areas of mutual contact when the valve is disassembled. The valve mechanism is reassembled, and the stem is then rotated in a cyclic motion such that the plug(s) grind into the seat(s), creating a matched fit. The precision of this fit may be checked by disassembling the valve, cleaning off all remaining lapping compound, applying a metal-staining compound such as *Prussian blue*, then reassembling. The stem is rotated once more such that the plug(s) will rub against the seat(s), wearing through the applied stain. Upon disassembly, the worn stain may be inspected to reveal the extend of metal-to-metal contact between the plug(s) and the seat(s). If the contact area is deemed insufficient, the lapping process may be repeated.

A more modern version of the globe valve design uses a piston-shaped plug inside a surrounding *cage* with ports cast or machined into it. These *cage-guided* globe valves throttle flow by uncovering more or less of the port area in the surrounding cage as the plug moves up and down. The cage also serves to guide the plug so the stem need not be subjected to lateral forces as in a stem-guided valve design. A photograph of a cut-away control valve shows the appearance of the cage (in this case, with the plug in the fully closed position). Note the “T”-shaped ports in the cage, through which fluid flows as the plug moves up and out of the way:




An advantage of the cage-guided design is that the valve’s flowing characteristics may be easily altered just by replacing the cage with another having different size or shape of holes. By contrast, stem-guided and port-guided globe valves are characterized by the shape of the plug, which requires further disassembly to replace than the cage in a cage-guided globe valve. With most cage-guided valves all that is needed to replace the cage is to separate the bonnet from the rest of the valve body, at which point the cage may be lifted out of the body and swapped with another cage. In order to change a globe valve’s plug, you must first separate the bonnet from the rest of the body and then de-couple the plug and plug stem from the actuator stem, being careful not to disturb the packing inside of the bonnet as you do so. After replacing a plug, the “bench-set” of the valve must be re-adjusted to ensure proper seating pressure and stroke calibration.

Cage-guided globe valves are available with both *balanced* and *unbalanced* plugs. A balanced plug has one or more ports drilled from top to bottom, allowing fluid pressure to equalize on both sides of the plug. This helps minimize the forces acting on the plug which must be overcome by the actuator:




Unbalanced plugs generate a force equal to the product of the differential pressure across the plug and the plug's area ($F = PA$), which may be quite substantial in some applications. Balanced plugs do not generate this same force because they equalize the pressure on both sides of the plug, however, they exhibit the disadvantage of one more leak path when the valve is in the fully closed position (through the balancing ports, past the piston ring, and out the cage ports):



Thus, balanced and unbalanced cage-guided globe valves exhibit similar characteristics to double-ported and single-ported stem- or port-guided globe valves, and for similar reasons. Balanced cage-guided valves are easy to position, just like double-ported stem-guided and port-guided globe valves. However, balanced cage-guided valves tend to leak more when in the shut position due to a greater number of leak paths, much the same as with double-ported stem-guided and port-guided globe valves.

Another style of globe valve body is the *three-way* body, sometimes called a *mixing* or a *diverting* valve. This valve design has three ports on it, with the plug (in this particular case, a cage-guided plug) controlling the degree to which two of the ports connect with the third port:



This dual illustration shows a three-way valve in its two extreme stem positions. If the stem is positioned between these two extremes, all three ports will be “connected” to varying degrees. Three-way valves are useful in services where a flow stream must be diverted (split) between two different directions, or where two flow streams must converge (mix) within the valve to form a single flow stream.


A photograph of a three-way globe valve mixing hot and cold water to control temperature is shown here:



27.1.2 Gate valves

Gate valves work by inserting a dam (“gate”) into the path of the flow to restrict it, in a manner similar to the action of a sliding door. Gate valves are more often used for on/off control than for throttling.

The following set of photographs shows a hand-operated gate valve (cut away and painted for use as an instructional tool) in three different positions, from full closed to full open (left to right):



27.1.3 Diaphragm valves


Diaphragm valves use a flexible sheet pressed close to the edge of a solid dam to narrow the flow path for fluid. Their operation is not unlike controlling the flow of water through a flexible hose by pinching the hose. These valves are well suited for flows containing solid particulate matter such as slurries, although precise throttling may be difficult to achieve due to the elasticity of the diaphragm. The next photograph shows a diaphragm valve actuated by an electric motor, used to control the flow of treated sewage:



The following photograph shows a hand-actuated diaphragm valve, the external shape of the valve body revealing the “dam” structure against which the flexible diaphragm is pressed to create a leak-tight seal when shut:




Some diaphragm valves are pneumatically actuated, using the force of compressed air on one side of the diaphragm to press it against the dam (on the other side) to shut off flow. This next example is of a small air-actuated diaphragm valve, controlling the flow of water through a 1-inch pipe:



The actuating air for this particular diaphragm valve comes through an electric solenoid valve. The solenoid valve in this photograph has a brass body and a green-painted solenoid coil.

27.2 Rotary-stem valves

A different strategy for controlling the flow of fluid is to insert a rotary element into the flow path. Instead of sliding a stem into and out of the valve body to actuate a throttling mechanism, rotary valves rely on the rotation of a shaft to actuate the trim. An important advantage of rotary control valves over sliding-stem designs such as the globe valve and diaphragm valve is a virtually obstructionless path for fluid when the valve is wide-open³.



³Of course, gate valves also offer obstructionless flow when wide-open, but their poor throttling characteristics give most rotary valve designs the overall advantage.

27.2.1 Ball valves

In the ball valve design, a spherical ball with a passageway cut through the center rotates to allow fluid more or less access to the passageway. When the passageway is parallel to the direction of fluid motion, the valve is wide open; when the passageway is aligned perpendicular to the direction of fluid motion, the valve is fully shut (closed).

The following set of photographs shows a hand-operated ball valve in three different positions, from nearly full closed to nearly full open (left to right):



Simple ball valves with full-sized bores in the rotating ball are generally better suited for on/off service than for throttling (partially-open) service. A better design of ball valve for throttling service is the *characterized* or *segmented* ball valve, shown in various stages of opening in the following set of photographs:



The V-shaped notch cut into the opening lip of the ball provides a narrower area for fluid flow at low opening angles, providing more precise flow control than a plain-bore ball valve.

27.2.2 Butterfly valves

Butterfly valves are quite simple to understand: the “butterfly” element is a disk that rotates perpendicular to the path of fluid flow. When parallel to the axis of flow, the disk presents minimal obstruction; when perpendicular to the axis, the disk completely blocks any flow. Fluid-tight shut-off is difficult to obtain in the classic butterfly design unless the seating area is lined with a soft (elastic) material.

27.2.3 Disk valves

Disk valves (often referred to as *eccentric disk valves*, or as *high-performance butterfly valves*) are a variation on the butterfly design intended to improve seat shut-off. The disk's center is offset from the shaft centerline, causing it to approach the seat with a “cam” action that results in high seating pressure. Thus, tight shut-off of flow is possible even when using metal seats and disks.

The following photograph shows the body of a Fisher E-plug control valve, with the disk in a partially-open position:




27.3 Dampers and louvres

A *damper* (otherwise known as a *louvre*) is a multi-element flow control device generally used to throttle large flows of air at low pressure. Dampers find common application in furnace and boiler draft control, and in HVAC (Heating, Ventilation, and Air Conditioning) systems.

Common damper designs include parallel and radial. Parallel-vane dampers resemble a Venetian blind, with multiple rectangular vanes synchronously rotated to throttle flow through a rectangular opening. A photograph of a parallel-vane damper is shown here, part of an induced-draft (suction) air fan system on a separator at a cement plant. The vanes are not visible in this photograph because they reside inside the metal air duct, but the actuator mechanism and linkages connecting seven vane shafts together are:



Radial-vane dampers use multiple vanes arranged like petals of a flower to throttle flow through a circular opening. A photograph of a radial-vane damper is shown here (note the levers and linkages on the periphery of the tube, synchronizing the motions of the eight vanes so they rotate at the same angle):




Dampers find use in many non-industrial applications as well. Take for instance these greenhouse vents, actuated by pneumatic (air-powered) piston actuators:



27.4 Valve packing

Regardless of valve type, all stem-actuated control valves require some form of seal allowing motion of the stem from some external device (an *actuator*) while sealing process fluid so no leaks occur between the moving stem and the body of the valve. The general term for this sealing mechanism is *packing*.

This mechanical feature is not unlike the *stuffing box* used to seal seawater from entering a boat or ship at the point where the propeller shaft penetrates the hull:




The fundamental problem is the same for the ship as it is for the control valve: how to allow a moving shaft to pass through what needs to be an impenetrable barrier to some fluid (in the case of the ship's hull, the fluid is seawater). The solution is to wrap the shaft in a flexible material that maintains a close fit to the shaft without binding its motion. A traditional packing material for ship propeller shafts is *flax rope*. Some form of lubrication is usually provided so this packing material does not impose excessive friction on the shaft's motion⁴.

Modern marine stuffing boxes use advanced materials such as Teflon (PTFE) or graphite instead of flax, which wear longer and leak less water. In the world of control valves, the traditional packing material used to be asbestos (shaped into rings or ropes, much like flax used to be shaped for use in stuffing boxes), but is now commonly Teflon or graphite as well.

In the case of a ship's stuffing box, a little bit of water leakage is not a problem since all ships are equipped with bilge pumps to pump out collected water over time. However, leakage is simply unacceptable in many industrial control valve applications where we must minimize *fugitive emissions*. A "fugitive emission" is any unwanted escape of process substance into the surrounding environment, usually from leaks around pump and valve shafts. Special "environmental" packing sets are available for control valve applications where this is a concern.

⁴Some packing materials, most notably Teflon and graphite, tend to be *self-lubricating*.


Packing in a sliding-stem valve fits in a section of the valve body called the *bonnet*, shown in this simplified diagram of a single-ported, stem-guided globe valve:



Here, the packing material takes the form of several concentric rings, stacked on the valve stem like washers on a bolt. These packing rings are forced down from above by the *packing flange* to apply a compressive force around the circumference of the valve stem. This compressive force is necessary to generate mechanical stress in the packing material to make it seal tightly against the stem of the valve and the interior wall of the bonnet.

Two nuts threaded onto studs maintain proper force on the packing rings. Care must be taken not to over-tighten these nuts and over-compress the packing material, or else the packing will create excessive friction on the valve stem. Not only will this friction impede precise valve stem motion, but it will also create undue wear on the stem and packing, increasing the likelihood of future packing leakage. Insufficient packing flange force will lead to poor sealing, with process fluid potentially leaking past the packing and out of the valve.

A closer look at the bonnet shows a multitude of components working together to form a low-friction, pressure-tight seal for the moving valve stem:




In this diagram, we see two sets of packing rings separated by a metal piece called a *lantern ring*. The lantern ring acts as a spacer allowing lubricant introduced through the lubrication port to enter into both packing sets from the middle of the bonnet.

The packing shown here is “loaded” by the compressive force exerted by the packing follower. The only elasticity in this particular system resides in the packing material itself. This is called *stationary loading*, otherwise known as *jam* packing. Over time, as the packing material wears and fatigues, the packing follower must be re-compressed by carefully tightening the packing nuts.

One must be very careful when torquing the packing nuts on a stationary-loaded packing set. Insufficient torque (which translates into insufficient stress applied to the packing) will result in process fluid leakage. Excessive torque (causing excessive stress on the packing) will result in high valve stem friction and premature packing failure. The latter scenario is what one usually finds in industrial settings, where well-intentioned but uninformed personnel over-tighten valve packing in an effort to prevent leaks. The proper remedy for a packing assembly that leaks despite having been properly torqued is replacement, not further tightening.

An alternative to “stationary” loading is to insert a metal spring into the packing assembly, so that the elasticity of the spring helps to maintain an appropriate amount of packing stress as the packing material wears and ages. This is called *live loading*, examples of which are shown here:

Live-loaded valve stem packing



In one of these examples we see a coil spring inside the bonnet used to live-load the packing. In the other example we see a set of spring-steel washers known as a *Belleville* spring. Belleville springs have a concave profile, giving them resistance to compression along the shaft axis. These spring washers are always stacked in opposed pairs (concave against concave, convex against convex) so the washers have room to compress.

Photographs taken of an actual valve packing assembly removed from the bonnet (left), and re-assembled on the valve stem (right) reveal the structure of the packing and associated components.



There is no lantern ring in this packing assembly, but there is a coil spring. This makes it a *live-loaded* packing as opposed to a *jam* packing.

In packing applications requiring external lubrication, a *stem packing lubricator* may be connected to the lubrication port on the bonnet. This device uses a long, threaded bolt as a piston to push a quantity grease into the packing assembly:



To operate a lubricator, the hand valve on the lubricator is first secured in the closed (shut) position, then the bolt is fully unscrewed until it falls out of the lubricator body. An appropriate lubricating grease is squeezed into the bolt hole in the lubricator body, and the bolt threaded back into place until hand-tight. Using a wrench or socket to tighten the bolt a bit more (generating pressure in the grease) and opening the hand valve allows grease to enter the packing chamber. The bolt is then screwed in fully, pushing the entire quantity of grease into the packing. As a final step, the hand valve is fully shut so there is no way for process liquid to leak out past the bolt threads.

The two most common packing materials in use today are Teflon (PTFE) and graphite. Teflon is the better of the two with regard to fluid sealing, stem friction, and stem wear⁵. Teflon is also quite resistant to attack from a wide variety of chemical substances. Unfortunately, it has a limited temperature range and cannot withstand intense nuclear radiation (making it unsuitable for use near reactors in nuclear power plants). Graphite is another self-lubricating packing material, and it has a far greater temperature range than Teflon⁶ as well as the ability to withstand harsh nuclear radiation, but creates much more stem friction than Teflon. Graphite packing also has the unfortunate property of permitting *galvanic corrosion* between the stem and bonnet metals due to its electrical conductivity. Sacrificial zinc washers are sometimes added to graphic packing assemblies to help mitigate this corrosion, but this only postpones rather than prevents corrosive damage to the stem.

The following photographs show samples of woven graphite (left) and Teflon (right) “rope” packing, longer pieces of which would normally be found bent around valve stems to form seals. The graphite packing has a shiny finish and flakes easily, while the Teflon packing is plain white in color and maintains its integrity. Both feel slippery to the touch:



Hybrid packing materials, such as carbon-reinforced Teflon, have been developed in an attempt to combine the best characteristics of both materials.


A legacy valve packing material is *asbestos*, woven into packing rings much the same way as graphite fibers are woven into modern packing rings. Asbestos is a mineral, which made it suitable for high-temperature process applications. Its electrical non-conductivity eliminated the galvanic corrosion problem inherent to graphite. Unfortunately, its classification as a hazardous substance⁷ precludes its use as a packing material for contemporary applications.

⁵Based on friction values shown on page 131 of Fisher's *Control Valve Handbook* (Third Edition), Teflon packing friction is typically 5 to 10 times less than graphite packing for the same stem size!

⁶Graphite packing is usable in services ranging from cryogenic temperatures to 1200 degrees Fahrenheit, as opposed to Teflon which is typically rated between -40 °F and 450 °F.

⁷Asbestos fibers have the ability to permanently lodge in the air sacs of human lungs, leading to long-term health problems if those fibers are inhaled.

A completely different approach to packing is a device called a *bellows seal*: an accordion-like metal tube fastened to the valve stem and to the bonnet, forming a leak-proof seal with negligible friction. The accordion ribs give the bellows seal an ability to stretch and compress with a sliding stem's linear motion. Since the bellows is an uninterrupted metal tube, there is no place at all for leaks to develop:



The port on the extended bonnet serves as a point of connection for process fluid leak detection sensors, to sound an alarm and/or take action in the event of a ruptured bellows. The sensor may be as simple as a pressure switch, calibrated to “trip” at some modest pressure value below that of the normal process operating pressure. When the bellows seal breaks⁸, the sensor will detect the leak and the standard packing assembly will maintain a reasonable seal until repairs are made on

⁸Bellows have a limited service life, which means the possibility of a rupture is likely. This is why a conventional packing assembly is always included in a bellows-equipped bonnet.

the valve.

An actual bellows seal unit appears in this photograph:



The accordion-shaped bellows is contained and protected inside the thick metal tube visible in this photograph. One end of the bellows is welded to the valve stem, and the other end is welded to the protective tube. With the wide flange of the tube firmly clamped in the bonnet of the valve, a leak-free seal exists.


27.5 Valve seat leakage

In some process applications, it is important that the control valve be able to completely stop fluid flow when placed in the “closed” position. Although this may seem to be a fundamental requirement of any valve, it is not necessarily so. Many control valves spend most of their operating lives in a partially-open state, rarely opening or closing fully. Additionally, some control valve designs are notorious for the inability to completely shut off (e.g. double-ported globe valves). Given the common installation of manual “block” valves upstream and downstream of a control valve, there is usually a way to secure zero flow through a pipe even if a control valve is incapable of tight shut-off. For some applications, however, tight control valve shut-off is mandatory.

For this reason we have several classifications for control valves, rating them in their ability to fully shut off. Seat leakage tolerances are given roman numeral designations, as shown in this table⁹:

Class	Maximum allowable leakage rate	Test pressure drop
I	(no specification given)	(no specification given)
II	0.5% of rated flow capacity, air or water	45-60 PSI or max. operating
III	0.1% of rated flow capacity, air or water	45-60 PSI or max. operating
IV	0.01% of rated flow capacity, air or water	45-60 PSI or max. operating
V	0.0005 ml/min water per inch orifice size per PSI	Max. operating
VI	Bubble test, air or nitrogen	50 PSI or max. operating

The “bubble test” used for Class VI seat leakage is based on the leakage rate of air or nitrogen gas past the closed valve seat as measured by counting the rate of gas bubbles escaping a bubble tube submerged under water. For a 6 inch valve, this maximum bubble rate is 27 bubbles per minute (or about 1 bubble every two seconds):



It is from this leakage test procedure that the term *bubble-tight shut-off* originates. Class VI shut-off is often achievable only through the use of “soft” seat materials such as Teflon rather than

⁹Data in this table taken from Fisher's *Control Valve Handbook*.

hard metal-to-metal contact between the valve plug and seat. Of course, this method of achieving bubble-tight shut-off comes at the price of limited operating temperature range and the inability to withstand nuclear radiation exposure.

Special test fixtures are typically used in control valve rebuild shops to test the leakage rates of a rebuilt valve. One such test bench appears in this photograph:




In the foreground of this photograph we see a special “vise” used to make quick, pressure-tight connections to the flange of any control valve placed within it. A movable flange sandwiches the control valve against a stationary flange, both flanges faced with high-density plastic for a pressure-tight fit against the valve body flange faces. A stainless-steel panel in the background provides a set of air pressure regulators, hand valves, rotameters, pressure gauges, and even a “bubble” flow indicator to measure leakage flow rate through the control valve under varying pressure conditions.

27.6 Control valve actuators

The purpose of a control valve *actuator* is to provide the motive force to operate a valve mechanism. Both sliding-stem and rotary control valves enjoy the same selection of actuators: *pneumatic*, *hydraulic*, *electric motor*, and *hand* (manual).

27.6.1 Pneumatic actuators

Pneumatic actuators use air pressure pushing against either a flexible diaphragm or a piston to move a valve mechanism. The following photograph shows a cut-away control valve, with a pneumatic diaphragm actuator mounted above the valve body. You can see the large coil spring providing default positioning of the valve (air pressure acting against the diaphragm moves the valve against the spring) and the rubber diaphragm at the very top. Air pressure applied to the bottom side of the diaphragm lifts the sliding stem of the valve in the upward direction, against the spring's force which tries to push the stem down:



The amount of force (F) in units of pounds generated by any fluid pressing against any surface is equal to the fluid's pressure (P) in units of PSI multiplied by the surface area (A) in units of square inches ($F = PA$). In the case of a circular diaphragm, with area equal to πr^2 , the complete formula for force is $F = P\pi r^2$. For example, a control valve diaphragm 14 inches in diameter (radius = 7 inches) with an applied air pressure of 15 PSI generates a linear force of 2309 pounds.

Air pressure required to motivate a pneumatic actuator may come directly from the output of a pneumatic process controller, or from a *signal transducer* (or *converter*) translating an electrical signal into an air pressure signal. Such transducers are commonly known as *I/P* or "I to P"

converters, since they typically translate an electric current signal (I) of 4 to 20 mA DC into an air pressure signal (P) of 3 to 15 PSI.

The following photographs show I/P transducers of different make and model. A Fisher model 846 appears in the upper-left photograph, while an older Fisher model 546 appears in the upper-right (with cover removed). A Foxboro model E69F I/P appears in the lower-left photograph, while a Moore Industries model IPT appears in the lower-right:



Despite their differing designs and appearances, they all function the same: accepting an analog DC current signal input and a clean supply air pressure of about 20 PSI, outputting a variable air pressure signal proportional to the electric current input. An interesting feature to compare between these four I/P transducers is their relative ruggedness. Every transducer shown except the Moore Industries model (lower-right) is built to withstand direct exposure to a process atmosphere, hence the heavy cast-metal housings and electrical conduit fittings. The Moore Industries unit is intended for a sheltered location, and may be plugged in to a “manifold” with several other I/P transducers to form a compact bank of transducers capable of driving air pressure signals to several valve actuators.


Some pneumatic valve actuators are equipped with *handwheels* which are used to manually position the valve in the event of air pressure failure. The next photograph shows a sliding-stem control valve with pneumatic diaphragm actuator and a “handwheel” on the top:



Note the three manual valves located around the control valve: two to *block* flow through the control valve and one to *bypass* flow around the control valve in the event of control valve failure or maintenance. These manual valves happen to be of the *gate* design, with *rising-stem* actuators to clearly show their status (stem protruding = open valve ; stem hidden = closed valve). Such block-and-bypass manual valve arrangements are quite common in the process industries where control valves fulfill critical roles and some form of manual control is needed as an emergency alternative.

Note also the air pressure tubing between the valve actuator and the air supply pipe, bent into a loop. This is called a *vibration loop*, and it exists to minimize strain on the metal tubing from vibration that may occur.

Pneumatic actuators may take the form of pistons rather than diaphragms. Illustrations of each type are shown here for comparison:




Piston actuators generally have longer stroke lengths than diaphragm actuators, and are able to operate on much greater air pressures¹⁰. Since actuator force is a function of fluid pressure and actuator area ($F = PA$), this means piston actuators are able to generate more force than diaphragm actuators of the same diameter. A 14 inch diaphragm operating at a maximum pressure of 35 PSI generates 5388 pounds of force, but the same size piston operating at a maximum pressure of 150 PSI generates 23091 pounds of force. The combination of greater force and greater displacement yields more work potential for piston actuators than diaphragm actuators of equivalent size, since mechanical work is the product of force and displacement ($W = Fx$).

Diaphragm actuators enjoy the definite advantage of less friction than piston actuators. Less friction means greater precision in positioning the valve stem, which gives diaphragm actuators an advantage over piston actuators where precise valve positioning is important, all other factors being equal.

¹⁰The greater pressure rating of a piston actuator comes from the fact that the only “soft” component (the sealing ring) has far less surface area exposed to the high pressure than a rolling diaphragm. This results in significantly less stress on the elastic ring than there would be on an elastic diaphragm exposed to the same pressure. There really is no limit to the stroke length of a piston actuator as there is with the stroke length of a diaphragm actuator. It is possible to build a piston actuator miles long, but such a feat would be impossible for a diaphragm actuator, where the diaphragm must stretch (or roll) the entire stroke length.

The following photograph of an ultra-high pressure oxygen valve shows a large pneumatic piston actuating a relatively tiny valve body:




Since the only rationale for selecting such a large piston actuator is to generate large actuating force, we may conclude that this relatively small valve body requires an unusually high force to actuate. This is indeed the case, as the process fluid pressure drop across the valve trim in this application happens to be *several thousand PSI*. This great of a pressure differential, dropped across even a small valve plug, generates substantial force. The actuator must generate even more force than this in order to successfully move the valve, and must do so while limited to the typical instrument air pressure value of 100 PSI. Thus, the only way for the actuator to generate a superior force to the valve plug while working with much less fluid pressure is for the actuator piston to have a much greater area than the plug.

A double-piston pneumatic actuator appears in the next photograph, providing the mechanical force needed to turn an on/off butterfly valve:





In this particular actuator design, a pair of pneumatically-actuated pistons move a rack-and-pinion mechanism to convert linear piston motion into rotary shaft motion to move the butterfly trim. Note the rotary indicator (yellow in color) at the end of the rotary valve stem, showing what position the butterfly valve is in. Note also the travel switch box (black in color) housing multiple limit switches providing remote indication of valve position to the control room.

A rack-and-pinion mechanism looks like this, as viewed looking into the axis of the rotary shaft:



Compressed air applied to the bottom tube (with the top tube vented) pushes both pistons toward the center, spinning the pinion gear counter-clockwise. Applying compressed air to the top tube (with the bottom tube vented) pushes both pistons outward, rotating the pinion gear clockwise.

An example of this actuator design, cut away to reveal its inner workings, appears here:



Another pneumatic piston actuator design uses a simple crank lever instead of a rack-and-pinion gear set to convert linear piston motion into rotary motion. This next photograph shows such a piston actuator connected to a ball valve:



Perhaps the greatest disadvantage of piston actuators as applied to control valves is friction between the piston's pressure-sealing ring and the cylinder wall. This is not a problem for on/off control valves, but it may be a significant problem for throttling valves where precise positioning is desired. Diaphragm actuators do not exhibit the same degree of friction as piston actuators because the elastic diaphragm rolls and flexes rather than rubs against a stationary surface as is the case with piston sealing rings.

27.6.2 Hydraulic actuators

Hydraulic actuators use liquid pressure rather than gas pressure to move the valve mechanism. Nearly all hydraulic actuator designs use a piston rather than a diaphragm to convert fluid pressure into mechanical force. The high pressure rating of piston actuators lends itself well to typical hydraulic system pressures, and the lubricating nature of hydraulic oil helps to overcome the characteristic friction of piston-type actuators. Given the high pressure ratings of most hydraulic pistons, it is possible to generate tremendous actuating forces with a hydraulic actuator, even if the piston area is modest. For example, an hydraulic pressure of 2000 PSI applied to one side of a 3 inch diameter piston will generate a linear thrust exceeding 14000 pounds (7 tons)!

In addition to the ability of hydraulic actuators to easily generate extremely large forces, they also exhibit very *stable* positioning owing to the non-compressibility of hydraulic oil. Unlike pneumatic actuators, where the actuating fluid (air) is “elastic,” the oil inside a hydraulic actuator cylinder does not yield appreciably under stress. If the passage of oil to and from a hydraulic cylinder is blocked by small valves, the actuator will become firmly “locked” into place. This is an important feature for certain valve-positioning applications where the actuator must firmly hold the valve position in one position.

Some hydraulic actuators contain their own electrically-controlled pumps to provide the fluid power, so the valve is actually controlled by an electric signal. Other hydraulic actuators rely on a separate fluid power system (pump, reservoir, cooler, hand or solenoid valves, etc.) to provide hydraulic pressure on which to operate. Hydraulic pressure supply systems, however, tend to be more limited in physical span than pneumatic distribution systems due to the need for thick-walled tubing (to contain the high oil pressure), the need to purge the system of all gas bubbles, and the problem of maintaining a leak-free distribution network. It is usually not practical to build a hydraulic oil supply and distribution system large enough to cover the entirety of an industrial facility. Another disadvantage of hydraulic systems compared to pneumatic is lack of intrinsic power storage. Compressed air systems, by virtue of air’s compressibility (elasticity), naturally store energy in any pressurized volumes, and so provide a certain degree of “reserve” power in the event that the main compressor shut down. Hydraulic systems do not naturally exhibit this desirable trait.

A hydraulic piston actuator attached to a large shut-off valve (used for on/off control rather than throttling) appears in the next photograph. Two hydraulic cylinders may be seen above the round valve body, mounted horizontally. Like the pneumatic piston valve shown earlier, this valve actuator uses a rack-and-pinion mechanism to convert the hydraulic pistons' linear motion into rotary motion to turn the valve trim:



A feature not evident in this photograph is a hydraulic hand pump that may be used to manually actuate the valve in the event of hydraulic system failure.


27.6.3 Self-operated valves

Although not a type of actuator itself, a form of actuation worthy of mention is where the process fluid pressure itself actuates a valve mechanism. This self-operating principle may be used in throttling applications or on/off applications, in gas or liquid services alike. The process fluid may be directly tubed to the actuating element (diaphragm or piston), or passed through a small mechanism called a *pilot* to modulate that pressure before reaching the valve actuator. This latter design allows the main valve's motion to be controlled by an adjustable device (the pilot).

A very common application for pilot-operated control valves is gas pressure regulation, especially for fuel gas such as propane or natural gas used to fuel large industrial burners. This next photograph shows a Fisher gas pressure regulator used for regulating the pressure of natural gas fueling an industrial burner:




The following diagram shows how a self-operated, spring-loaded gas pressure regulating valve functions:



A spring tries to force the plug off the seat, while “feedback” gas pressure from the downstream side of the valve acts against a flexible diaphragm to move the plug toward the seat. The less

downstream pressure, the more the trim opens up; the more downstream pressure, the more the trim shuts off. This spring establishes the pressure-regulating “setpoint” value for the regulator. If a different setpoint is desired, the spring must be replaced with one having a different stiffness.

A variation on this theme is the *pilot-loaded* or *externally-loaded* pressure regulator, using a source of external gas pressure to establish the pressure regulation setpoint. Here, a simple manual-adjustment pressure regulator serves as the “pilot” device to send a loading pressure to the top of the main regulator’s actuating diaphragm:



Instead of a stiff internal spring establishing the regulator’s pressure setpoint, the externally-supplied loading pressure does that. Since this loading pressure is easily adjusted by turning the knob on the manual-set pressure regulator, the main regulator now becomes adjustable as well. The pilot mechanism controls the main gas throttling mechanism, hence the name *pilot*.

This next pilot-operated valve is used in a liquid (wastewater) service rather than gas:



A consumer-grade application for pilot-operated valves is irrigation system control, where the solenoid valves used to switch water flow on and off to sprinkler heads use pilot mechanisms rather than operate the valve mechanism directly with magnetic force. A small solenoid valve opens and closes to send water pressure to an actuating diaphragm, which then operates the larger valve mechanism to start and stop the flow of water to the sprinkler. The use of a pilot allows a relatively small amount of electrical power to control the valve, compared to the amount of electrical power that would be necessary if the solenoid coil were built large enough to actuate the main water valve directly.


A special case of self-operated valve is the *Pressure Relief Valve (PRV)* or *Pressure Safety Valve (PSV)*. These valves are normally shut, opening only when sufficient fluid pressure develops across them to relieve that process fluid pressure and thereby protect the pipes and vessels upstream. Like the other self-operated valves, these safety valves may directly actuate using process fluid pressure or they may be triggered by a pilot mechanism sending process fluid pressure to the actuator only above certain pressures. Relief valves using pilots have the advantage of being widely adjustable, whereas non-pilot safety valves usually have limited adjustment ranges.

For more information on overpressure protection devices (including PRVs and PSVs) refer to section [32.5](#) beginning on page [2655](#).

27.6.4 Electric actuators

Electric motors have long been used to actuate large valves, especially valves operated as on/off (“shutoff”) devices. Advances in motor design and motor control circuitry have brought motor-operated valve (MOV) technology to the point where it now competes with legacy actuator technologies such as pneumatic in actuating *throttling* valves as well.

Most electric valve actuators use a *worm gear* set to reduce the high rotational speed of the electric motor to a slow rotation suitable for moving a large valve mechanism. An illustration of a worm gear set appears here:



The *worm screw* looks much like a threaded fastener, with its “threads” properly pitched to engage with the teeth of the *worm wheel* gear. As the worm screw turns, it slowly pushes or pulls the circumference of the worm wheel, resulting in a large gear ratio (i.e. many turns of the screw are required to produce a single turn of the wheel). This slow-turning wheel may then be used to move a sliding-stem valve by means of a threaded shaft (another screw) or used to directly turn a rotary valve (e.g. butterfly, ball, plug).

An electric actuator appears in the next photograph, providing on/off rotary actuation to a ball valve. This particular electric actuator comes with a hand crank for manual operation, in the event that the electric motor (or the power provided to it) fails:



A small lever to the left of the hand crank actuates a *clutch* mechanism to engage or disengage the valve mechanism from the electric motor and the hand wheel. This clutch “selects” either the motor or the hand wheel as the prime mover for the valve, to avoid having the hand wheel spin as the motor turns. Unless this lever is first moved to the “manual” position, turning the hand wheel accomplishes nothing.

The next photograph shows an electric valve actuator coupled to a large butterfly valve. Although nothing visible in this photograph betrays the nature of this actuator's signaling, it happens to be *digital* rather than analog, receiving position commands through a *Profibus* digital network rather than an analog 4-20 mA current signal:



The shape of the actuator's metal casing reveals the final gear drive of this actuator as a *worm gear set*. Note the round shape of the casing where the open/close indicator is located (at the centerline of the butterfly valve stem): this contains the worm wheel. Note immediately to the left of that round casing is a vertically-oriented cylinder shape: this contains the worm screw, which engages with the teeth of the worm wheel. Above the worm screw is a parallel-shaft *spur gear* set which acts to further reduce the speed of the actuator's electric motor: the motor shaft terminates in a small gear, which meshes with a larger gear that turns the worm screw shaft. Above the parallel-shaft gear set is the main casing of the actuator, which actually contains its own internal worm gear set. This multiple-stage gear reduction means that the motor spins very fast in comparison to the butterfly element inside the valve, and that the butterfly is capable of exerting a fantastic amount of torque in comparison to the torque rating of the electric motor.

Electric motors require no external fluid power system to function, unlike pneumatic or hydraulic actuators. All they require is a source of electrical power (often 480 volts AC, three-phase). Some electric valve actuators even have the capability of operating from the power of an electric battery pack, for reliable operation in the event of a power system outage.

Virtually all electric valve actuators require some form of *feedback* to indicate the valve's position.

At minimum, this consists of *limit switches* to indicate when the valve is fully shut and fully open. For throttling services, an electric actuator requires an actual valve position sensor so that it may precisely adjust the valve to any desired state. This sensor may take the form of a potentiometer, or a variable differential transformer (LVDT or RVDT), or pulse encoder.


This Rotork brand MOV has a digital display at one end showing its closed status both in text (“Closed Limit”) and symbolically (by the vertical line, which is supposed to represent a closed butterfly element):



In addition to visual indicators of status, electric valve actuators commonly provide auxiliary electrical contacts signaling fully-open and fully-closed positions, which may be used to energize remote indicator lights or discrete input channels of control systems (e.g. PLC). Throttling-service MOVs also provide analog (and/or digital) signaling of valve stem position for remote indication or feedback to an electronic control system.

27.6.5 Hand (manual) actuators

Valves may also be actuated by hand power alone. The following valves are all “manual” valves, requiring the intervention of a human operator to actuate:



Note the threaded stem of the lower-left valve. This stem rises and falls with the handle's turning, providing visual indication of the valve's status. Such an actuator is called a *rising-stem* design. Also note the chain on the upper-right actuator wheel, allowing operation of the high-mounted valve from a ground-level position.

A hybrid of hand and pneumatic valve actuation is seen on this Valtek brand control valve, where the control valve assembly is actuated by a pneumatic piston actuator, but is also equipped with a manually-operated “handwheel”:



A handwheel mechanism may be used to override the pneumatic actuator simply by overpowering it in either direction (i.e. providing a greater force on the valve stem than the piston actuator provides) or it may be left in a neutral position to allow the pneumatic actuator full control over valve stem position. Handwheels may be used to override the control valve’s pneumatic actuator to either the full-open or full-closed positions, or it may simply be used to assert a high- or low-limit “stop” to the valve stem to prohibit stem motion beyond a certain position.

Note the “lockout” tab flipped to the horizontal position near the handwheel, located between two of the handwheel’s spokes. This simple mechanism permits the handwheel to be locked out to prevent accidental turning.


27.7 Valve failure mode

An important design parameter of a control valve is the position it will “fail” to if it loses motive power. For electrically actuated valves, this is typically the last position the valve was in before loss of electric power. For pneumatic and hydraulic actuated valves, the option exists of having a large spring provide a known “fail-safe” position (either open or closed) in the event of fluid pressure (pneumatic air pressure or hydraulic oil pressure) loss.


27.7.1 Direct/reverse actions

The fail-safe mode of a pneumatic/spring valve is a function of both the actuator's action and the valve body's action. For sliding-stem valves, a *direct-acting* actuator pushes down on the stem with increasing pressure while a *reverse-acting* actuator pulls up on the stem with increasing pressure. Sliding-stem valve bodies are classified as *direct-acting* if they open up when the stem is lifted, and classified as *reverse-acting* if they shut off (close) when the stem is lifted. Thus, a sliding-stem, pneumatically actuated control valve may be made *air-to-open* or *air-to-close* simply by matching the appropriate actuator and body types.

The most common combinations mix a direct-acting valve body with either a reverse- or direct-acting valve actuator, as shown in this illustration:




Reverse-acting valve bodies may also be used, with opposite results:



The reverse-acting gate valve body shown in the left-hand illustration is open, with fluid flowing *around* the stem while the wide plug sits well below the seat area. Reverse-acting valve bodies tend to be more complex in construction than direct-acting valve bodies, and so they are less common in control valve applications. An interesting exception to this trend – although not technically a *control* valve but rather a self-actuated device – is the Fisher model 1098EGR pilot-operated pressure regulator which uses a reverse-acting valve body to throttle the flow of gas through it.

27.7.2 Available failure modes

Valve fail mode may be shown in instrument diagrams by either an arrow pointing in the direction of failure (assuming a direct-acting valve body where stem motion toward the body closes and stem motion away from the body opens the valve trim) and/or the abbreviations “FC” (fail closed) and “FO” (fail open). Other failure modes are possible, as indicated by this set of valve symbols:



In order for a pneumatic or hydraulic valve to fail in the *locked* state, an external device must trap fluid pressure in the actuator’s diaphragm or piston chamber in the event of supply pressure loss.


Valves that fail in place and drift in a particular direction are usually actuated by double-acting pneumatic piston actuators. These actuators do not use a spring to provide a definite fail mode, but rather use air pressure both to open and to close the valve. In the event of an air pressure loss, the actuator will neither be able to open nor close the valve, and so it will tend to remain in position. If the valve is of the globe design with unbalanced trim, forces exerted on the valve plug will move it in one direction (causing drift).

27.7.3 Selecting the proper failure mode

It is important to note how the failure mode of a valve is often linked to its control action (air-to-open, air-to-close)¹¹. That is, an air-to-open pneumatic control valve will fail closed on loss of air pressure, and vice-versa. This is an important fact because good safety engineering demands that the risk factors of the process determine proper valve failure mode rather than control system convention or habit. People may find it easier to understand the operation of an air-to-open control valve than an air-to-close valve (more signal = more process fluid flow), but this should not be a guiding principle in valve selection. Air-to-open control valves naturally fail closed which means they are appropriate for a particular process control application *only* if that process is safer with a failed-closed valve than with a failed-open valve. If the process is safer with a fail-open valve, then the pneumatically-actuated control valve specified for that application needs to be air-to-close.

In fact, this basic principle forms the basis – or at least it *should* form the basis – of decisions made for all instrument actions in critical control loops: first determine the safest mode of valve failure, then select and/or configure instrument actions in such a way that the most probable modes of signal path failure will result in the control valve consistently moving to that (safest) position.

For example, consider this automated cooling system for a large power-generating engine:



Clearly, it is more hazardous to the engine for the valve to fail closed than it would be for the valve to fail open. If the valve fails closed, the engine will surely overheat from lack of cooling. If it fails open, the engine will merely run cooler than designed, the only negative consequence being decreased efficiency. With this in mind, the only sensible choice for a control valve is one that fails open (air-to-close).

However, our choices in instrument action do not end with the control valve. How should the temperature transmitter, the controller, and the I/P transducer be configured to act? In each case,

¹¹Exceptions exist for valves designed to fail in place, where a valve may be engineered to “lock” in position through the action of an external device whether the valve itself is air-to-open or air-to-close.

the answer should be to act in such a way that the valve will default to its fail-safe position (wide open) in the event of the most likely input signal fault.

Stepping “backward” through the control system from the valve to the temperature sensor, the next instrument we encounter is the I/P transducer. Its job, of course, is to convert a 4-20 mA current signal into a corresponding pneumatic pressure that the valve actuator can use. Since we know that the valve’s failure mode is based on a loss of actuating air pressure, we want the I/P to be configured in such a way that it outputs minimum pressure in the event of an electrical fault in its 4-20 mA input signal wiring. Whether the wiring fails shorted or fails open, the result will be 0 milliamps at the I/P input terminals. Thus, the configuration of the I/P transducer should be *direct*, such that a 4 to 20 mA input signal produces a 3 to 15 PSI output pressure, respectively (i.e. minimum input current yields minimum output pressure).

The next instrument in the loop is the controller. Here, we want the most likely input signal failure to result in a minimum output signal, so the valve will (once again) default to its “fail safe” position. Consequently, we should configure the controller for *direct* action just like we did with the I/P transducer (i.e. a decreasing PV signal from a broken wire or loose connection in the input circuit results in a decreasing output signal).

Finally, we come to the last instrument in the control loop: the temperature transmitter (TT). As with most instruments, we have the option of configuring it for direct or reverse action. Should we choose direct (i.e. hotter engine = more mA output) or reverse (hotter engine = less mA output)? Here, our choice needs to be made in such a way that the overall effect of the control system is *negative feedback*. In other words, we need to configure the transmitter such that a hotter engine results in increased coolant flow (a wider-open control valve). Since we know the rest of the system has been designed so a minimum signal anywhere tends to drive the valve to its fail-safe mode (wide open), we must choose a *reverse-acting* transmitter, so a hotter engine results in a decreased milliamp signal from the transmitter. If the transmitter has a sensor “burnout” mode switch, we should flip this switch into the low-scale burnout position, so a burned-out sensor will result in 4 mA output (low end of the 4-20 mA scale), thus driving the valve into its safest (wide-open) position.

Such a configuration – with its air-to-close control valve and a reverse-acting transmitter – may seem strange and counter-intuitive, but it is the safest design for this engine cooling system. We arrived at this “odd” configuration of instruments by first choosing the safest control valve failure mode, then choosing instrument actions in such a way that the most likely signal-path failures anywhere in the system would result in the same, consistent valve response. Of course it should go without saying that accurate documentation in the form of a loop diagram with instrument actions clearly shown is an absolutely essential piece of the whole system. If the safety of a control system depends on using any “non-standard” instrument configurations, those configurations had better be documented so those maintaining the system in the future will know what to expect!

Another important detail in this system is to configure the controller such that the operator display for the output signal still registers in an intuitive way: 0% should still represent a shut control valve, while 100% should still represent a wide-open valve. With the valve being air-to-close (signal-to-close from the controller’s perspective), this means the controller should be configured for *reverse indication*¹² on the output display, so that an output of 4 mA (wide-open valve) reads 100% open, and an output of 20 mA (fully shut valve) reads 0%. As confusing as this might be for

¹²Note that reverse *indication* is not the same thing as reverse *action* for a loop controller. Reverse indication simply means the output display shows 100% valve position at 4 mA output, and 0% valve position at 20 mA output. Reverse action means the output decreases when the input (process variable) increases.

the technician who must service the controller, it is more important that the operator using this controller every working day sees something that makes intuitive sense. “Minor” details such as this become critically important if an emergency ever occurs, and the operator must make split-second decisions based on the indications they see!

27.8 Actuator bench-set

Valve actuators provide force to move control valve trim. For precise positioning of a control valve, there must be a calibrated relationship between applied force and valve position. Most pneumatic actuators exploit *Hooke’s Law* to translate applied air pressure to valve stem position.

$$F = kx$$

Where,

F = Force applied to spring in newtons (metric) or pounds (British)

k = Constant of elasticity, or “spring constant” in newtons per meter (metric) or pounds per foot (British)

x = Displacement of spring in meters (metric) or feet (British)


Hooke’s Law is a linear function, which means that spring motion will be linearly related to applied force from the actuator element (piston or diaphragm). Since the working area of a piston or diaphragm is constant, the relationship between actuating fluid pressure and force will be a simple proportion ($F = PA$). By algebraic substitution, we may alter Hooke’s Law to include pressure and area:

$$F = kx$$

$$PA = kx$$

Solving for spring compression as a function of pressure, area, and spring constant:


$$x = \frac{PA}{k}$$



When a control valve is assembled from an actuator and a valve body, the two mechanisms must be coupled together in such a way that the valve moves between its fully closed and fully open positions with an expected range of air pressures. A common standard for pneumatic control valve actuators is 3 to 15 PSI¹³.

¹³3 PSI could mean fully closed and 15 PSI fully open, or vice-versa, depending on what form of actuator is coupled to what form of valve body. A direct-acting actuator coupled to a direct-acting valve body will be open at low pressure and closed at high pressure (increasing pressure pushing the valve stem toward the body, closing off the valve trim),


There are really only two mechanical adjustments that need to be made when coupling a pneumatic diaphragm actuator to a sliding-stem valve: the *stem connector* and the *spring adjuster*. The stem connector mechanically joins the sliding stems of both actuator and valve body so they move together as one stem. This connector must be adjusted so neither the actuator nor the valve trim prevents full travel of the valve trim:



Note how the plug rests fully on the seat when the valve is closed, and how the travel indicator indicates fully open at the point where the actuator diaphragm nears its fully upward travel limit. This is how things should be when the stem connector is properly adjusted.


resulting in *air-to-close* action. Reversing either actuator or valve type (e.g. reverse actuator with direct valve or direct actuator with reverse valve) will result in *air-to-open* action.

If the stem connector is set with the actuator and valve stems spaced too far apart (i.e. the total stem length is too long), the actuator diaphragm will bind travel at the upper end and the valve plug will bind travel at the lower end. The result is a valve that cannot ever fully open:



A control valve improperly adjusted in this manner will never achieve full-flow capacity, which may have an adverse impact on control system performance.

If the stem connector is set with the actuator and valve stems too closely coupled (i.e. the total stem length is too short), the actuator diaphragm will bind travel at the lower end and the valve plug will bind travel at the upper end. The result is a valve that cannot ever fully close:



This is a very dangerous condition: a control valve that lacks the ability to fully shut off. The process in which this valve is installed may be placed in jeopardy if the valve lacks the ability to stop the flow of fluid through it!


Once the stem length has been properly set by adjusting the stem connector, the spring adjuster must be set for the proper *bench set* pressure. This is the pneumatic signal pressure required to lift the plug off the seat. For an air-to-open control valve with a 3 to 15 PSI signal range, the “bench set” pressure would be 3 PSI.

Bench set is a very important parameter for a control valve because it establishes the seating force (*seat load*) of the plug against the seat when the valve is fully closed. Proper seating pressure is critical for tight shut-off, which carries safety implications in some process services. Consult the manufacturer’s instructions when adjusting the bench set pressure for any sliding-stem control valve.

These instructions will typically guide you through both the stem connector and the spring adjuster procedures, to ensure both parameters are correctly set.


27.9 Pneumatic actuator response

A limitation inherent to pneumatic valve actuators is the amount of air flow required to or from the actuator to cause rapid valve motion. This is an especially acute problem in all-pneumatic control systems, where the distance separating a control valve from the controller may be substantial:



The combined effect of air-flow friction in the tube, flow limitations inherent to the controller mechanism, and volume inside the valve actuator conspire to create a sluggish valve response to sudden changes in controller output signal, not unlike the response of an RC (resistor-capacitor) time-delay circuit where a step-change in voltage input results in an inverse exponential output signal.

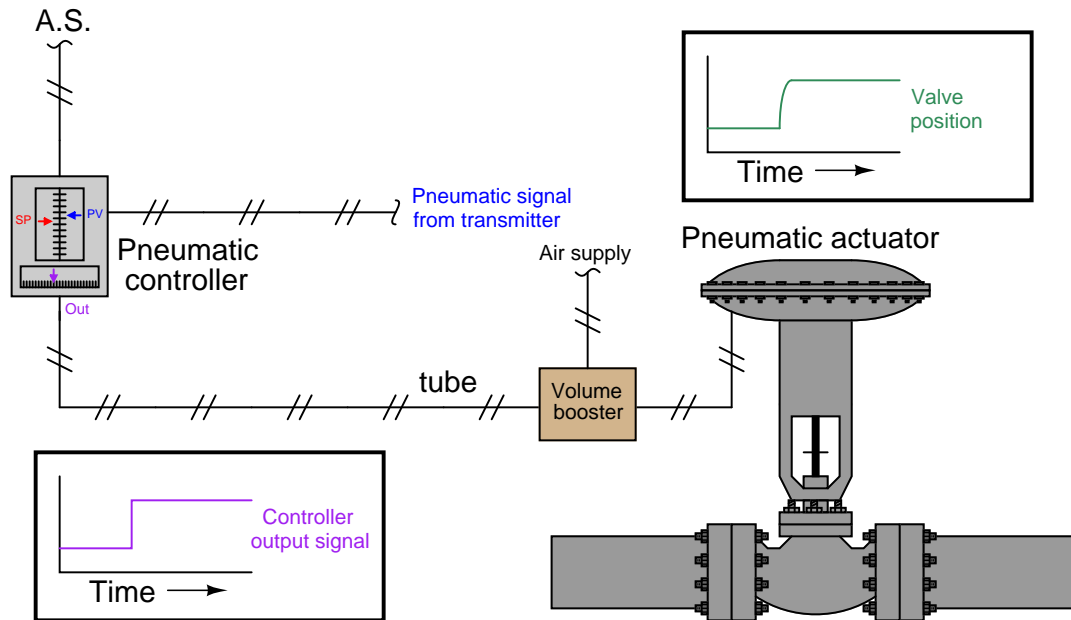
If the pneumatic valve actuator is driven by an I/P transducer instead of directly by a pneumatic controller, the problem is lessened by the ability to locate the I/P close to the actuator, thus greatly minimizing tube friction and thus minimizing the “time constant” (τ) of the control valve’s response:



Still, if the pneumatic actuator is particularly large in volume, an I/P transducer may experience trouble supplying the necessary air flow rate to rapidly actuate the control valve. Certainly the problem of time delay is reduced, but not eliminated, by the close-coupled location of the I/P transducer to the actuator.


One way to improve valve response in either type of system (full-pneumatic or I/P-driven) is to use a device known as a *volume booster* to source and vent compressed air for the valve actuator. A “volume booster” is a pneumatic device designed to reproduce a pneumatic pressure signal (1:1 ratio), but with far greater output flow capacity. In electrical terms, a volume booster is analogous to a *voltage follower*: a circuit designed to boost current to a load, without boosting or diminishing voltage. A 3 to 15 PSI pneumatic pressure signal applied to the input of a volume booster will result in an identical output signal (3 to 15 PSI), but with greatly enhanced flow capacity.

A pneumatic control system equipped with a volume booster would look something like this:



Of course, enhanced air flow to and from the actuator does not completely eliminate time delays in valve response. So long as the flow rate into or out of an actuator is finite, some time will be required to change pressure inside the actuator and thus change valve position. However, if the actuator volume cannot be reduced for practical reasons of actuating force (larger diaphragm or piston area needed for more force, also resulting in more volume for any given stroke length), then the only variable capable of reducing time lag is increased air flow rate, and a volume booster directly addresses that deficiency.

Internally, a volume booster's construction is not unlike a manually-adjusted pressure regulator¹⁴:



In either mechanism, an internal diaphragm senses output pressure and acts against a restraining force (either a spring preloaded by a hand adjustment screw or an external pressure signal acting on another diaphragm) to position an air flow throttling/venting mechanism. If the output pressure is less than desired, the diaphragm moves down to open the air sourcing plug and supply additional air volume to the output. If the output pressure is greater than desired, the diaphragm moves up to shut off the sourcing plug and open up the venting port to relieve air pressure to atmosphere.

¹⁴The volume booster design shown here is loosely based on the Fisher model 2625 volume boosting relay.


27.10 Valve positioners

The reason why a pneumatic control valve's stem position corresponds linearly to the amount of air pressure applied to the actuator is because mechanical springs tend to follow *Hooke's Law*, where the amount of spring motion (x) is directly proportional to applied force ($F = kx$). A pneumatic actuator applies force as a function of air pressure and piston/diaphragm area ($F = PA$), and the spring in turn compresses or stretches to generate an equal and opposite reaction force. The end-result is that actuator pressure linearly translates into valve stem motion ($x = \frac{PA}{k}$).

This linear and repeatable relationship between pneumatic signal pressure and valve stem position holds true if and only if the actuating diaphragm/piston and spring are the sole forces at work on the valve stem. If any *other* force acts upon this mechanism, the relationship between signal pressure and valve stem position will no longer be ideal.

Unfortunately, there exist many other forces acting on a valve stem besides the actuator force and the spring's reaction force. Friction from the stem packing is one such force, and reaction force at the valve plug caused by differential pressure across the plug's area is another¹⁵. These forces conspire to re-position the valve stem so stem travel does not precisely correlate to actuating fluid pressure.

A common solution to this dilemma is to add a *positioner* to the control valve assembly. A positioner is a motion-control device designed to actively compare stem position against the control signal, adjusting pressure to the actuator diaphragm or piston until the correct stem position is reached:



Positioners essentially act as control systems within themselves¹⁶: the valve's stem position is the process variable (PV), the command signal to the positioner is the setpoint (SP), and the positioner's signal to the valve actuator is the manipulated variable (MV) or output. Thus, when a process controller sends a command signal to a valve equipped with a positioner, the positioner receives that command signal and applies as much or as little air pressure to the actuator as needed

¹⁵One way to minimize dynamic forces on a globe valve plug is to use a *double-ported* plug design, or to use a *balanced* plug on a cage-guided globe valve. A disadvantage to both these valve plug designs, though, is greater difficulty achieving tight shut-off.

¹⁶The technical term for this type of control system is *cascade*, where one controller's output becomes the setpoint for a different controller. In the case of a valve positioner, the positioner receives a valve stem position setpoint from the main process controller. We could say that the main process controller in this case is the *primary* or *master* controller, while the valve positioner is the *secondary* or *slave* controller.

in order to achieve that desired stem position. Thus, the positioner will “fight” against any other forces acting on the valve stem to achieve crisp and accurate stem positioning according to the command signal. A properly functioning positioner ensures the control valve will be “well-behaved” and obedient to the command signal.

The following photograph shows a Fisher model 3582 pneumatic positioner mounted to a control valve. The positioner is the grey-colored box with three pressure gauges on its right-hand side:



On the left-hand side of this positioner may be seen part of the feedback mechanism: a metal bracket bolted to the valve stem connector, linking to an arm coming out of the positioner’s side. Every control valve positioner must be equipped with some means to sense the position of the valve’s stem, otherwise the positioner could not compare the valve stem’s position against the command signal.

A more modern positioner appears in the next photograph, the Fisher DVC6000 (again, the grey-colored box with pressure gauges on its right-hand side):



Like the older model 3582 positioner, this DVC6000 model uses a feedback linkage on the left-hand side to sense the valve stem's position. The even newer model DVC6200 uses a magnetic *Hall Effect* sensor to sense the position of a magnet bolted to the valve stem. This non-mechanical position feedback design eliminates backlash, wear, interference, and other potential problems associated with mechanical links. Better feedback is essential to better valve positioning.

Control valve positioners are typically constructed in such a way to source and vent high air flow rates, such that the positioner also fulfills the functionality of a volume booster¹⁷. Thus, a positioner not only ensures more precise valve stem positioning, but also faster stem velocity (and shorter time delays) than if the valve actuator were directly “powered” by an I/P transducer.

¹⁷This is not to say valve positioners have no need for external volume boosters, just that the actuating air flow capacity of a typical positioner greatly exceeds the air flow capacity of a typical I/P transducer.


Another advantage of adding a positioner to a pneumatically actuated control valve is superior valve seating (tight shutoff). This benefit is not obvious at first inspection, and so some explanation is in order.

First, one must understand that mere contact between the plug and seat within a sliding-stem valve is not enough to ensure tight shut-off. Rather, the plug must be *forcefully* pressed down onto the seat in order to fully shut off all flow through the valve. Anyone who has ever tightened the handle on a leaking hose bib (garden spigot) intuitively understands this principle: a certain amount of contact force between the plug and the seat is necessary in order to slightly deform and thereby mold those two components to a perfect fluid-tight fit. The technical term for this mechanical requirement is *seat load*.


Imagine if you will a diaphragm-actuated, sliding-stem, air-to-open control valve with a bench set range of 3 to 15 PSI. At an applied actuator pressure of 3 PSI, the diaphragm generates just enough force to exactly overcome the actuator spring's pre-load force, but not enough force to actually move the plug off the seat. In other words, at 3 PSI diaphragm pressure, the plug is touching the seat but with little or no force to provide a tight shut-off seal. If this control valve is directly powered by an I/P transducer with a 3-15 PSI calibrated range, it means the valve will be barely shut at a 0% signal value (3 PSI) rather than tightly shut off. In order to fully force the valve plug against the valve seat to achieve a tight seal, *all* air pressure would have to be vented from the diaphragm to ensure no diaphragm force opposing the spring. This is impossible with an I/P having a calibrated range of 3-15 PSI.

Now imagine that exact same valve equipped with a positioner, taking the 3-15 PSI signal from the I/P and using it as a command (setpoint) for valve stem position, applying as much or as little pressure to the diaphragm as necessary to achieve the desired stem position. Proper positioner calibration is such that the valve stem does not begin to lift until the signal has risen slightly *above* 0%, which means at 0% (4 mA) the positioner will be trying to force the valve to a slightly *negative* stem position. In attempting to achieve this impossible demand, the positioner's output will saturate low, applying no pressure whatsoever to the actuating diaphragm, resulting in full spring force applied by the plug against the seat. A comparison of the two scenarios is shown here:

*Air-to-open control valve at 0% signal,
positioner versus non-positioner*



While positioners are beneficial on spring-equipped valve actuators, they are absolutely essential for some other styles of actuators. Consider the following double-acting pneumatic piston actuator which has no spring:




Without a spring providing a restraining force to return the valve to a “fail-safe” position, there exists no Hooke’s Law relationship between applied air pressure and stem position. A positioner must alternately apply air pressure to *both* surfaces of the piston to raise and lower the valve stem.

Electric control valve actuators are another class of actuator design absolutely requiring some form of positioner system, because an electric motor is not “aware” of its own shaft position in order that it may precisely move a control valve. Thus, a positioner circuit using a potentiometer or LVDT/RVDT sensor to detect valve stem position and a set of transistor outputs to drive the motor is necessary to make an electric actuator responsive to an analog control signal.

27.10.1 Force-balance pneumatic positioners

A simple force-balance pneumatic valve positioner design appears in the following cutaway illustration:



The control signal for this valve is a 3 to 15 PSI pneumatic signal, coming from either an I/P

transducer or a pneumatic controller (neither one shown in the illustration). This control signal pressure applies an upward force on the force beam, such that the baffle tries to approach the nozzle. Increasing backpressure in the nozzle causes the pneumatic amplifying relay to output a greater air pressure to the valve actuator, which in turn lifts the valve stem up (opening up the valve). As the valve stem lifts up, the spring connecting the force beam to the valve stem becomes further stretched, applying additional force to the right-hand side of the force beam. When this additional force balances the bellows' force, the system stabilizes at a new equilibrium.


Like all force-balance systems, the force beam motion is constrained by the balancing forces, such that its motion is negligible for all practical purposes. In the end, equilibrium is achieved by one force balancing another, like two teams of people pulling oppositely on a length of rope¹⁸: so long as the two teams' forces remain equal in magnitude and opposite in direction, the rope will not deviate from its original position.

¹⁸In an earlier chapter of this book, force- and motion-balance pneumatic mechanisms were likened to “tug-of-war” contestants versus ballroom dancers, respectively. Force-balance mechanisms pit force against force to achieve mechanical balance, like two teams competing in a tug-of-war where opposing forces are perfectly balanced and no motion takes place. Motion-balance mechanisms match one motion with another motion to achieve mechanical balance, like two ballroom dancers moving across a dance floor while maintaining a constant distance between each other. All valve positioner mechanisms require motion on the part of the valve stem, and so it is natural to assume all valve positioner mechanisms will be motion-balance because unlike a tug-of-war something is definitely moving. However, if we examine the simple force-balance positioner mechanism closely we will see that *only* the valve stem moves in this mechanism, while nothing else does. To apply the tug-of-war analogy to this application, it is as if one team of contestants pulls on a stiff rope while the other team pulls on an elastic rope, the two ropes tied together in a knot at the center. In order to achieve a perfect balance of forces so the knot won't move to one side or the other, the team holding the elastic rope must *stretch their rope* further in order to balance an increased force from the team holding the stiff rope. The fact that one team is moving does not negate the fact that balance between the two teams is still a matter of force against force. To illustrate this point more vividly, we may ask the question: if the elastic rope is replaced by one that is even more elastic than before, will it advantage one team of contestants over the other? The answer to this question is *no*, as the two teams will still be equally matched if they were equally matched before. The only difference now is that the team holding the elastic rope will have to stretch the rope further than before to apply the same force as before. In a true motion-balance system, a greater motion imparted by one portion of the mechanism must be matched by a greater motion in the other portion of the mechanism.


The following photograph shows a PMV model 1500 force-balance positioner used to position a rotary valve actuator, with the cover on (left) and removed (right):



The 3-15 PSI pneumatic control signal enters into the bellows, pushing downward on the horizontal force beam (colored black). A pneumatic pilot valve assembly at the left-hand side of the force beam detects any motion, increasing air pressure to the valve actuating diaphragm if any downward motion is detected and releasing air pressure from the actuator if any upward motion is detected:



As compressed air is admitted to the valve actuator by this pilot valve assembly, the rotary valve will begin to rotate in the open direction. The shaft's rotary motion is converted into a linear motion inside the positioner by means of a *cam*: a disk with an irregular radius designed to produce linear displacement from angular displacement:



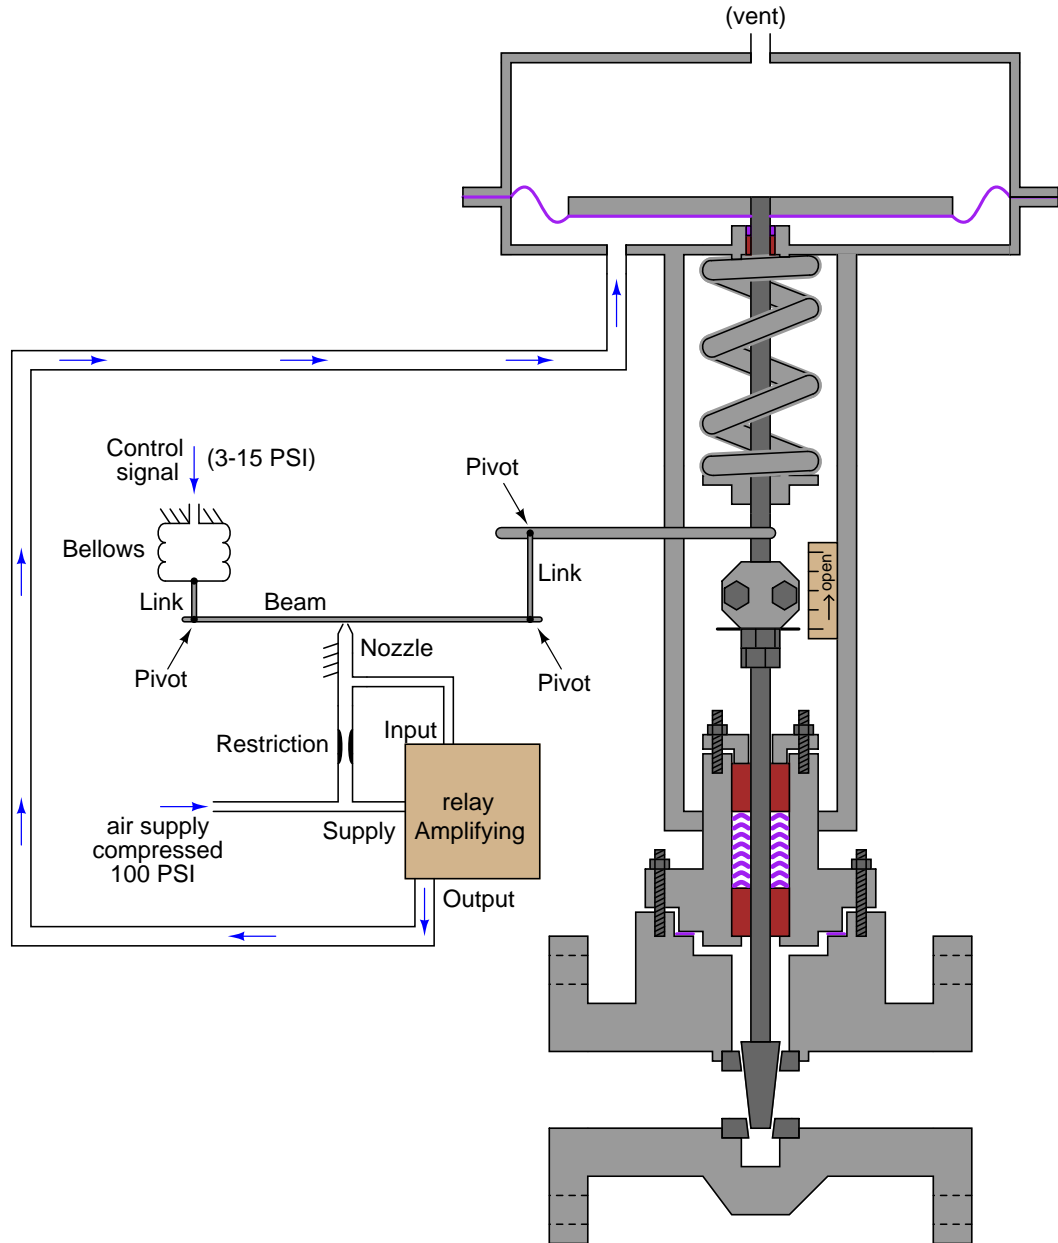
A roller-tipped *follower* at the end of a gold-colored beam rides along the cam's circumference. Cam motion is translated into linear force by the compression of a coil spring directly against the force of the pneumatic bellows on the force beam. When the cam moves far enough to compress the spring enough to balance the additional force generated by the bellows, the force beam return to its equilibrium position (very nearly where it began) and the valve will stop moving.

If you closely examine this last photograph, you will see the positioner's *zero* screw adjustment: the threaded rod extending below the gold-colored beam. This screw adjustment biases the amount of spring compression, making the positioner mechanism “think” the cam is in a different position. For example, turning this threaded rod clockwise (as viewed from the slotted end where a screwdriver would engage) further compresses the spring, pushing up with greater force on the dark-colored bar, achieving the same effect as if the cam had rotated counter-clockwise slightly. This makes the positioner take action to rotate the cam clockwise to compensate, closer toward the 0% valve stem position.

Even though the cam and follower in this positioner mechanism actually do move with valve stem motion, it is still considered a *force-balance* mechanism because the beam connected to the pilot valve does not move appreciably. The pilot valve always comes to rest at its equilibrium position through a balancing of *forces* on the beam.

27.10.2 Motion-balance pneumatic positioners

Motion-balance pneumatic valve positioner designs also exist, whereby the motion of the valve stem counteracts motion (not force) from another element. The following cutaway illustration shows how a simple motion-balance positioner would work:




In this mechanism, an increasing signal pressure causes the beam to advance toward the nozzle, generating increased nozzle backpressure which then causes the pneumatic amplifying relay to send more air pressure to the valve actuator. As the valve stem lifts up, the upward motion imparted to the right-hand end of the beam counters the beam's previous advance toward the nozzle. When equilibrium is reached, the beam will be in an angled position with the bellows' motion balanced by valve stem motion.

The following photograph shows a close view of a Fisher model 3582 pneumatic motion-balance positioner's mechanism:



At the heart of this mechanism is a D-shaped metal ring translating bellows motion and valve stem motion into flapper (baffle) motion. As the bellows (located underneath the upper-right corner of the D-ring) expands with increasing pneumatic signal pressure, it rocks the beam along its vertical axis. With the positioner set for direct-acting operation, this rocking motion drives the flapper closer to the nozzle, increasing backpressure and sending more compressed air to the valve actuator:




As the valve stem moves, a feedback lever rotates a cam underneath the bottom-most portion of the D-ring. A roller-tipped “follower” riding on that cam translates the valve stem’s motion to another rocking motion on the beam, this time along the horizontal axis. Depending on how the cam has been fixed to the feedback shaft, this motion may rock the flapper farther away from the nozzle or closer toward the nozzle. This selection of cam orientation must match the action of the actuator: either direct (air to extend the stem) or reverse (air to retract the stem).

The D-ring mechanism is rather ingenious, as it allows convenient adjustment of span by angling the flapper (baffle) assembly at different points along the ring’s circumference. If the flapper assembly is set close to horizontal, it will be maximally sensitive to bellows motion and minimally sensitive to valve stem motion, forcing the valve to move farther to balance small motions of the bellows (long stroke length). Conversely, if the flapper assembly is set close to vertical, it will be maximally sensitive to valve stem motion and minimally sensitive to bellows motion, resulting in little valve stroke (i.e. the bellows needs to expand greatly in order to balance a small amount of stem motion).

27.10.3 Electronic positioners

Recall that the purpose of a *valve positioner* is to ensure the mechanical valve's position matches the command signal at all times. Thus, a valve positioner is actually a closed-loop control system in its own right: applying as much or as little pressure to the actuator in order to achieve the commanded valve stem position at all times. Mechanical valve positioners use levers, cams, and other physical components to achieve this closed-loop control.

Electronic valve positioners, such as the Fisher model DVC6000, use an electronic sensor to detect valve stem position, a microprocessor to compare that sensed stem position against the control signal by mathematical subtraction (error = position – signal), then a pneumatic signal converter and relay(s) to send air pressure to the valve actuator. A simplified diagram of a generic electronic valve positioner is shown here:



As you can see from this diagram, there is a lot going on inside an electronic positioner. We have not just one, but *two* control algorithms working together to maintain proper valve position: one monitoring and controlling pressure applied to the actuator (compensating for changes in air supply pressure that might otherwise affect the valve's position) and the other monitoring and controlling stem position itself, sending a cascaded control signal to the pressure control components.

The command signal (sent from the process loop controller, PLC, or other control system) tells the positioner where the valve stem *should* be positioned. The first controller inside the positioner (PI) calculates how much air pressure at the actuator should be needed to achieve the requested stem position. The next controller (PID) drives the I/P (current-to-pressure) converter as much as necessary to achieve that pressure. If anything causes the valve stem to not be at the commanded position, the two controllers inside the positioner work together to force the valve to its proper position.

Not only do electronic valve positioners achieve superior position control when compared to mechanical valve positioners, but their array of sensors and digital communication ability provides a new level of diagnostic data both to maintenance personnel and the supervising control system (if programmed to monitor and act on this data). Examples of diagnostic data provided by electronic positioners include:

- Supply air pressure
- Actuator air pressure
- Ambient temperature
- Position and pressure errors
- Total valve stem travel (like an odometer in an automobile)

Additionally, the microprocessor embedded within an electronic valve positioner is capable of performing self-tests, self-calibrations, and other routine procedures traditionally performed by instrument technicians on mechanical valve positioners. Having access to such measurements as total valve stem travel even allows an electronic positioner to predictively calculate packing wear-out time, automatically flagging a maintenance alarm notifying operators and/or instrument technicians when the valve's stem packing will need to be replaced!

A useful capability of some “smart” valve positioners – since they monitor actuator air pressure in addition to stem position – is the ability to maintain a respectable degree of valve control in the event of a stem position sensor failure. If the microprocessor detects a failed (off-scale) position feedback signal, it may be programmed to continue operating the valve based on pressure alone: adjusting the applied air pressure to the valve actuator according to the pressure/position function it has recorded in the past. While not strictly functioning as a *positioner* any more since it cannot sense valve stem position, it may still fulfill its role as a volume booster (compared to the flow capacity of a typical I/P) and give reasonable control over the valve where any other (non-smart) valve positioner would actually make matters worse in the event of losing its stem position feedback. With any purely mechanical positioner, the control valve will typically “saturate” either fully open or fully closed if the stem position feedback linkage falls off. Not so with the best “smart” positioners!

Perhaps the most significant diagnostic data provided by an electronic positioner is the comparison of actuator pressure versus stem position, usually expressed in the form of a graph. Actuator pressure is a direct reflection of force applied by the actuator to the valve stem, since the relationship between force and pressure for either a piston or a diaphragm is simply $F = PA$, where area (A) is constant. Thus, a comparison of actuator air pressure versus stem position is really an expression of force-versus-position for the valve. This so-called *valve signature* is incredibly useful in identifying and correcting such problems as excessive packing friction, valve trim interference, and plug/seat fit problems.

A screenshot showing a “valve signature” (taken from an Emerson software product called *ValveLink*, as part of their *AMS* suite) appears here, showing the behavior of an air-to-open Fisher E-body globe valve:



Two plots of actuator pressure versus stem position are shown in this graph, one red and one blue. The red graph shows the valve’s response in the *opening* direction where additional pressure is required to overcome packing friction as the valve moves open (up). The blue graph shows the


valve as it *closes*, less pressure applied to the diaphragm now to allow the spring's compression to overcome¹⁹ packing friction as the valve moves closed (down) to its resting state. The sharp turns at each end of this graph show where the valve stem reaches its end positions and cannot move farther despite further changes in actuator pressure.

Each plot is roughly linear in accordance with Hooke's Law describing the behavior of the valve spring, where the force applied to a spring is directly proportional to the displacement (compression) of that spring: $F = kx$. Any departure from a single linear plot indicates some *other* force(s) besides spring compression and pneumatic force acting on the valve stem. This is why we see two plots vertically offset from each other: packing friction is another force acting on the valve stem in addition to spring compression and the force exerted by air pressure on the actuator diaphragm. The relatively small magnitude of this offset as well as its consistency indicates that packing friction in this valve is "healthy." The more packing friction this valve experiences, the more vertically-offset the two plots will be.

The sharp down-turn at the left-hand end of the graph where the valve plug contacts the seat is called the *seating profile*. Located at the end of the plot where the valve closes off, the *seating profile* holds much useful information about the physical condition of the plug and seat. As these trim parts wear in a control valve, the shape of the seating profile changes accordingly. Irregular seating profiles may diagnose seat erosion, galling, or a number of other maladies.

¹⁹Recall from basic physics that friction force always *opposes* the direction of motion. Thus, when the valve is opening, friction works against the actuator's air pressure (assuming an air-to-open valve), requiring additional air pressure to maintain motion. When the valve is closing, though, packing friction works in the same direction as the actuator's air pressure "helping" the valve stay more open than it should. This is why the positioner must maintain *less* actuator air pressure for any given position while moving closed than while the valve moves open. The difference in air pressure moving open versus moving closed at any given stem position is proportional to *twice* the dynamic packing friction. Stated mathematically, $F_{packing} = \frac{1}{2}(P_{opening} - P_{closing})A$.

By zooming in on the lower-left end of the valve signature graph, the seating profile may be examined in fine detail. A seating profile taken of a Fisher E-body globe valve in pristine condition appears here:



If the maintenance staff of a particular facility are diligent enough to record the valve signatures of its control valves after assembly or rebuild, the “original” signature of any particular control valve may be compared against the signature of the same control valve taken at any later date, allowing qualitative determinations of wear without having to disassemble the valves for inspection.

Interestingly, this relationship of actuator pressure (force) to stem position is also available in the electronic positioners used with some modern electrically-actuated valves. In the case of an electric actuator, force applied to the valve stem directly relates to *motor current*, which is easily measured and interpreted by the electronic positioner. Thus, even with a different actuator technology, the same kind of diagnostic data may be presented in graphical form for the purpose of more easily diagnosing valve problems. These diagnostics apply even to open/close motor-operated valves not

used for throttling service²⁰, and are especially useful on gate, plug, and ball-type shut-off valves where seat engagement is substantial for tight shut-off.

27.11 Split-ranging

There are many process control applications in industry where it is desirable to have multiple control valves respond to the output of a common controller. Control valves configured to follow the command of the same controller are said to be *split-ranged*, or *sequenced*.


Split-ranged control valves may take different forms of sequencing. A few different modes of control valve sequencing are commonly seen in industry: *complementary*, *exclusive*, and *progressive*²¹.

²⁰Prior to the advent of motor-actuated valves, practically all shutoff valves in industrial facilities were manually operated. While this is an inconvenience for operations personnel, it did carry one advantage: the human operators tasked with closing these valves by hand could feel how each valve seated. The amount of effort and the onset of closing torque sensed while turning the valve handle shut gave operators tactile feedback on the condition of each valve seat. Motor-powered valve actuators eliminated the need for this routine manual labor, but also eliminated the routine collection of this valuable diagnostic information. Modern electric valve actuators now provide the best of both worlds: convenient and fast valve operation with accurate self-diagnostic assessment of valve seating.

²¹I have searched in vain for standardized names to categorize different forms of control valve sequencing. The names “complementary,” “exclusive,” and “progressive” are my own invention. If I have missed someone else’s categorization of split-ranging in my research, I sincerely apologize.

27.11.1 Complementary valve sequencing


The first type of control valve sequencing is a mode where two valves serve to proportion a mixture of two fluid streams, such as this example where base and pigment liquids are mixed together to form colored paint:



Both base and pigment valves operate from the same 3 to 15 PSI pneumatic signal output by the I/P transducer (AY), but one of the valves is Air-To-Open while the other is Air-To-Close. The following table shows the relationship between valve opening for each control valve and the controller's output:


Controller output (%)	I/P output (PSI)	Pigment valve (stem position)	Base valve (stem position)
0 %	3 PSI	fully shut	fully open
25 %	6 PSI	25% open	75% open
50 %	9 PSI	half-open	half-open
75 %	12 PSI	75% open	25% open
100 %	15 PSI	fully open	fully shut

An alternative expression for this split-range valve behavior is a graph showing each valve opening as a colored stripe of varying width (wider representing further open). For this particular mode of split-ranging, the graph would look like this:



With this form of split-ranging, there is never a condition in the controller's output range where both valves are fully open or fully shut. Rather, each valve *complements* the other's position²².

An alternative to complementary valve sequencing in a process where two fluid streams mix (or diverge) is to use one *three-way* valve²³ rather than a pair of two-way valves:



²²In mathematics, a "complement" is a value whose sum with another quantity always results in a fixed total. Complementary angles, for instance, always add to 90° (a right angle).

²³Also known as a *mixing valve* or a *diverting valve*, depending on how it is applied to process service.

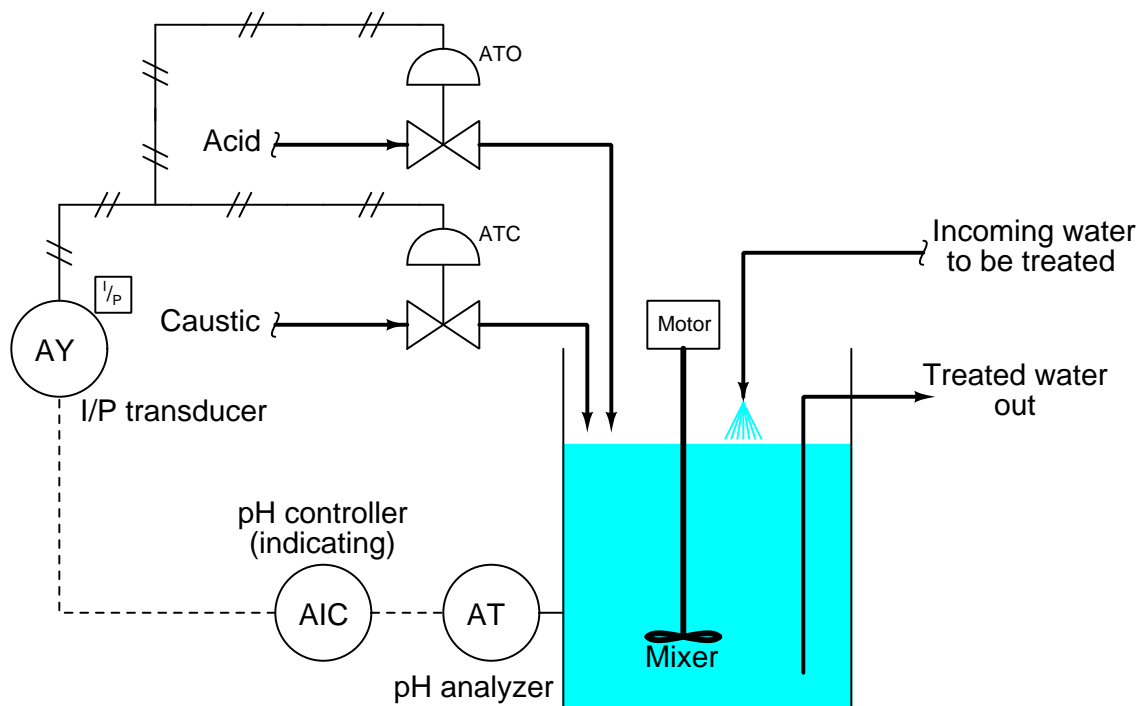
A photograph of a three-way globe valve mixing hot and cold water to control temperature is shown here:



27.11.2 Exclusive valve sequencing

Other applications for split-ranged control valves call for a form of valve sequencing where both valves are fully closed at a 50% controller output signal, with one valve opening fully as the controller output drives toward 100% and the other valve opening fully as the controller output goes to 0%. The nature of this valve sequencing is to have an “either-or” throttled path for process fluid. That is, *either* process fluid flows through one valve *or* through the other, but never through both at the same time.

A practical example of this form of split-ranging is reagent feed for a pH neutralization process, where the pH value of process liquid is brought closer to neutral by the addition of either acid or caustic reagent liquids:




Here, a pH analyzer monitors the pH value of the liquid solution and a single pH controller commands two reagent valves to open when needed. If the process pH begins to increase, the controller output signal increases as well (direct action) to open up the acid valve. The addition of acid to the mixture will have the effect of lowering the mixture’s pH value. Conversely, if the process pH begins to decrease, the controller output signal will decrease as well, closing the acid valve and then opening the caustic valve. The addition of caustic to the mixture will have the effect of raising the mixture’s pH value.

Both reagent control valves operate from the same 3 to 15 PSI pneumatic signal output by the I/P transducer (AY), but the two valves' calibrated ranges are not the same. The Air-To-Open acid valve has an operating range of 9 to 15 PSI, while the Air-To-Close caustic valve has an operating range of 9 to 3 PSI. The following table shows the relationship between valve opening for each control valve and the controller's output:

Controller output (%)	I/P output (PSI)	Acid valve (stem position)	Caustic valve (stem position)
0 %	3 PSI	fully shut	fully open
25 %	6 PSI	fully shut	half-open
50 %	9 PSI	fully shut	fully shut
75 %	12 PSI	half-open	fully shut
100 %	15 PSI	fully open	fully shut

Again, we may express the two valves' exclusive relationship in the form of a graph, with colored stripes representing valve opening:



Exclusive-sequenced control valves are used in applications where it would be undesirable to have both valves open simultaneously. In the example given of a pH neutralization process, the goal here is for the controller to add either acid reagent or caustic reagent to “push” the pH value either direction as needed. However, simultaneously adding both acid and caustic to the process would be wasteful, as one reagent would simply neutralize the other with no benefit to the process liquid itself.

27.11.3 Progressive valve sequencing


A third form of control valve sequencing is used to expand the operating range of flow control for some fluid beyond that which a single control valve could muster. Once again pH control provides a suitable example to illustrate an application of this form of sequencing.

pH is an especially challenging application of process control because the dynamic range of the process is enormous. Each unit of pH value change represents a *ten-fold* change in hydrogen ion concentration within the process liquid. This means the difference in ion concentration between a process liquid having a value of 10 pH and a process liquid having a value of 7 pH (a pH difference of 3) is a factor of *one thousand* (10^3)! Consequently, the flow rate of reagent necessary to neutralize a process liquid stream may vary widely. It is quite possible that a control valve sized to throttle minimum flow will simply be too small to meet the demands of high flow when needed. Yet, a control valve sized large enough to meet the maximum flow rate may be too large to precisely “turn down” when just a trickle of reagent is needed.

This same general control problem was encountered by automotive engineers in the days when *carburetors* were used to mix gasoline with air prior to combustion in an engine. A carburetor is a mechanical air flow control device using a “butterfly” valve element to throttle air flow into the engine, and a venturi element producing vacuum to aspirate fuel droplets into the air stream to create an air-fuel mixture. A carburetor with a butterfly valve and flow tube sized to idle well and respond to the needs of in-town driving would not flow enough air to provide good high-speed performance. Conversely, a large carburetor suitable for driving at racing speeds would offer poor control at low-speed and idling operation. Their solution to this problem was the *progressive carburetor*, having two butterfly valves to throttle the flow of air into the engine. One butterfly valve passed low amounts of air flow only, while a larger butterfly valve opened up only when the accelerator pedal was nearly at its maximum position. The combination of two differently-sized butterfly valves – progressively opened – gave drivers the best of both worlds. Now, an automobile engine could perform well both at low power levels and at high power levels.


On a fundamental level, the problem faced in pH control as well as by early automotive engineers is the same thing: insufficient *rangeability*. Some processes demand a greater range of control than any single valve can deliver, and it is within these processes that a pair of progressively-sequenced control valves is a valid solution.

Applying this solution to a pH control process where the incoming liquid always has a high pH value, and must be neutralized with acid:



Proper sequencing of the small and large acid control valves is shown in the table and the graph:

Controller output (%)	I/P output (PSI)	Small acid valve (stem position)	Large acid valve (stem position)
0 %	3 PSI	fully shut	fully shut
25 %	6 PSI	half-open	fully shut
50 %	9 PSI	fully open	fully shut
75 %	12 PSI	fully open	half-open
100 %	15 PSI	fully open	fully open



With the two acid control valves sequenced progressively, the control system will have significantly more rangeability necessary to regulate pH under widely varying process conditions.

A photograph showing a pair of progressively sequenced control valves used to control boiler feedwater to the high-pressure (“HP”) drum on a large power boiler is shown here:



One small control valve (horizontal, with a green actuator) is the first to open from 4 to 12 mA. A larger control valve (vertical, with a red actuator) is the next to open from 12 mA to 20 mA. Under normal (full-power) operating conditions, the green valve is wide-open while the larger red valve throttles feedwater to the steam drum. When the boiler is operating in standby (low-power) mode, the red valve is shut while the smaller green valve throttles feedwater to the steam drum. The combination of these two control valves provides better steam drum level control over the boiler’s full range of operation than any single control valve could do alone.


27.11.4 Valve sequencing implementations

In all previous control valve sequencing examples shown, both control valves received the same pneumatic signal from a common I/P (current-to-pressure) converter. This means each valve received the exact same pressure signal from the transducer for any given controller output value. Sequencing of the two valves (i.e. making each one respond differently to the same air pressure signal), therefore, was a matter of setting each valve to a different bench-set pressure range.


It should be understood, however, that setting up two control valves with different bench-set ranges is not the only way to split-range a pair of valves. Other ways exist as well, each with its own advantages and disadvantages.

The following illustrations depict several alternative methods for control valve sequencing (split-ranging):


Common pneumatic signal




Common electrical signal



Dual controller outputs



Dual controllers



The common pneumatic signal approach (one controller, one I/P transducer) is simple but suffers from the disadvantage of slow response, since one I/P transducer must drive two pneumatic actuators. Response time may be improved by adding a pneumatic *volume booster* between the I/P and the valve actuators, or by adding a positioner to at least one of the valves. Either of these solutions works by the same principle: reducing the air volume demand on the one common I/P transducer.

Wiring two I/P transducers in series so they share a common signal current is another way to split-range two control valves. This approach does not suffer from slow response, since each valve has its own dedicated I/P transducer to supply it with actuating air. We now have a choice where we implement the split ranges: we can do it in the I/P transducers (i.e. each I/P transducer having a different calibration) *or* in the control valves (i.e. each control valve having a different bench-set pressure range). Since it is generally easier to re-range an I/P than it is to rebuild a control valve with a different spring (to give it a different actuating pressure range), this approach has the advantage of convenient configuration.

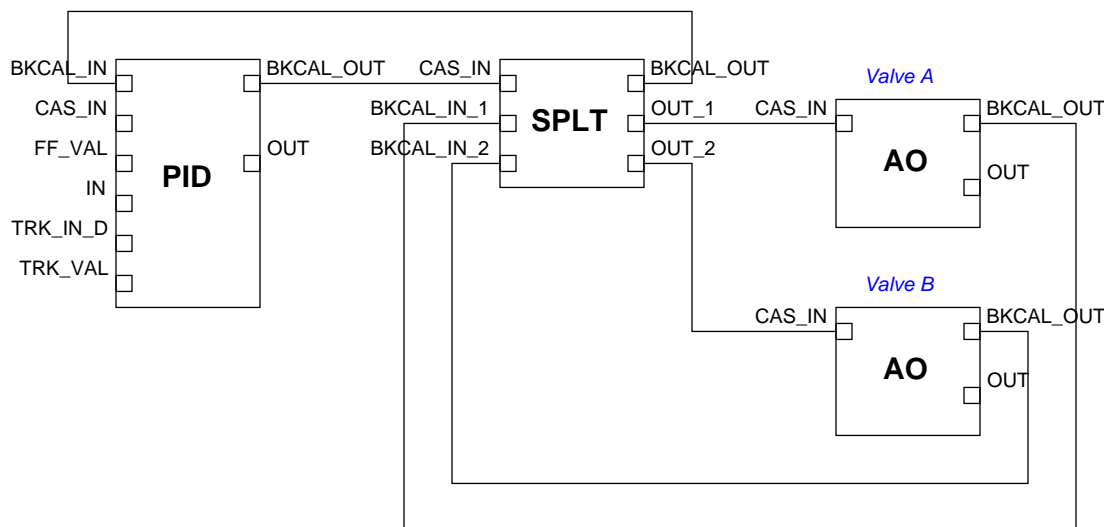
A disadvantage of the series-wired I/P strategy is the extra burden placed on the controller's output signal circuitry: one must be careful to ensure the two series-connected I/P converters do not drop too much voltage at full current, or else the controller may have difficulty driving both devices in series. Another (potential) disadvantage of series-connected valve devices in one current loop is the inability to install "smart" instruments communicating with the HART protocol, since multiple devices on the same loop will experience address conflicts²⁴. HART devices can only work in hybrid analog/digital mode when there is one device per 4-20 mA circuit.

A popular way to implement split-ranging is to use multiple 4-20 mA outputs on the same controller. This is very easy to do if the controller is part of a large system (e.g. a DCS or a PLC) with multiple analog output channels. If multiple outputs are configured on one controller, each valve will have its own dedicated wire pair for control. This tends to result in simpler wiring than series-wired I/P transducers or positioners, since each valve loop is a standard 4-20 mA circuit just like any other (non-split-ranged) control valve loop circuit. It may also be the most practical way to implement split-ranging when "smart" valve positioners are used, since the dedicated loop circuits allow for normal operation of the HART protocol with no address conflicts.

An advantage of dual controller outputs is the ability to perform the split-range sequencing within the controller itself, which is often easier than re-ranging an I/P or calibrating a valve positioner. This way, the 4-20 mA signals going to each valve will be unique for any given controller output value. If sequenced as such, the I/P transducer calibration and valve bench-set values may be standard rather than customized. Of course, just because the controller is capable of performing the necessary sequencing doesn't mean the sequencing must be done within the controller. It is possible to program the controller's dual analog outputs to send the exact same current signal to each valve, configuring each valve (or each positioner, or each I/P transducer) to respond differently to the identical current signals.

²⁴Although the HART standard does support "multidrop" mode where multiple devices exist on the same current loop, this mode is digital-only with no analog signal support. Not only do many host systems not support HART multidrop mode, but the relatively slow data communication rate of HART makes this choice unwise for most process control applications. If analog control of multiple HART valve positioner devices from the same 4-20 mA signal is desired, the address conflict problem may be resolved through the use of one or more *isolator* devices, allowing all devices to share the same analog current signal but isolating each other from HART signals.

A digital adaptation of the dual-output controller sequencing method is seen in FOUNDATION Fieldbus systems²⁵, where a special software function block called “SPLT” exists to provide split-ranged sequencing to two valves. The “SPLT” function block takes in a single control signal and outputs two signals, one output signal for each valve in a split-ranged pair. The function block diagram for such a system appears here:



In this Fieldbus system, a single PID control block outputs a signal to the SPLT block, which is programmed to drive two unique positioning signals to the two valves’ AO (analog output) blocks. It should be noted that while each AO block is unique to its own control valve, the SPLT and even PID blocks may be located in any capable device within the Fieldbus network. With FOUNDATION Fieldbus, control system functions are not necessarily relegated to separate devices. It is possible, for example, to have a control valve equipped with a Fieldbus positioner actually perform its own PID control calculations and split-ranged sequencing by locating those function blocks in that one physical device!

Dual controllers are an option only for specialized applications requiring different degrees of responsiveness for each valve, usually for exclusive or progressive split-ranging applications only. Care must be taken to ensure the controllers’ output signals do not wander outside of their intended ranges, or that the controllers do not begin to “fight” each other in trying to control the same process variable²⁶.


²⁵To review, Fieldbus is an all-digital industrial control protocol, where instruments connect to a control system and to each other by means of a single network cable. Signals are routed not by specific wire connections, but rather by software entities called *function blocks* whereby the engineer or technician programs the instruments and control system what to do with those signals. The function blocks shown in this example would typically be accessed through the graphic display of a DCS in a real Fieldbus system, lines drawn between the blocks instructing the system where each of the instrument signals need to go.

²⁶Both controllers should be equipped with provisions for reset windup control (or have no integral action at all), such that the output signal values are predictable enough that they behave as a synchronized pair rather than as two separate controllers.

An important consideration – and one that is easily overlooked – in split-range valve systems is *fail-safe mode*. As discussed in a previous section of this chapter (on section 27.7.3), the basis of fail-safe control system design is that the control valve(s) must be chosen to fail in the mode that is safest for the process in the event of actuating power loss or control signal loss. The actions of all other instruments in the loop should then be selected to complement the valves' natural operating mode.

In control systems where valves are split-ranged in either complementary or exclusive fashion, one control valve will be fully closed and the other will be fully open at each extreme end of the signal range (e.g. at 4 mA and at 20 mA). If the sequencing for a set of complementary or exclusive split-ranged control valves happens after the controller (e.g. different actuator actions) the valves must fail in opposite modes upon loss of controller signal. However, if it is deemed safer for the process to have the two valves fail in the same state – for example, to both fail closed in the event of air pressure or signal loss – we may use dual sequenced controller outputs, achieving either complementary or exclusive control action by driving the two valves with two different output signals. In other words, split-ranging two control valves so they normally behave in opposite fashion does *not* necessarily mean the two valves must fail in opposite states. The secret to achieving proper failure mode *and* proper split-range sequencing is to carefully locate where the sequencing takes place in the control system.


As an example of a split-ranged system with opposite valve failure modes, consider the following temperature control system supplying either hot water or chilled water to a “jacket” surrounding a chemical reactor vessel. The purpose of this system is to either add or remove heat from the reactor as needed to control the temperature of its contents. Chemical piping in and out of the reactor vessel has been omitted from this P&ID for simplicity, so we can focus just on the reactor’s temperature control system:



Here, the controller has been configured for dual-output operation, where the output value drives two identical 4-20 mA signals to the control valve positioners, which directly input the current signals from the controller without the need for I/P transducers in between. The hot water valve (TV-37a) is fail-closed (FC) while the cold water valve (TV-37b) is fail-open (FO). Half-range positioner calibrations provide the exclusive sequencing necessary to ensure the two valves are never open simultaneously – TV-37b operates on the lower half of the 4-20 mA signal range (4-12 mA), while TV-37a operates on the upper half (12-20 mA).

Consider the effects from the controller (TIC-37) losing power. Both 4-20 mA signals will go dead, driving both valves to their fail-safe modes: hot water valve TV-37a will fully close, while cold water valve TV-37b will fully open. Now consider the effects of air pressure loss to both valves. With no air pressure to operate, the actuators will likewise spring-return to their fail-safe modes: once again hot water valve TV-37a will fully close, while cold water valve TV-37b will fully open. In both failure events, the two control valves assume consistent states, ensuring maximum cooling to the reactor in the event of an output signal or instrument air failure.

However, suppose we desired both of these valves fail in the closed position in the event of an output signal or instrument air failure, rather than have the cooling valve fail open while the heating valve fails closed. Clearly this would require both TV-37a and TV-37b to be fail-closed (FC), which would mean we must find some other way to sequence their operation to achieve split ranging. Examine this reconfiguration of the reactor temperature control system, using identical control valves (signal-to-open, fail-closed) for both hot and cold water supply, and a controller with exclusively-sequenced 4-20 mA output signals:




Consider the effects from the controller (TIC-37) losing power. Both 4-20 mA signals will go dead, driving both valves to their fail-safe modes: fully closed. Now consider the effects of air pressure loss to both valves. With no air pressure to operate, the actuators will spring-return to their fail-safe modes: once again both control valves fully close. In both failure events, the two control valves consistently close. The failure modes of both valves are still consistent regardless of the nature of the fault, but note how this scheme allows both valves to fail in the same mode if that is what we deem safest for the process.

As with all fail-safe system designs, we begin by choosing the proper fail-safe mode for each control valve *as determined by the safety requirements of the process*, not by what we would consider the simplest or easiest-to-understand instrument configurations. Only after we have chosen each valve's failure mode do we choose the other instruments' configurations. This includes split-range sequencing: where and how we sequence the valves is a decision to be made only after the valves' fail-safe states are chosen based on process safety.

27.12 Control valve sizing

When control valves operate between fully open and fully shut, they serve much the same purpose in process systems as resistors do in electric circuits: to dissipate energy. Like resistors, the form that this dissipated energy takes is mostly heat, although some of the dissipated energy manifests in the form of vibration and noise²⁷.

In most control valves, the dominant mechanism of energy dissipation comes as a result of turbulence introduced to the fluid as it travels through constrictive portions of the valve trim. The following illustration shows these constrictive points within two different control valve types (shown by arrows):




Note: red arrows show points of maximum fluid energy loss

The act of choosing an appropriate control valve for the expected energy dissipation is called *valve sizing*.

²⁷Valve noise may be *severe* in some cases, especially in certain gas flow applications. An important performance metric for control valves is noise production expressed in decibels (dB).

27.12.1 Physics of energy dissipation in a turbulent fluid stream

Control valves are rated in their ability to throttle fluid flow much in the same way resistors are rated in their ability to throttle the flow of electrons in a circuit. For resistors, the unit of measurement for electron flow restriction is the *ohm*: 1 ohm of resistance results in a voltage drop of 1 volt across that resistance given a current through the resistance equal to 1 ampere:



*Arrows point in direction
of conventional flow*

The mathematical relationship between current, voltage, and resistance for any resistor is *Ohm's Law*:

$$R = \frac{V}{I}$$

Where,

R = Electrical resistance in ohms

V = Electrical voltage drop in volts

I = Electrical current in amperes

Ohm's Law is a simple, linear relationship, expressing the "friction" encountered by electric charge carriers as they slowly drift through a solid object.

When a fluid moves turbulently through any restriction, energy is inevitably dissipated in that turbulence. The amount of energy dissipated is proportional to the kinetic energy of the turbulent motion, which is proportional to the square of velocity according to the classic kinetic energy equation for moving objects:

$$E_k = \frac{1}{2}mv^2$$

If we were to re-write this equation to express the amount of kinetic energy represented by a volume of moving fluid with velocity v , it would look like this:

$$\text{Kinetic energy per unit volume} = \frac{1}{2}\rho v^2$$

We know that the amount of energy dissipated by turbulence in such a fluid stream will be some proportion (k) of the total kinetic energy, so:

$$\text{Energy dissipated per unit volume} = \frac{1}{2}k\rho v^2$$

Any energy lost in turbulence eventually manifests as a loss in fluid pressure downstream of that turbulence. Thus, a control valve throttling a fluid flowstream will have a greater upstream pressure than downstream pressure (assuming all other factors such as pipe size and height above ground level being the same downstream as upstream):

$$\begin{array}{c}
 \Delta P = P_1 - P_2 \\
 \begin{array}{c}
 P_1 \quad + \quad | \quad - \quad P_2 \\
 \underbrace{\hspace{1.5cm}} \\
 R
 \end{array} \\
 Q \longrightarrow \frac{\hspace{1.5cm}}{\hspace{1.5cm}} \longrightarrow Q
 \end{array}$$

This pressure drop ($P_1 - P_2$, or ΔP) is equivalent to the voltage drop seen across any current-carrying resistor, and may be substituted for dissipated energy per unit volume in the previous equation²⁸. We may also substitute $\frac{Q}{A}$ for velocity v because we know volumetric flow rate (Q) is the product of fluid velocity and pipe cross-section area ($Q = Av$) for incompressible fluids such as liquids:

$$P_1 - P_2 = \frac{1}{2}k\rho \left(\frac{Q}{A}\right)^2$$

Next, we will solve for a quotient with pressure drop ($P_1 - P_2$) in the numerator and flow rate (Q) in the denominator so the equation bears a resemblance to Ohm's Law ($R = \frac{V}{I}$):

$$\begin{aligned}
 P_1 - P_2 &= \frac{1}{2}k\rho \frac{Q^2}{A^2} \\
 \frac{P_1 - P_2}{Q^2} &= \frac{k\rho}{2A^2} \\
 \frac{\sqrt{P_1 - P_2}}{Q} &= \sqrt{\frac{k\rho}{2A^2}}
 \end{aligned}$$

²⁸In case you were wondering, it is appropriate to express energy loss per unit volume in the same units of measurement as pressure. For a more detailed discussion of dimensional analysis, see section 2.11.13 beginning on page 208 where Bernoulli's equation is examined and you will see how the units of $\frac{1}{2}\rho v^2$ and P are actually the same.

Either side of the last equation represents a sort of “Ohm’s Law” for turbulent liquid restrictions: the left-hand side expressing fluid “resistance” in the state variables of pressure drop and volumetric flow, and the right-hand term expressing fluid “resistance” as a function of fluid density and restriction geometry. We can see how pressure drop ($P_1 - P_2$) and volumetric flow rate (Q) are not linearly related as voltage and current are for resistors, but that nevertheless we still have a quantity that acts like a “resistance” term:

$$R = \frac{\sqrt{P_1 - P_2}}{Q} \qquad R = \sqrt{\frac{k\rho}{2A^2}}$$

Where,

- R = Fluid “resistance”
- P_1 = Upstream fluid pressure
- P_2 = Downstream fluid pressure
- Q = Volumetric fluid flow rate
- k = Turbulent energy dissipation factor
- ρ = Mass density of fluid
- A = Cross-sectional area of restriction

The fluid “resistance” of a restriction depends on several variables: the proportion of kinetic energy lost due to turbulence (k), the density of the fluid (ρ), and the cross-sectional area of the restriction (A). In a control valve throttling a liquid flow stream, only the first and last variables are subject to change with stem position, fluid density remaining relatively constant.

In a wide-open control valve, especially valves offering a nearly unrestricted path for moving fluid (e.g. ball valves, eccentric disk valves), the value of A will be at a maximum value essentially equal to the pipe’s area, and k will be nearly zero²⁹. In a fully shut control valve, A is zero, creating a condition of infinite “resistance” to fluid flow.

It is customary in control valve engineering to express the “restrictiveness” of any valve in terms of how much flow it will pass given a certain pressure drop and fluid specific gravity (G_f). This measure of valve performance is called *flow capacity* or *flow coefficient*, symbolized as C_v . A greater flow capacity value represents a less restrictive (less “resistive”) valve, able to pass greater rates of flow for the same pressure drop. This is analogous to expressing an electrical resistor’s rating in terms of conductance (G) rather than resistance (R): how many amperes of current it will pass with 1 volt of potential drop ($I = GV$ instead of $I = \frac{V}{R}$).

If we return to one of our earlier equations expressing pressure drop in terms of flow rate, restriction area, dissipation factor, and density, we will be able to manipulate it into a form expressing flow rate (Q) in terms of pressure drop and density, collecting k and A into a third term which will become flow capacity (C_v):

$$P_1 - P_2 = \frac{1}{2}k\rho\frac{Q^2}{A^2}$$

²⁹In a case of minimal throttling, almost none of the fluid’s kinetic energy is lost to turbulence, but rather passes right through the valve unrestricted.

First, we must substitute specific gravity (G_f) for mass density (ρ) using the following definition of specific gravity:

$$G_f = \frac{\rho}{\rho_{water}}$$

$$\rho_{water}G_f = \rho$$

Substituting and continuing with the algebraic manipulation:

$$P_1 - P_2 = \frac{1}{2}k\rho_{water}G_f\frac{Q^2}{A^2}$$

$$\frac{P_1 - P_2}{G_f} = \frac{1}{2}k\rho_{water}\frac{Q^2}{A^2}$$

$$\left(\frac{2A^2}{k\rho_{water}}\right)\left(\frac{P_1 - P_2}{G_f}\right) = Q^2$$

$$Q = \sqrt{\frac{2A^2}{k\rho_{water}}}\sqrt{\frac{P_1 - P_2}{G_f}}$$

The first square-rooted term in the equation, $\sqrt{\frac{2A^2}{k\rho_{water}}}$, is the valve capacity or C_v factor. Substituting C_v for this term results in the simplest form of valve sizing equation (for incompressible fluids):

$$Q = C_v\sqrt{\frac{P_1 - P_2}{G_f}}$$



In the United States of America, C_v is defined as the number of gallons per minute of water that will flow through a valve with 1 PSI of pressure drop³⁰. A similar valve capacity expression used with metric units rates valves in terms of how many cubic meters per hour of water will flow through a valve with a pressure drop of 1 bar. This latter flow capacity is symbolized as K_v .

For the best results predicting required C_v values for control valves in any service, it is recommended that you use valve sizing software provided by control valve manufacturers. The formulae shown here do not account for all factors³¹ influencing fluid flow rate and pressure drop, and therefore yield approximate values only. Modern valve sizing software is easy to use, especially when referenced to specific models of control valve sold by that manufacturer, and is able to account for a diverse multitude of factors affecting proper sizing.

³⁰The specification of certain British units of measurement for flow and pressure drop means that there is more to C_v than just $\sqrt{\frac{2A^2}{k\rho_{water}}}$. C_v also incorporates a factor necessary to account for the arbitrary choice of units.

³¹Such factors include fluid compressibility, viscosity, specific heat, vapor pressure to name a few. Not only will modern valve sizing software more accurately predict valve sizes for particular applications than these simple formulae, but this software may also provide estimations of *noise levels* produced by the valve.

Control valve sizing is complex enough that some valve manufacturers used to give away “slide rule” calculator devices so customers could choose the C_v values they needed with relative ease. Photographs of a two-sided valve sizing slide rule are shown here for historical reference:



27.12.2 Importance of proper valve sizing

The flow coefficient of a control valve (C_v) is a numerical value expressing the number of gallons per minute flow of water the valve will pass with a constant pressure drop of 1 PSI from inlet to outlet. This rating is usually given for the valve in its wide-open state. For example, a control valve with a C_v rating of 45 should flow 45 gallons per minute of water through it with a 1 PSI pressure drop *when wide open*. The flow coefficient value for this same control valve will be less than 45 when the valve position is anything less than fully open. When the control valve is in the fully shut position, its C_v value will be zero. Thus, it should be understood that C_v is truly a variable – not a constant – for any control valve, even though control valves are often specified simply by their maximum flow capacity.

As previously shown, the basic liquid-flow equation relating volumetric flow rate to pressure drop and specific gravity is:

$$Q = C_v \sqrt{\frac{P_1 - P_2}{G_f}}$$

Where,

Q = Volumetric flow rate of liquid (gallons per minute, GPM)

C_v = Flow coefficient of valve

P_1 = Upstream pressure of liquid (PSI)

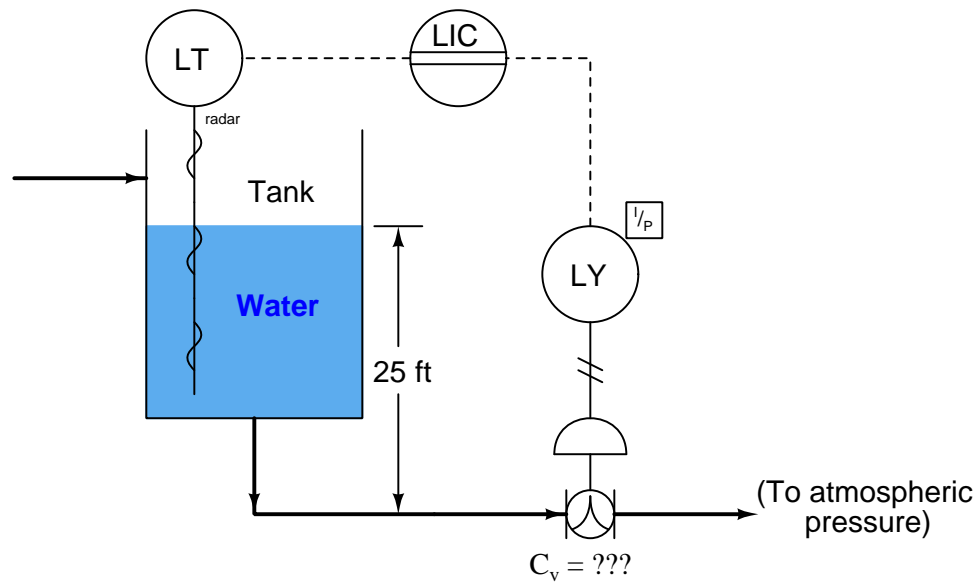
P_2 = Downstream pressure of liquid (PSI)

G_f = Specific gravity of liquid (ratio of liquid density to standard water density)

A common misconception is that C_v represents an absolute maximum flow rating for a valve. This is not true, any more than the resistance value of an electrical resistor reflects some maximum current value for that component. It is possible – at least in theory – to shove *any* rate of flow through *any* open valve given enough pressure to motivate that flow. C_v , like electrical resistance measured in ohms, merely states the proportion between flow rate and pressure drop. C_v is simply the maximum flow rate of water through a control valve *given a 1 PSI pressure difference from upstream to downstream*.

It should be obvious that any control valve must be sized large enough (i.e. possess sufficient maximum C_v capacity) to flow the greatest expected flow rate in any given process installation. A valve that is too small for an application will not be able to pass enough process fluid through it given the fluid pressure available.

Given this fact, it may seem safe to choose a valve sized much larger than what is needed, just to avoid the possibility of not having enough flow capacity. For instance, consider this control valve sizing problem, where a characterized ball valve controls the flow rate of water out of a surge tank to maintain a constant water level 25 feet higher than the height of the valve:



Suppose the process engineer designing this system specifies a maximum flow rate for this valve of 470 GPM. What should the maximum C_v rating be for this valve? To begin, we must know the expected pressure drop across the valve. The 25 foot water column height upstream provides us with the means to calculate P_1 by using the formula for calculating hydrostatic pressure (the pressure generated by a vertical column of liquid under the influence of gravity):

$$P = \gamma h$$

$$P_1 = (62.4 \text{ lb/ft}^3)(25 \text{ feet})$$

$$P_1 = 1560 \text{ PSF} = 10.8 \text{ PSI}$$

There is no need to calculate P_2 , since the P&ID reveals the downstream side of the valve to be vented to atmosphere, and thus is guaranteed to be 0 PSI gauge pressure. This gives us a pressure drop of 10.8 PSI across the control valve, with an expected maximum flow rate of 470 GPM. Manipulating our flow capacity equation to solve for C_v :

$$Q = C_v \sqrt{\frac{P_1 - P_2}{G_f}}$$

$$C_v = \frac{Q}{\sqrt{\frac{P_1 - P_2}{G_f}}}$$

$$C_v = \frac{470 \text{ GPM}}{\sqrt{\frac{10.8 \text{ PSI}}{1}}}$$

$$C_v = 143$$

This tells us we need a control valve with a C_v value of *at least* 143 to meet the specified (maximum) flow rate. A valve with insufficient C_v would not be able to flow the required 470 gallons per minute of water with only 10.8 PSI of pressure drop. As mentioned previously, these calculations yield approximate results only. An analysis using computer-based valve sizing software may require an even larger control valve for this application.

Does this mean we should over-size the valve, just to ensure we don't install one that is too small? For the sake of argument, would there be any problem with installing a control valve with a C_v value of 300? The general answer to these questions is that over-sized valves may create other problems. Not only is there the possibility of allowing too much flow under wide-open conditions (consider whatever process vessels and equipment lie downstream of the oversized valve), but also that the process will be difficult to control under low-flow conditions.

In order to understand how an over-sized control valve leads to unstable control, an exaggerated example is helpful to consider: imagine installing a fire hydrant valve on your kitchen sink faucet³². Certainly, a wide-open hydrant valve would allow sufficient water flow into your kitchen sink. However, due to the gross mismatch between a fire hydrant valve's flow capacity versus that of relatively skinny tubing providing water to the sink faucet, most of this valve's usable range of throttling will be limited to the first *percent* of stem travel. After the valve is opened just a few percent from fully shut, restrictions in the piping of your house's water system will have limited the flow rate to its maximum, thus rendering the rest of the valve's stem travel capacity utterly useless. It would be challenging indeed to try filling a drinking cup with water from this hydrant valve: just a little bit too much stem motion and the cup would be subjected to a full-flow stream of water!

An over-sized valve is therefore an overly-*sensitive* valve from the perspective of the control system driving it. With the upper end of the valve's travel being useless for control (having little or no effect), all the throttling action must take place within the first few percent of stem motion.

³²This is a good example of a general problem-solving strategy in action: making some dramatic change to the scenario and then reasoning the consequences of that change to better understand general principles. For those readers who may be unfamiliar with American terminology, a *fire hydrant* is a large hand valve installed at intervals along public roadways, allowing connection of fire hoses to an underground water supply pipe in the event of an emergency fire. These valves are quite large, and would be comically oversized if installed inside a person's house, for any purpose.


This makes precise control of flow rate more challenging than it should be. Typical valve problems such as stem friction, hysteresis, and calibration error therefore become amplified when the valve is over-sized, because any amount of imprecision in stem positioning becomes a greater *percentage* of the valve's useful travel range than if the valve were properly sized and able to use its full range of motion.

Control valve over-sizing is a common problem in industry, often created by future planning for expanded process flow. “If we buy a large valve now,” so the reasoning goes, “we won't have to replace a smaller valve with a large valve when the time comes to increase our production rate.” In the interim period when that larger valve must serve to control a meager flow rate, however, operational problems caused by poor control quality may end up costing the business more than the cost of an additional valve. The key here, as in so many other applications in business, is to carefully consider costs over the life-span of the device, not just the initial (capital) expense.


A practical solution to the problem of valve over-sizing – especially when larger flow rates will be required of the valve in the future – is to initially order the control valve with a body size suitable for the largest flow capacity it will need in the future, but equipped internally with *reduced-port* (or *restricted-capacity*) trim. This means trim having smaller holes (“ports”) through which the fluid must flow. Such “reduced” trim is under-sized for the valve body, making the control valve's C_v rating significantly less than it would be with normal-sized trim installed. The benefit of installing reduced-port trim in a control valve is that the flow capacity of the valve may be upgraded simply by removing the reduced trim components and replacing them with *full-port* (full-sized) trim. Upgrading a control valve's trim to full-port size is significantly less expensive than replacing the entire control valve with a larger one³³.

³³This is particularly true when one considers the piping changes usually necessary to accommodate a valve size change. Undersized valves installed in a pipe often require *reducer* fittings to “narrow” the full-bore size of the pipe down to the flange size of the control valve body. Upon replacement of the under-sized valve, these reducers must be removed to accommodate the larger valve body. The piping itself may need to be cut and re-welded to match the flange-to-flange dimensions of the new (larger) control valve. All of this requires time, labor, and material investment. If a large valve body with reduced-port trim were initially installed, however, most of this time, labor, and expense could be avoided when the time comes to replace the reduced-port trim with full-port trim.

Reduced-port trim for a stem-guided or port-guided globe valve takes the form of a new (smaller) plug and seat assembly. The seat is specially designed to match the plug for tight shutoff and good throttling behavior, while having the necessary external dimensions to fit the larger valve body casting:



Reduced-port trim for a cage-guided globe valve often consists of nothing more than a new cage, having smaller ports in the cage than standard. Thus, reduced-port trim for a cage-guided globe valve may³⁴ utilize the exact same plug and seat as the full-port trim for the same cage-guided globe valve:



³⁴Reduced-port cage-guided trim may also take the form of a cage, plug, and seat of reduced diameter, with flanges attached in such a way that this smaller trim still fits inside the larger valve body. The example illustrated here, with a full-diameter cage having narrower ports on it, is just one way of achieving reduced flow capacity in a cage-guided design but certainly not the only way.

27.12.3 Gas valve sizing

Sizing a control valve for gas or vapor service is more complicated than for liquid service, due to the compressibility of gases and vapors. As a gas or vapor compresses with changes in pressure, its density changes correspondingly. In previous mathematical analyses of fluid flow restriction, one of our assumptions was that fluid density (ρ) remained constant. This assumption may hold true for some flowing gas conditions as well, provided minimal pressure changes within the path of flow. However, for most gas control valve applications where the very purpose of the valve is to introduce substantial pressure changes, the assumption of constant gas density is unrealistic.

Shown here is one of the simpler gas valve sizing equations you will encounter:

$$Q = 963 C_v \sqrt{\frac{\Delta P (P_1 + P_2)}{G_g T}}$$

Where,

Q = Gas flow rate, in units of Standard Cubic Feet per Hour (SCFH)

C_v = Valve capacity coefficient

ΔP = Pressure dropped across valve, pounds per square inch differential (PSID)

P_1 = Upstream valve pressure, pounds per square inch absolute (PSIA)

P_2 = Downstream valve pressure, pounds per square inch absolute (PSIA)

G_g = Specific gravity of gas (ratio of gas density to standard air density)

T = Absolute temperature of gas in degrees Rankine ($^{\circ}\text{R}$), equal to degrees Fahrenheit plus 459.67

This equation holds true only for “subcritical” flow, where the moving gas stream velocity never approaches the speed of sound³⁵. Other equations exist for calculating flow rates of gas through control valves in the presence of sonic flow regimes. Note the inclusion of absolute pressures in this equation, and not just differential pressure (ΔP , or $P_1 - P_2$). This is intended to correct for effects related to compression of the gas under pressure.


Valve sizing is complicated enough, both for liquid and gas service, that the use of valve sizing computer software is strongly recommended as opposed to hand-calculations. The number of important parameters, nonlinear factors, and alternative equations relevant to control valve sizing are numerous enough to bewilder most technicians (and more than a few engineers). Valve sizing software will also predict noise levels generated by the valve, and in many cases specify actual valve trim styles offered by the manufacturer for mitigating problems such as noise.

³⁵The *ISA Handbook of Control Valves* cites this equation as being valid for conditions where the valve’s downstream pressure (P_2) is equal to or greater than one-half the upstream pressure (P_1), with both pressures expressed in absolute units. In other words, $P_2 \geq 0.5P_1$ or $P_1 \leq 2P_2$. An upstream:downstream pressure ratio in excess of 2:1 usually means flow through a valve will become *choked*.

27.12.4 Relative flow capacity

The flow capacity of a valve (C_v) is a quantitative rating of its ability to pass a fluid flow for a set of given pressure and fluid density conditions, defined as the flow rate of water in gallons per minute through the valve given a constant pressure drop of 1 PSID. C_v may be predicted, or empirically measured, for any type of control valve given the proper information.

Not all control valve types exhibit the same C_v coefficients, however, for the same pipe size. A 4 inch butterfly valve, for example, has a much greater full-open C_v rating than a 4 inch globe valve, due to the much more direct path it offers to a moving fluid. A simple comparison of these two valve types clearly shows why this is true (note the “constriction” points labeled with arrows):



Note: red arrows show points of maximum fluid energy loss

A globe valve is simply more effective at generating fluid turbulence – and therefore dissipating fluid kinetic energy – than a butterfly valve of the same pipe size, because the globe valve design forces the fluid to change direction more often and in different ways.

One way to help quantify a particular valve design’s ability to throttle fluid flow is to express this ability as a ratio of flow coefficient (C_v) versus cross-sectional pipe area. The basic principle here is that we should expect the C_v of any particular valve design to be proportional to pipe area (e.g. a ball valve with twice the pipe area should have twice the flow capacity, all other factors being equal), and therefore a ratio of these two quantities should be fairly constant for any valve design. Since we know the area of a pipe is proportional to the square of either radius or diameter ($A = \frac{\pi d^2}{4}$ or $A = \pi r^2$), we may simplify this ratio by omitting all constants such as π and simply relating C_v factor to the square of pipe diameter (d^2). This ratio is called the *relative flow capacity*, or C_d :

$$C_d = \frac{C_v}{d^2}$$

Several approximate valve capacity factors (C_d) for different control valve types are shown here³⁶, assuming full-area (not reduced-port) trim, a wide-open position, and diameter measured in inches:

Valve design type	C_d
Single-port globe valve, ported plug	9.5
Single-port globe valve, contoured plug	11
Single-port globe valve, characterized cage	15
Double-port globe valve, ported plug	12.5
Double-port globe valve, contoured plug	13
Rotary ball valve, segmented	25
Rotary ball valve, standard port (diameter $\approx 0.8d$)	30
Rotary butterfly valve, 60°, no offset seat	17.5
Rotary butterfly valve, 90°, offset seat	29
Rotary butterfly valve, 90°, no offset seat	40

To calculate the approximate C_v for any valve, all we need to do is square that valve's pipe diameter (in inches) and multiply by the valve type's relative flow capacity (C_d):

$$C_v \approx d^2 C_d$$

Based on the figures in this table, for example, we may predict that a segmented ball valve with a pipe size of 3 inches will have a flow capacity (C_v) of approximately 225, or that a single-port cage-guided globe valve with an 8 inch pipe size will have a C_v of approximately 960.

As you can see from a comparison of C_d values, a no-offset butterfly valve has nearly 4 times the flow capacity of a single-ported contoured-plug globe valve of the same pipe size ($C_d = 40$ versus $C_d = 11$). This makes butterfly valves advantageous in applications where large flow capacities must be achieved at minimal cost, such as in air handling (HVAC) systems for commercial buildings and combustion air controls for large industrial burners.

At first glance this may seem to make larger C_d valve types the superior choice for control valve applications, yielding the greatest C_v values for the smallest pipe sizes. However, there are other factors to consider such as ease of maintenance, noise and cavitation abatement, characterization, and valve seat leakage. For these reasons some of the lowest- C_d valve types (e.g. globe valves) remain popular choices for industrial control applications even though they require larger pipe sizes to achieve the same amount of flow compared to other valve types (e.g. ball valves).

27.13 Control valve characterization

Control valves are supposed to deliver reliable, repeatable control of process fluid flow rate over a wide range of operating conditions. As we will soon see, this is something of a challenge, as the rate of fluid flow through a control valve depends on more than just the position of its stem. This section discusses the problem of control valve behavior in real process applications, and explores the concept of *characterization* as a solution to the problem.

³⁶Source for C_d factors: [Chapter 4.17: Valve Sizing](#) of Béla Lipták's *Instrument Engineers' Handbook, Process Control (Volume II), Third Edition*, page 590.

27.13.1 Inherent versus installed characteristics

When control valves are tested in a laboratory setting, they are connected to a piping system providing a nearly constant pressure difference between upstream and downstream ($P_1 - P_2 = \text{constant}$). With a fluid of constant density and a constant pressure drop across the valve, flow rate becomes directly proportional to the valve's flow coefficient (C_v). This is clear from an examination of the basic valve capacity equation, if we replace the pressure and specific gravity terms with a single constant k :


$$Q = C_v \sqrt{\frac{P_1 - P_2}{G_f}}$$

(If pressures and specific gravity are constant . . .)

$$Q = kC_v$$


As discussed in an earlier section of this chapter (see section 27.12.1), the amount of “resistance” offered by a restriction of any kind to a turbulent fluid depends on the cross-sectional area of that restriction and also the proportion of fluid kinetic energy dissipated in turbulence. If a control valve is designed such that the combined effect of these two parameters vary linearly with stem motion, the C_v of the valve will likewise be proportional to stem position. That is to say, the C_v of a “linear” control valve will be approximately half its maximum rating with the stem position at 50%; approximately one-quarter its maximum rating with the stem position at 25%; and so on.

If such a valve is placed in a laboratory flow test piping system with constant differential pressure and constant fluid density, the relationship of flow rate to stem position will be linear. With constant pressure drop, the flow rate through any valve is directly proportional to that valve's C_v , and with a “linear” valve design the C_v (and therefore the flow rate as well) must be directly proportional to stem position:



However, most real-life valve installations do not provide the control valve with a constant pressure drop. Due to frictional pressure losses in piping and changes in supply/demand pressures that vary with flow rate, a typical control valve “sees” substantial changes in differential pressure as its controlled flow rate changes. Generally speaking, the pressure drop available to the control valve *decreases* as flow rate *increases*.

The result of this pressure drop versus flow relationship is that the actual flow rate of the same valve installed in a real process will *not* linearly track valve stem position. Instead, it will “droop” as the valve is further opened. This “drooping” graph is called the valve’s *installed characteristic*, in contrast to the *inherent characteristic* exhibited in the laboratory with constant pressure drop:




Each time the stem lifts up a bit more to open the valve trim further, flow increases, but not as much as at lower-opening positions. It is a situation of *diminishing returns*, where we still see increases in flow as the stem lifts up, but to a lesser and lesser degree.


In my years of teaching, I have found this concept of “installed characteristic” to be especially challenging for many students. In the interest of clarifying the concept, the next two subsections will present a pair of contrasting valve performance scenarios.

27.13.2 Control valve performance with constant pressure

First, let us imagine a control valve installed at the base of a dam, releasing water from the reservoir. Given a constant height of water in the reservoir, the upstream (hydrostatic) pressure at the valve will likewise be constant. Let's assume this constant upstream pressure will be 20 PSI (corresponding to approximately 46 feet of water column above the valve inlet). With the valve discharging into the air, downstream pressure will essentially be zero. This set of upstream and downstream conditions guarantees a constant pressure drop of 20 PSI across our control valve at all times, for all flow conditions:




Furthermore, let us assume the control valve has a “linear” inherent characteristic and a maximum flow capacity (C_v rating) of 18. This means the valve’s C_v will be 18 at 100% open, 13.5 at 75% open, 9 at 50% open, 4.5 at 25% open, and 0 at fully closed (0% open). We may plot the behavior of this control valve at these four stem positions by graphing the amount of flow through the valve for varying degrees of pressure drop across the valve. The result is a set of *characteristic curves*³⁷ for our hypothetical control valve:



Each curve on the graph traces the amount of flow through the valve at a constant stem position, for different amounts of applied pressure drop. For example, looking at the curve representing 50% open ($C_v = 9$), we can see the valve should flow about 42 GPM at 22 PSI, about 35 GPM at 15 PSI, about 20 GPM at 5 PSI, and so on. Of course, we can obtain these same flow figures simply by evaluating the formula $Q = C_v \sqrt{\Delta P}$ (which is in fact what I used to plot these curves), but the point here is to learn how to interpret the graph.

³⁷For those readers with an electronics background, the concept of “characteristic curves” for a control valve is *exactly* the same as that of characteristic curves for transistors. Instead of plotting the amount of current a bipolar transistor will pass through its collector terminal (I_C) given varying amounts of collector-emitter voltage drop (V_{CE}), we are plotting the rate of water flow through the valve (Q) given varying amounts of supply pressure (ΔP).

We may use this set of characteristic curves to determine how this valve will respond in *any* installation by superimposing another curve on the graph called a *load line*³⁸, describing the pressure drop available to the valve at different flow rates. Since we know our hypothetical dam supplies a constant 20 PSI across the control valve for all flow conditions, the load line for the dam will be a vertical line at 20 PSI:



By noting the points of intersection³⁹ between the valve's characteristic curves and the load line, we may determine the flow rates from the dam at those stem positions:

Opening (%)	C_v	Flow rate (GPM)
0	0	0
25	4.5	20.1
50	9	40.2
75	13.5	60.4
100	18	80.5


If we were to graph *this* table, plotting flow versus stem position, we would obtain a very linear graph. Note how 50% open gives us twice as much flow as 25% open, and 100% open nearly twice as much flow as 50% open. This tells us our control valve will respond linearly when operated under these conditions (i.e. operating with a constant pressure drop).

³⁸Once again, the exact same concept applied in transistor circuit analysis finds application here in control valve behavior! The load line for a transistor circuit describes the amount of voltage available to the transistor under different current conditions, just like the load line here describes the amount of pressure available to the valve under different flow conditions.

³⁹Load line plots are a graphical method of solving nonlinear, simultaneous equations. Since each curve represents a set of solutions to a particular equation, the intersection of two curves represents values uniquely satisfying both equations at the same time.


27.13.3 Control valve performance with varying pressure

Now let us consider a scenario where the pressure drop across the valve changes as the rate of flow through the valve changes. We may modify the previous example of a control valve releasing water from a dam to generate this effect. Suppose the valve is not closely coupled to the dam, but rather receives water through a narrow (restrictive) pipe:



In this installation, the narrow pipe generates a flow-dependent pressure drop due to friction between the turbulent water and the pipe walls, leaving less and less upstream pressure at the valve as flow increases. The control valve still drains to atmosphere, so its downstream pressure is still a constant 0 PSIG, but now its upstream pressure diminishes with increasing flow. How will this affect the valve's performance?

We may turn to the same set of characteristic curves to answer this question. All we need is a new load line describing the pressure available to the valve at different flow rates, then we may look for the points of intersection between this load line and the valve's characteristic curves. For the sake of our hypothetical example, I have sketched an arbitrary "load line" (actually a load *curve*) showing how the valve's pressure falls off as flow rises⁴⁰:




Now we see a definite nonlinearity in the control valve's behavior. No longer does a doubling of stem position (from 25% to 50%, or from 50% to 100%) result in a doubling of flow rate⁴¹:

Opening (%)	C_v	Flow rate (GPM)
0	0	0
25	4.5	18.7
50	9	35.5
75	13.5	48
100	18	57

⁴⁰The precise determination of this curve is based on a model of the narrow pipe as a flow-restricting element, similar in behavior to an orifice, or to a control valve with a fixed stem position. Since pressure is dropped along the pipe's length as a function of turbulence (velocity), the load "line" curves for the exact reason the valve's own characteristic plots are curved: the relationship between fluid velocity and turbulent pressure loss is naturally non-linear.

⁴¹Not only is the response of the valve altered by this degradation of upstream pressure, but we can also see from the load line that a certain maximum flow rate has been asserted by the narrow pipe which did not previously exist: 75 GPM. Even if we unbolted the control valve from the pipe and let water gush freely into the atmosphere, the flow rate would saturate at only 75 GPM because that is the amount of flow where all 20 PSI of hydrostatic "head" is lost to friction in the pipe. Contrast this against the close-coupled scenario, where the load line was vertical on the graph, implying no theoretical limit to flow at all! With an absolutely constant upstream pressure, the only limit on flow rate was the maximum C_v of the valve (analogous to a perfect electrical voltage source with zero internal resistance, capable of sourcing any amount of current to a load).

If we plot the valve's performance in both scenarios (close-coupled to the dam, versus at the end of a restrictive pipe), we see the difference very clearly:



The “drooping” graph shows how the valve responds when it does not receive a constant pressure drop throughout the flow range. This is how the valve responds when *installed* in a non-ideal process, compared to the straight-line response it exhibits under *ideal* conditions of constant pressure. This is what we mean by “installed” characteristic versus “ideal” or “inherent” characteristic.

Pressure losses due to fluid friction as it travels down pipe is just one cause of valve pressure changing with flow. Other causes exist as well, including pump curves⁴² and frictional losses in other system components such as filters and heat exchangers. Whatever the cause, any piping system that fails to provide constant pressure across a control valve will “distort” the valve’s inherent characteristic in the same “drooping” manner, and this must be compensated in some way if we desire linear response from the valve.

Not only does the diminishing pressure drop across the valve mean we cannot achieve the same full-open flow rate as in the laboratory (with a constant pressure drop), but it also means the control valve responds with different amounts of sensitivity at various points along its range. Note how the installed characteristic graph is relatively steep at the beginning where the valve is nearly closed, and how the graph grows “flatter” at the end where the valve is nearly full-open. The rate of response (rate-of-change of flow Q compared to stem position x , which may be expressed as the derivative $\frac{dQ}{dx}$) is much greater at low flow rates than it is at high flow rates, all due to diminished pressure

⁴²The amount of fluid pressure output by any pump tends to vary with the fluid flow rate through the pump as well as the pump speed. This is especially true for centrifugal pumps, the most common pump design in process industries. Generally speaking, the discharge (output) pressure of a pump rises as flow rate decreases, and falls as flow rate increases. Variations in system fluid pressure caused by the pump constitutes one more variable for control valves to contend with.

drop at higher flow rates. This means the valve will respond more “sensitively” at the low end of its travel and more “sluggishly” at the high end of its travel.


From the perspective of a feedback control system, this varying valve responsiveness means the system will be unstable at low flow rates and unresponsive at high flow rates. At low flow rates – where the valve is nearly closed – any small movement of the valve stem will have a relatively large effect on fluid flow. However, at high flow rates, a much greater stem motion will be required to achieve a comparable effect on fluid flow. Thus, the control system will tend to over-react at low flow rates and under-react at high flow rates, simply because the control valve fails to exert the same degree of control over process flow at different flow rates. Oscillations may occur at low flow rates, and excessive deviations from setpoint at high flow rates as a result of this “distorted” valve behavior.

27.13.4 Characterized valve trim

The root cause of the problem – a varying pressure drop caused by frictional losses in the piping and other factors – generally cannot be eliminated. This means there is no way to regain maximum flow capacity short of replacing the control valve with one having a greater C_v rating⁴³. However, there is a clever way to flatten the valve’s responsiveness to achieve a more linear characteristic, and that is to purposely design the valve such that its inherent characteristic complements the process “distortion” caused by changing pressure drop. In other words, we design the control valve trim so it opens up gradually during the initial stem travel (near the closed position), then opens up more aggressively during the final stages of stem travel (near the full-open position). With the valve made to open up in a nonlinear fashion inverse to the “droop” caused by the installed pressure changes, the two non-linearities should cancel each other and yield a more linear response.

⁴³Even then, achieving the ideal maximum flow rate may be impossible. Our previous 100% flow rate for the valve was 80.5 GPM, but this goal has been rendered impossible by the narrow pipe, which according to the load line limits flow to an absolute maximum of 75 GPM (even with an infinitely large control valve).

This re-design will give the valve a nonlinear characteristic when tested in the laboratory with constant pressure drop, but the installed behavior should be more linear:



Now, control system response will be consistent at all points within the controlled flow range, which is a significant improvement over the original state of affairs.

Control valve trim is manufactured in a variety of different “characteristics” to provide the desired installed behavior. The two most common inherent characteristics are *linear* and *equal percentage*. “Linear” valve trim exhibits a fairly proportional relationship between valve stem travel and flow capacity (C_v), while “equal percentage” trim is decidedly nonlinear. A control valve with “linear” trim will exhibit consistent responsiveness only with a constant pressure drop, while “equal percentage” trim is designed to counter-act the “droop” caused by changing pressure drop when installed in a process system.

The C_v for linear and equal-percentage control valve trims are given by the following formulae⁴⁴:

$$C_v = xC_{vm} \quad \text{Linear trim}$$

$$C_v = C_{vm}R^{(x-1)} \quad \text{Equal percentage trim}$$

Where,

C_v = Flow coefficient of control valve at stem position x

C_{vm} = Flow coefficient of control valve while wide-open ($x = 100\%$)


x = Stem position, as a *per unit* value (ranging from 0 to 1) inclusive

R = Rangeability coefficient of equal-percentage trim

⁴⁴Note that the equal percentage formula given here can never achieve a C_v value of zero, regardless of stem position. This is untrue for real control valves, which of course achieve $C_v = 0$ when the stem is in the fully closed position. Therefore, the equal percentage formula shown here cannot be precisely trusted at small stem position values.

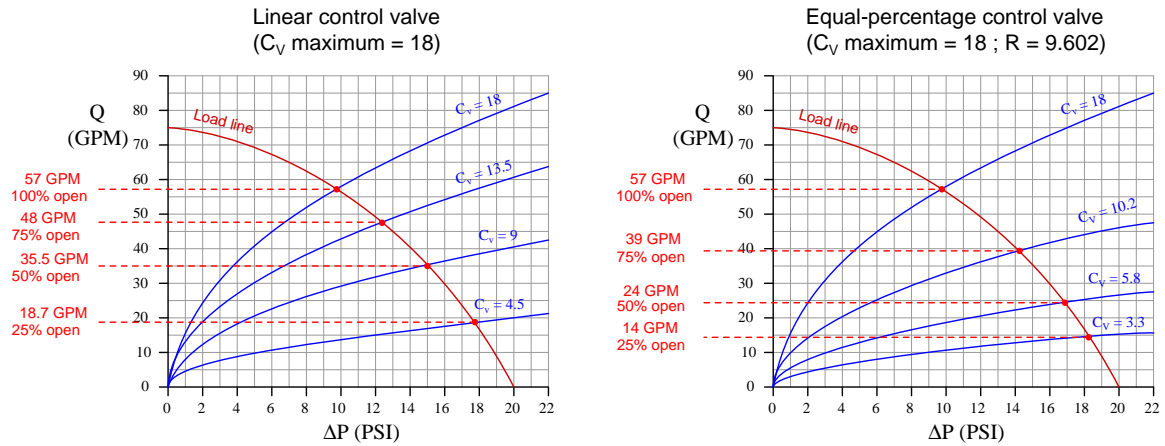
Another common inherent valve characteristic available from manufacturers is *quick-opening*, where the valve's C_v increases dramatically during the initial stages of opening, but then increases at a much slower rate for the rest of the travel. Quick-opening valves are often used in pressure-relief applications, where it is important to rapidly establish flow rate during the initial portions of valve stem travel.

The following pair of graphs show quick-opening, linear, and equal-percentage valve characteristics both as they are commonly presented in textbooks as well as based on real⁴⁵ control valve data from manufacturer's datasheets:



⁴⁵Data for the three graphs were derived from actual C_v factors published in Fisher's ED, EAD, and EDR sliding-stem control valve product bulletin (51.1:ED). I did not copy the exact data, however; I "normalized" the data so all three valves would have the exact same full-open C_v rating of 50.


When we compare the performance of an equal-percentage control valve against the linear control valve from the previous scenario (where water flowed from a dam through a long, narrow pipe) using the “load line” plot to determine flow rates, we see that the equal-percentage valve yields a more linear installed response than the inherently linear valve. You can see how the blue curves on these graphs (representing each control valve’s C_v at 25%, 50%, 75%, and 100% stem positions) are identical only at the wide-open position and differ at all other positions:



Stem position (%)	Flow rate through linear valve (GPM)	Flow rate through equal-percent valve (GPM)	Ideal flow rate (GPM)
0	0	0	0
25	18.7 (+4.45% from ideal)	14 (-0.25% from ideal)	14.25
50	35.5 (+7% from ideal)	24 (-4.5% from ideal)	28.5
75	48 (+5.25% from ideal)	39 (-3.75% from ideal)	42.75
100	57	57	57

Note that the equal-percentage valve with a maximum C_v of 18 does not yield any greater water flow rate at full-open than the linear valve with the same maximum C_v . No amount or type of valve characterization can make up for the pressure lost in restrictive piping. What equal-percentage characterization *does* accomplish is to make the relationship between flow rate and stem position closer to linear than it would be otherwise.

Different valve characterizations are achieved by different valve trim shapes. For instance, the plug profiles of a single-ported, stem-guided globe valve may be modified to achieve the common quick-opening, linear, and equal-percentage characteristics:




Photographs of linear (left) and equal-percentage (right) globe valve plugs having the same port size are shown side-by-side for comparison:




It should be clear⁴⁶ how the equal-percentage plug on the right-hand side retains more of its width along its length than the linear plug on the left-hand side. This means the equal-percentage plug is more restrictive than the linear plug for a greater portion of its withdrawal out of the seat. As each plug is drawn out of the seat's port by the actuator motion, the linear plug "opens up" more aggressively than the equal-percentage plug, even though both plugs are equally open when drawn fully out of the seat's port.

⁴⁶Astute readers will also note how the stem diameter of the left-hand (linear) plug is significantly greater than the stem diameter of the right-hand (equal-percentage) plug. This has nothing to do with characterization, and is simply an irrelevant difference between the two plugs. The truth of the matter is, dear reader, that these are the only two valve plugs I had on hand suitable for illustrating the difference between linear and equal-percentage trim. One just happened to have a thicker stem than the other.

Cage-guided globe valve trim characteristic is a function of port shape. As the plug rises up, the amount of port area uncovered determines the shape of the characteristic graph:



Ball valve trim characteristic is a function of notch shape. As the ball rotates, the amount of notch area opened to the fluid determines the shape of the characteristic graph. All valve trim in the following illustration is shown approximately half-open (50% stem rotation):



A different approach to valve characterization is to use a non-linear positioner function instead of a non-linear trim. That is, by “programming” a valve positioner to respond in a characterized fashion to command signals, it is possible to make an inherently linear valve behave as though it

were quick-opening, equal-percentage, or anywhere in between. All the positioner does is modify the valve stem position as per the desired characteristic function instead of proportionally follow the signal as it normally would.

This approach has the distinct advantage of convenience (especially if the valve is already equipped with a positioner) over changing the actual valve trim. However, if valve stem friction ever becomes a problem, its effects will be disproportionate along the valve travel range, as the positioner must position the valve more precisely in some areas of travel than others when pressed into service as a characterizer.

It should be noted that not all process control loops benefit from control valves with equal-percentage inherent characteristics. There are some process applications, for example, where the pressure drop across a control valve holds relatively constant over a wide range of valve flow rates⁴⁷, in which case an inherently linear control valve will indeed yield a linear installed characteristic. In other words, if the pressure drop across the control valve never changes much, what we have is an “ideal” scenario (i.e. no “droop” or “distortion”) that doesn’t need special valve characterization to behave linearly.

Some other process applications actually work quite well⁴⁸ with a quick-opening installed valve characteristic, in which case an inherently linear valve with varying pressure drop at different flow rates suffices. Ultimately, what matters is the valve stem position’s effect on the process variable and whether that relationship is consistent (linear) over a wide range of operating conditions.

27.14 Control valve problems

Control valves are subject to a number of common problems. This section is dedicated to an exploration of the more common control valve problems, and potential remedies.

⁴⁷Such applications are typically found when the purpose of the control valve is to regulate process fluid pressure. Consider, for example, a control valve regulating upstream gas pressure in a vessel by venting gas from that vessel to atmosphere. In such an application, the valve’s upstream pressure (P_1) will be nearly constant due to the control loop’s action, and the valve’s downstream pressure (P_2) will be constant due to it being atmospheric pressure.

⁴⁸An example of such a process is temperature control through a heat exchanger where the controlled fluid flow regime happens to transition from laminar to turbulent as the control valve opens further: at low stem positions (nearly shut) where the flow is laminar and heat transfer is impeded, large changes in flow rate may be necessary to effect modest changes in temperature; at high stem positions (nearly open) where the flow is turbulent and heat transfer is efficient, only small changes in flow rate are necessary to create modest changes in temperature. In such an application a quick-opening installed characteristic may actually yield more consistent behavior than a linear installed characteristic.

27.14.1 Mechanical friction

Control valves are mechanical devices with moving parts, and as such they are subject to *friction*, primarily between the valve stem and the stem packing. Some degree of friction is inevitable in valve packing⁴⁹ and in some types of trim where components must move past each other throughout the full range of valve stem travel (e.g. cage-guided globe valves, rotary ball valves) the trim itself adds additional friction beyond that which is imposed by the packing.


In physics, friction is classified as either *static* or *dynamic*. Static friction is defined as frictional force holding two stationary objects together. Dynamic friction is defined as frictional force impeding the motion of two objects sliding past each other. Static friction is always greater in magnitude than dynamic friction. Anyone who has ever pulled a sled through snow or ice knows that more force is required to “break” the sled loose from a stand-still (static friction) than is required to keep it moving (dynamic friction). The same holds true for friction in a control valve: the amount of force required to initially overcome static friction (i.e. initiate valve stem motion) usually exceeds the amount of force required to maintain the valve stem in motion.

The presence of friction in a control valve increases the force necessary from the actuator to cause valve movement. If the actuator is electric or hydraulic, the only real problem with increased force is the additional energy required from the actuator to move the valve (recall that mechanical *work* is the product of force and parallel displacement). If the actuator is pneumatic, however, a more serious problem arises from the combined effects of static and dynamic friction.

A simple “thought experiment” illustrates the problem. Imagine an air-to-open, sliding-stem control valve with the lower bench-set pressure applied to the pneumatic actuator. This should be the amount of pressure where the valve is just about to open from a fully-closed position. Now imagine slowly increasing the air pressure applied to the actuator. What should this valve do? If the spring tension is set properly, and there is negligible friction in the valve, the stem should smoothly rise from the fully-closed position as pressure increases beyond the bench-set pressure. However, what will this valve do if there is substantial friction present in the valve assembly? Instead of the stem smoothly lifting immediately as pressure exceeds the bench-set value, this valve will remain fully closed until enough *extra* pressure has accumulated in the actuator to generate a force large enough to overcome spring tension *plus* valve friction. Then, once the stem “breaks free” from static friction and begins to move, the stem will begin to accelerate because the actuator force now *exceeds* the sum of spring tension and friction, since dynamic friction is less than static friction. Compressed air trapped inside the actuator acts like a spring of its own, releasing stored energy. As the stem moves, however, the chamber volume in the diaphragm or piston actuator increases, causing pressure to drop, which causes the actuating force to decrease. When the force decreases sufficiently, the stem stops moving and static friction “grabs” it again. The stem will remain stationary until the applied pressure increases sufficiently again to overcome static friction, then the “slip-stick” cycle repeats.

⁴⁹Bellows seals are theoretically frictionless, but in practice bellows seals are almost always combined with standard packing to prevent catastrophic blow-out in the event of the bellows rupturing, and so the theoretical advantage of low friction is never realized.

If we graph the mechanical response of a pneumatic actuator with substantial stem friction, we see something like this:



What should be a straight, smooth line is reduced to a series of “stair-steps” as the combined effect of static and dynamic friction, plus the dynamic effects of a pneumatic actuator, conspire to make precise stem positioning nearly impossible. This effect is commonly referred to as *stiction*.


Even worse is the effect friction has on valve position when we *reverse* the direction of pressure change. Suppose that after we have reached some new valve position in the opening direction, we begin to ramp the pneumatic pressure downward. Due to static friction (again), the valve will *not* immediately respond by moving in the closed direction. Instead, it will hold still until enough pressure has been released to diminish actuator force to the point where there is enough unbalanced spring force to overcome static friction in the downward direction. Once this static friction is overcome, the stem will begin to accelerate downward because (lesser) dynamic friction will have replaced (greater) static friction. As the stem moves, however, air volume inside the actuating diaphragm or piston chamber will decrease, causing the contained air pressure to rise. Once this pressure rises enough that the stem stops moving downward, static friction will again “grab” the stem and hold it still until enough of a pressure change is applied to the actuator to overcome static friction.

What may not be immediately apparent in this second “thought experiment” is the amount of pressure change required to cause a reversal in stem motion compared to the amount of pressure change required to provoke continued stem motion in the same direction. In order to reverse the direction of stem motion, not only does the static friction have to be “relaxed” from the last movement, but additional static friction must be overcome in the opposite direction before the stem is able to move that way. We may clarify this concept by applying the problem-solving strategy of adding numerical quantities to the thought experiment. Suppose we have a control valve with a static packing friction of 50 pounds in either direction, and a diaphragm with a 12-inch diameter (113.1 square inches of area). According to the force/pressure/area formula ($F = PA$), an applied

pressure of 0.442 PSI will be required to overcome this static friction in either direction. That is to say, when moving in the upward direction, we must apply 0.442 PSI *more* pressure than is ideally required to achieve any given valve position. If after achieving this stem position we desire to move the valve downward to some new stem position, we must not only decrease the pressure by 0.442 PSI to relax the tension on the packing, but we must also decrease it *another 0.442 PSI* to overcome packing friction in the downward direction before the valve moves at all from its last position. Thus, a pressure reversal of 0.884 PSI (i.e. *twice* the equivalent value of packing friction) is required to make the valve reverse its direction of motion. This constitutes “deadband” in the control valve’s action, which degrades control behavior.

Thus, the effects of friction on a pneumatic control valve actuator may be quantified by subjecting the valve to small reversals in applied actuator pressure and measuring the resulting stem position. The largest increment of actuator pressure reversal resulting in zero stem motion represents the total amount of friction within the valve mechanism.

The following graph shows such a test, plotting actuator pressure over time as well as valve stem position over time. As the actuator pressure is stepped up and down in successively smaller intervals, the control valve’s stem position is seen to respond with less and less motion until it fails to respond at all:




Short of rebuilding a “sticky” control valve to replace damaged or worn components, there is not much that may be done to improve valve stiction other than regular lubrication of the packing (if appropriate). Lubrication is applied to the packing by means of a special *lubricator* device threaded into the bonnet of the valve:



As one of the common sources of excessive packing friction is over-tightening of the packing nuts by maintenance personnel eager to prevent process fluid leaks, a great deal of trouble may be avoided simply by educating the maintenance staff on the “care and feeding” of control valve packing for long service life.

Many modern digital valve positioners have the ability to monitor the drive force applied by an actuator on a valve stem, and correlate that force against stem motion. Consequently, it is possible to perform highly informative diagnostic tests on a control valve's mechanical "health," at least with regard to friction⁵⁰. For pneumatic and hydraulic actuators, actuator force is a simple and direct function of fluid pressure applied to the piston or diaphragm. For electric actuators, actuator force is an indirect function of electric motor current, or may be directly measured using load cells or springs and displacement sensors in the gear mechanism.

The following *valve signature* illustrates the kind of diagnostic "audit" that may be obtained from a digital control valve positioner based on actuator force (pneumatic air pressure) and stem motion:



This same diagnostic tool is useful for detecting trim seating problems in valve designs where there is sliding contact between the throttling element and the seat near the position of full closure (e.g. gate valves, ball valves, butterfly valves, plug valves, etc.). The force required to "seat" the valve into the fully-closed position will naturally be greater than the force required to move the throttling element during the rest of its travel, but this additional force should be smooth and consistent on the graph. A "jagged" force/travel graph near the fully-closed position indicates

⁵⁰Other measures of a control valve's mechanical status, such as flow capacity, flow characterization, and seat shut-off, cannot be inferred from measurements of actuator force and stem position.


interference between the moving element and the stationary seat, providing information valuable for predicting the remaining service life of the valve before the next rebuild.

27.14.2 Flashing

When a fluid passes through the constrictive passageways of a control valve, its average velocity increases. This is predicted by the Law of Continuity, which states that the product of fluid density (ρ), cross-sectional area of flow (A), and average velocity (\bar{v}) must remain constant for any flowstream:

$$\rho_1 A_1 \bar{v}_1 = \rho_2 A_2 \bar{v}_2$$

As fluid velocity increases through the constrictive passages of a control valve, the fluid molecules' kinetic energy increases. In accordance with the Law of Energy Conservation, potential energy in the form of fluid pressure must decrease correspondingly. Thus, fluid pressure decreases within the constriction of a control valve's trim as it throttles the flow, then increases (recovers) after leaving the constrictive passageways of the trim and entering the wider areas of the valve body:



If the fluid being throttled by the valve is a liquid (as opposed to a gas or vapor), and its absolute pressure ever falls below the vapor pressure⁵¹ of that substance, the liquid will begin to boil. This phenomenon, when it happens inside a control valve, is called *flashing*. As the graph shows, the point of lowest pressure inside the valve (called the *vena contracta* pressure, or P_{vc}) is the location where flashing will first occur, if it occurs at all.

Flashing is almost universally undesirable in control valves. The effect of boiling liquid at the point of maximum constriction is that flow through the valve becomes “choked” by the rapid expansion of liquid to vapor as it boils, degrading the valve’s flow capacity (i.e. decreasing the effective C_v). Flashing is also destructive to the valve trim, as boiling action propels tiny droplets of liquid at extremely high velocities past the plug and seat faces, eroding the metal over time.

⁵¹It should be noted that vapor pressure is a strong function of *temperature*. The warmer a liquid is, the more vapor pressure it will exhibit and thus the more prone it will be to flashing within a control valve.

A photograph showing a severely eroded valve plug (from a cage-guided globe valve) reveals just how destructive flashing can be:



A characteristic effect of flashing in a control valve is a “hissing” sound, reminiscent of what *sand* might sound like if it were flowing through the valve.

An important parameter predicting flashing in a control valve is the valve’s *pressure recovery factor*, based on a comparison of the valve’s total pressure drop from inlet to outlet versus the pressure drop from inlet to the point of minimum pressure within the valve.

$$F_L = \sqrt{\frac{P_1 - P_2}{P_1 - P_{vc}}}$$

Where,

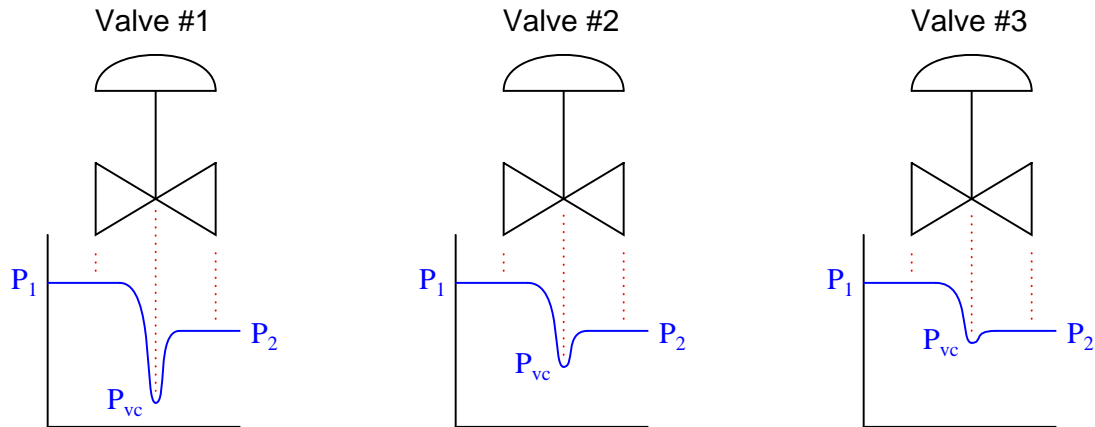
F_L = Pressure recovery factor (unitless)

P_1 = Absolute fluid pressure upstream of the valve

P_2 = Absolute fluid pressure downstream of the valve

P_{vc} = Absolute fluid pressure at the *vena contracta* (point of minimum fluid pressure within the valve)

The following set of illustrations shows three different control valves exhibiting the same permanent pressure drop ($P_1 - P_2$), but having different values of F_L :




Valve #1 exhibits the greatest pressure recovery (i.e. the amount that fluid pressure *increases* from the minimum pressure at the vena contracta to the downstream pressure: $P_2 - P_{vc}$) and the lowest F_L value. It is also the valve most prone to flashing in liquid service, because the vena contracta pressure is so much lower (all other factors being equal) than in the other two valves. If any of these valves will experience flashing in liquid service, it would be valve #1.

Valve #3, by contrast, has very little pressure recovery, and a large F_L value (nearly equal to 1). From the perspective of avoiding flashing, it is the best of the three valves to use for liquid service.

The style of valve (ball, butterfly, globe, etc.) is very influential on pressure recovery factor. The more convoluted the path for fluid within a control valve, the more opportunities that fluid will have to dissipate energy in turbulent motion, resulting in the greatest permanent pressure drop for the least amount of restriction at any single point in the flow's path.


Compare these two styles of valve to see which will have lowest pressure recovery factor and therefore be most prone to flashing:



Clearly, the globe valve does a better job of evenly distributing pressure losses throughout the path of flow. By contrast, the butterfly valve can only drop pressure at the points of constriction between the disk and the valve body, because the rest of the valve body is a straight-through path for fluid offering little restriction at all. As a consequence, the butterfly valve experiences a much lower vena contracta pressure (i.e. greater pressure recovery, and a lower F_L value) than the globe valve for any given amount of permanent pressure loss, making the butterfly valve more prone to flashing than the globe valve with all other factors being equal.

27.14.3 Cavitation

Fluid passing through a control valve experiences changes in velocity as it enters the narrow constriction of the valve trim (increasing velocity) then enters the widening area of the valve body downstream of the trim (decreasing velocity). These changes in velocity result in the fluid molecules' kinetic energies changing as well, in accordance with the kinetic energy equation $E_k = \frac{1}{2}mv^2$. In order that energy be conserved in a moving fluid stream, any increase in kinetic energy due to increased velocity must be accompanied by a complementary decrease in potential energy, usually in the form of fluid pressure. This means the fluid's pressure will fall at the point of maximum constriction in the valve (the *vena contracta*, at the point where the trim throttles the flow) and rise again (or *recover*) downstream of the trim:



If fluid being throttled is a liquid, and the pressure at the vena contracta is less than the vapor pressure of that liquid at the flowing temperature, the liquid will spontaneously boil. This is the phenomenon of *flashing* previously described. If, however, the pressure recovers to a point greater than the vapor pressure of the liquid, the vapor will re-condense back into liquid again. This is called *cavitation*.

As destructive as flashing is to a control valve, cavitation is worse. When vapor bubbles re-condense into liquid they often do so asymmetrically, one side of the bubble collapsing before the rest of the bubble. This has the effect of translating the kinetic energy of the bubble's collapse into a high-speed "jet" of liquid in the direction of the asymmetrical collapse. These liquid "microjets" have been experimentally measured at speeds up to 100 meters per second (over 320 feet per second). What is more, the pressure applied to the surface of control valve components in the path of these microjets is intense. Each microjet strikes the valve component surface over a very small surface area, resulting in a very high pressure ($P = \frac{F}{A}$) applied to that small area. Pressure estimates as high as 1500 newtons per square millimeter (1.5 *giga*-pascals, or about 220000 PSI!) have been calculated for cavitating control valve applications involving water.

No substance known is able to continuously withstand this form of abuse, meaning that cavitation *will destroy* any control valve given enough time. The effect of each microjet impinging on a metal surface is to carve out a small pocket in that metal surface. Over time, the metal will begin to take on a “pock-marked” look over the area where cavitation occurs. This stands in stark contrast to the visual appearance of flashing damage, which is smooth and polished.

Photographs of a fluted valve plug and its matching seat are shown here as evidence of flashing and cavitation damage, respectively:



The plug of this valve has been severely worn by flashing and cavitation. The flashing damage is responsible for the relatively smooth wear areas seen on the plug. Cavitation damage is most prominent inside the seat, where almost all the damage is in the form of pitting. The mouth of the seat exhibits smooth wear caused by flashing, but deeper inside you can see the pock-marked surface characteristic of cavitation, where liquid microjets literally blasted away pieces of metal. This trim set belongs to a Fisher Micro-Flat Cavitation valve, designed with process liquid flow passing down instead of up (i.e. first past the wide body of the plug and then down through the seat, rather than first up through the seat and then past the wide body of the plug). This trim design does not prevent cavitation (as clearly evidenced by the photos), but it does “move” the area of cavitation damage down below the seat’s sealing surface into a long tube extending below the seat. Although the ravages of flashing clearly took their toll on this valve’s trim, the valve would have been rendered inoperable much sooner had cavitation been at work along the plug’s length and at the sealing area where the plug contacts the seat.

The sound made by substantial liquid cavitation also contrasts starkly against the sound made by flashing. Whereas flashing sounds as though sand were flowing through the valve, cavitation produces a much louder “crackling” sound comprised of distinct impact pulses, reminiscent of what *gravel* or *rocks* might sound like if they were somehow forced to flow through the valve.

Sustained cavitation also has the detrimental effect of accelerating corrosion in certain process services. Bare metal surfaces are highly reactive with many chemical fluids, but become more resistant to further attack when a thin layer of reacted metal on the surface (the so-called *passivation layer*) acts as a sort of chemical barrier. Rust on steel, or the powdery-white oxide of aluminum are good examples: the initially bare metal surfaces react with their surrounding environment to form a protective outer layer, impeding further degradation of the metal beneath that layer. Cavitation works to blast away any protective layer that might otherwise accumulate, allowing corrosion to work at full speed until the entire thickness of the metal is corroded through. The

complementary destructive actions of cavitation and corrosion together is sometimes referred to as *cavitation corrosion*.


Several methods exist for abating cavitation in control valves:

1. Prevent flashing in the first place
2. Cushion with introduced gas
3. Sustain the flashing action (i.e. delay cavitation)


Cavitation abatement method #1 is quite simple to understand: if we prevent flashing from ever happening in a control valve, cavitation cannot follow. The key to doing this is making sure the vena contracta pressure never falls below the vapor pressure for the liquid. Several techniques exist for doing this:

- Select a control valve type having less pressure recovery (i.e. greater F_L value)
- Increase both upstream and downstream pressures by relocating the valve to a higher-pressure location in the process.
- Use multiple control valves in series to reduce the lowest pressure at either one
- Decrease the liquid's temperature (this decreases vapor pressure)
- Use cavitation-control valve trim


The last suggestion in this list deserves further exploration. Valve trim may be specially designed for cavitation abatement by providing multiple stages of pressure drop for the fluid as it passes through the trim. The following is a pressure versus location graph for a cavitating control valve. The liquid's vapor pressure is shown here as a dashed line marked P_{vapor} :



A valve equipped with cavitation-control trim will have a different pressure profile, with multiple *vena contracta* points where the fluid passes through a series of constrictions within the trim itself:



This way, the same final permanent pressure drop ($P_1 - P_2$) may be achieved without the lowest pressure ever falling below the liquid's vapor pressure limit. An example of cavitation-control design applied to cage-guided globe valve trim is shown in the following illustration:



Ball-style control valves, with their relatively high pressure recovery (low pressure recovery factor F_L values) are more prone to cavitation than globe valves, all other factors being equal. Special ball trim designed to help distribute pressure drops over a longer flow path is available, an example of this shown in the next photograph:




The round (ball-shaped) portion of the trim is on the far side of this piece, with the cavitation-controlling structure visible in the foreground. Fluid flow passing through the gap between the ball's edge and the valve seat spills into this multi-chambered structure where turbulence helps develop pressure drops at several locations. In a normal ball valve, there is only one location for any substantial pressure drop to develop, and that is at the narrow gap between the ball's edge and the seat. Here, multiple regions of pressure drop exist, with the intent of avoiding the liquid's vapor pressure limit at any one location, thus eliminating flashing and consequently eliminating cavitation.

Cavitation abatement method #2 is practical only in some process applications, where a non-reacting gas may be injected into the liquid stream to provide some "cushioning" within the cavitating region. The presence of non-condensable gas bubbles in the liquid stream disturbs the microjets' pathways, helping to dissipate their energy before striking the valve body walls.

Cavitation abatement method #3 involves a strategy opposite that of method #1. If, for whatever reason, we cannot avoid falling below the vapor pressure of the liquid as the flow stream moves through the valve, we may have the option of ensuring the downstream liquid pressure never rises above the liquid's vapor pressure, at least until the fluid clears past the valuable control valve and into an area of the system where cavitation damage will not be so expensive. This avoids cavitation at the cost of guaranteed flashing within the control valve, which is generally not as destructive as cavitation.

A pressure diagram shows how this method works:



Of course, flashing is not good for a control valve either. Not only does it damage the valve over time, but it also causes problems with flow capacity, as we will explore next.

27.14.4 Choked flow

Both gas and liquid control valves may experience what is generally known as *choked flow*. Simply put, “choked flow” is a condition where the rate of flow through a valve does not change substantially as downstream pressure is reduced.

Ideally, turbulent fluid flow rate through a control valve is a simple function of valve flow capacity (C_v) and differential pressure drop ($P_1 - P_2$), as described by the basic valve flow equation:


$$Q = C_v \sqrt{\frac{P_1 - P_2}{G_f}}$$

This equation simply does not apply for choked-flow conditions.

In a gas control valve, choking occurs when the velocity of the gas reaches the speed of sound for that gas. This is often referred to as *critical* or *sonic* flow. In a liquid control valve, choking occurs with the onset of flashing⁵². The reason sonic velocity is relevant to flow capacity for a control valve has to do with the propagation of pressure changes in fluids. Pascal’s principle tells us that changes in pressure within a closed fluid system will manifest at all points in the fluid system, but this never happens instantaneously. Instead, pressure changes propagate through any fluid at the speed of sound within that fluid. If a fluid stream happens to move at or above the speed of sound, pressure changes downstream are simply not able to overcome the stream’s velocity to affect anything upstream, which explains why the flow rate through a control valve experiencing sonic (critical) flow velocities does not change with changes in downstream pressure: those downstream pressure changes cannot propagate upstream against the fast-moving flow, and so will have no effect on the flow as it accelerates to sonic velocity at the point(s) of constriction.

Choked flow conditions become readily apparent if the flow-versus-pressure function of a control valve at any fixed opening value is graphed. The basic valve flow equation predicts a perfectly straight line at constant slope with flow rate (Q) as the vertical variable and the square root of pressure drop ($\sqrt{P_1 - P_2}$) as the horizontal variable. However, if we actually test a control valve by holding its upstream liquid pressure (P_1) constant and varying its downstream pressure (P_2) while maintaining a fixed stem position, we notice a point where flow reaches a maximum limit value:

⁵²The *Control Valve Sourcebook – Power & Severe Service* on page 6-3 and the *ISA Handbook of Control Valves* on page 211 both suggest that the mechanism for choking in liquid service may be related to the speed of sound just as it is for choked flow in gas services. Normally, liquids have higher sonic velocities than gases due to their far greater bulk moduli (incompressibility). This makes choking due to sonic velocity very unlikely in liquid flowstreams. However, when a liquid flashes into vapor, the speed of sound for that two-phase mixture of liquid and vapor will be much less than it is for the liquid itself, opening up the possibility of sonic velocity choking.



In a choked flow condition, further reductions in downstream pressure achieve no greater flow of liquid through the valve. This is not to say that the valve has reached a maximum flow – we may still increase flow rate through a choked valve by increasing its upstream pressure. We simply cannot coax any more flow through a choked valve by decreasing its *downstream* pressure.

An approximate predictor of choked flow conditions for gas valve service is the upstream-to-minimum absolute pressure ratio. When the vena contracta pressure is less than one-half the upstream pressure, both measured in absolute pressure units, choked flow is virtually guaranteed. One should bear in mind that this is merely an approximation and not a precise prediction for choked flow. Much more information is needed about the valve design, the particular process gas, and other factors in order to reliably predict the presence of choking.

Choked flow in liquid services is predicted when the vena contracta pressure equals the liquid's vapor pressure, since choking is a function of flashing for liquid flowstreams.

No attempt will be made in this book to explain sizing procedures for control valves in choked-flow service, due to the complexity of the subject.

An interesting and useful application of choked flow in gases is a device called a *critical velocity nozzle*. This is a nozzle designed to allow a fixed flow rate of gas through it given a known upstream pressure, and a downstream pressure that is sufficiently low to ensure sonic velocities in the nozzle throat. One practical use for critical velocity nozzles is in the flow testing of compressed air systems. One or more of these nozzles are connected to the main header line of an air compressor system and allowed to vent to atmosphere. So long as the compressor(s) are able to maintain constant header pressure, the flow rate of air through the nozzles(s) is guaranteed to be fixed, allowing a technician to monitor compressor parameters under precisely known load conditions.

27.14.5 Valve noise

A troublesome phenomenon in severe services is the audible noise produced by turbulence as the fluid moves through a control valve. Noise output is worse for gas services experiencing sonic (critical) flow and for liquid services experiencing cavitation, although it is possible for a control valve to produce substantial noise even when avoiding these operating conditions. Noise produced by a control valve also translates into vibration imposed on the piping, which may cause problems such as loosening of threaded fasteners over time.

One way to reduce noise output is to use special valve trim resembling the trim used to mitigate cavitation. A common cage-guided globe valve trim design for noise reduction uses a special cage designed with numerous, small holes for process gas to flow through. One of the ways in which these small holes reduce audible noise is by shifting the *frequency* of that noise upward. This increase in frequency places the sound outside the range where the human ear is most sensitive to noise, and it also helps to reduce noise coupling to the piping, confining most of the noise “power” to the internal volume of the process fluid rather than radiating outward into the air.

Fisher manufactures a series of noise-abatement trim for process gas service called *Whisper* trim. A “Whisper” plug and cage set is shown in this next photograph:



In some versions, the holes are merely straight through the cage wall. In more sophisticated versions of *Whisper* trim (particularly the “*WhisperFlo*”), the small holes lead to a labyrinth of passages designed to dissipate energy by forcing the fluid to take several sharp turns as it passes through the wall of the cage. This allows relatively large pressure drops to develop without high fluid velocities, which is the primary causal factor for noise in control valves.

27.14.6 Erosion

A problem common to control valves used in *slurry* service (where the process fluid is a liquid containing a substantial quantity of hard, solid particles) is *erosion*, where the valve trim and body are worn by the passage of solid particles. Erosion produces some of the most striking examples of valve damage, as shown by the following photographs⁵³:



Here we see large holes worn in a globe valve plug, and substantial damage done to the seat as well. The process service in this case was water with “coke fines” (small particles of coke, a solid petroleum product). Even ceramic valve seat components are not immune to damage from slurry service, as revealed by this photograph of a valve seat with a notch worn by slurry flow:



There really is no good way to reduce the effects of erosion damage from slurry flows, other than to use exceptionally hard valve trim materials. Even then, the control valve must be considered a fast-wearing component (along with pumps and any other components in harm's way of the slurry flowstream), rebuilt or replaced at regular intervals.

⁵³A colleague of mine humorously refers to these valve trim samples as “shock and awe,” because they so dramatically reveal the damaging nature of certain process fluid services.

Another cause of erosion in control valves is *wet steam*, where steam contains droplets of liquid water propelled at high velocity by the steam flow. A dramatic example of wet steam damage appears in this next photograph, where the cage from a Fisher valve has been literally cut in half from the flow:



Steam may also “cut” other parts of a valve if allowed to leak past. Here, we see a valve bonnet with considerable damage caused by steam leaking past the *outside* of the packing, between the packing rings and the bonnet’s bore:



Any fluid with sufficient velocity may cause extensive damage to valve components. Small holes developing in the body of a valve may become large holes over time, as fluid rushes through the hole. The initial cause of the hole may be a manufacturing defect (such as *porosity* in the metal casting) or damage inflicted by the user (e.g. a crack in the valve body caused by some traumatic event). An example of such a hole in a valve body becoming worse over time appears in this next photograph, taken of a control valve removed after *40 years* of continuous service:



A close-up photograph of this same hole shows the leak path worn larger by passing fluid, allowing fluid to flow by the seat even with the valve in the fully-closed position:



In more severe process services, such holes rapidly grow in size. This next photograph shows a rather extreme example of a hole near the seat of a valve body, enlarged to the point where the valve is hardly capable of restricting fluid flow at all because the hole provides a bypass route for flow around the valve plug and seat:



27.14.7 Chemical attack

Corrosive chemicals may attack the metal components of control valves if those components are not carefully selected for the proper service. A close-up photograph of a chemically-pitted valve plug shows pitting characteristic of chemical attack:



As mentioned previously in this chapter, the effects of corrosion are multiplied when combined with the effects of cavitation. Most metals develop what is known as a *passivation layer* in response to chemical attack. The outer layer of metal corrodes, but the byproduct of that corrosion is a relatively inert compound acting to shield the rest of the metal from further attack. Rust on steel, or aluminum oxide on aluminum, are both common examples of passivation layers in response to oxidation of the metal. When cavitation happens inside a valve, however, the extremely high pressures caused by the liquid microjets will blast away any protection afforded by the passivation layer, allowing chemical attack to begin anew. The result is rapid degradation of the valve components.

27.15 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Definition of pressure:** $P = \frac{F}{A}$ (pressure is the amount of force applied over a specified area by a fluid. Relevant to fluid-powered valve actuators: the amount of force developed by an actuator is proportional to the fluid pressure applied to it and the area of its piston or diaphragm.
- **Pascal's principle:** changes in fluid pressure are transmitted evenly throughout an enclosed fluid volume. Relevant to fluid-powered valve actuators.
- **Fluid seals:** accomplished by maintaining tight contact between solid surfaces. Relevant to control valve plug and seat design, where even seating force is critical to maintaining tight shut-off. Also relevant to valve packing design.
- **Conservation of energy:** energy cannot be created or destroyed, only converted between different forms. Relevant to fluid velocities and pressures inside of a control valve.
- **Bernoulli's equation:** $z_1\rho g + \frac{v_1^2\rho}{2} + P_1 = z_2\rho g + \frac{v_2^2\rho}{2} + P_2$, which is an application of the Law of Energy Conservation, stating that the sum of all forms of energy in a moving fluid stream (height, kinetic, and pressure) must remain the same. Relevant to calculations of pressure drop and pressure recovery across control valve trim.
- **Conservation of mass:** mass is an intrinsic property of matter, and as such cannot be created or destroyed. Relevant to the Continuity Principle for moving fluids, where the mass flow rate of a fluid entering a control valve must equal the mass flow rate exiting the valve.
- **Linear equations:** any function represented by a straight line on a graph may be represented symbolically by the slope-intercept formula $y = mx + b$. Relevant to control valve positioners and split-ranging.
- **Zero shift:** any shift in the offset of an instrument is fundamentally additive, being represented by the "intercept" (b) variable of the slope-intercept linear formula $y = mx + b$. Relevant to control valve calibration: adjusting the "zero" of a valve positioner always adds to or subtracts from the valve stem position.
- **Span shift:** any shift in the gain of an instrument is fundamentally multiplicative, being represented by the "slope" (m) variable of the slope-intercept linear formula $y = mx + b$. Relevant to control valve calibration: adjusting the "span" of a valve positioner always multiplies or divides the stem stroke (compared to signal span).
- **Deadband and hysteresis:** the difference in response with the independent variable increasing versus decreasing. Usually caused by friction in a mechanism. Relevant to control valve response: a control valve will not go to the same position with its command signal at

some value increasing as it does at that same value decreasing, due to friction in the trim and packing.

- **Inverse mathematical functions:** an inverse function, when applied to the result of its counterpart function, “un-does” the operation and leaves you with the original quantity. Relevant to control valve characterization, where the inherently nonlinear response of a control valve installed in a process with variable pressure drop is made to behave more linearly by skewing the shape of the valve trim in an inverse manner.
- **Negative feedback:** when the output of a system is degeneratively fed back to the input of that same system, the result is decreased (overall) gain and greater stability. Relevant to control valve positioners as well as self-actuated regulators.
- **Self-balancing pneumatic mechanisms:** all self-balancing pneumatic instruments work on the principle of negative feedback maintaining a nearly constant baffle-nozzle gap. Force-balance mechanisms maintain this constant gap by balancing force against force with negligible motion, like a tug-of-war. Motion-balance mechanisms maintain this constant gap by balancing one motion with another motion, like two dancers moving in unison.

References

Baumann, Hans D., *Control Valve Primer, A User's Guide*, Second Edition, Instrument Society of America, Research Triangle Park, NC, 1994.

Brestel, Ronald; Hutchens, Wilbur; Wood, Charles, *Control Valve Packing Systems*, Technical Monograph 38, Fisher Controls International Inc., Marshalltown, IA, 1992.

“Cavitation in Control Valves”, document L351 EN, Samson AG, Frankfurt, Germany.

Control Valve Handbook, Third Edition, Fisher Controls International, Inc., Marshalltown, IA, 1999.

Control Valve Sourcebook – Power & Severe Service, Fisher Controls International, Inc., Marshalltown, IA, 1988.

“Design ED, EAD, and EDR Sliding-Stem Control Valves”, Product Bulletin 51.1:ED, Fisher, Marshalltown, IA, 2006.

Grumstrup, Bruce, *Considerations in the Design and Selection of Bellows Seal Equipment Valves*, Technical Monograph 37, Fisher Controls International Inc., Marshalltown, IA, 1991.

Hutchison, J.W., *ISA Handbook of Control Valves*, Second Edition, Instrument Society of America, Research Triangle Park, NC, 1976.

Jury, Floyd D., *Fundamentals of Aerodynamic Noise in Control Valves*, Technical Monograph 43, Fisher Controls International Inc., Marshalltown, IA, 1999.

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Control Volume II*, Third Edition, CRC Press, Boca Raton, FL, 1999.

“Micro Trims for Globe and Angle Valve Applications”, Product Bulletin 80.4:010, Emerson Process Management, Marshalltown, IA, 2005.

“Packing Selection Guidelines for Sliding-Stem Valves”, Product Bulletin 59.1:062, Emerson Process Management, Marshalltown, IA, 2007.

“Pipeline Accident Report – Pipeline Rupture and Subsequent Fire in Bellingham, Washington June 10, 1999”, NTSB/PAR-02/02, PB2002-916502, Notation 7264A, National Transportation Safety Board, Washington DC, 2002.

Richardson, Jonathan W., *Primary Seat Shutoff*, Technical Monograph 47, Fisher Controls International LLC, Marshalltown, IA, 2005.

Riveland, Marc, *Fundamentals of Valve Sizing for Liquids*, Technical Monograph 30, Fisher Controls International Inc., Marshalltown, IA, 1985.

Schafbuch, Paul, *Fundamentals of Flow Characterization*, Technical Monograph 29, Fisher Controls International Inc., Marshalltown, IA, 1985.

Warnett, Chris, *Using Valve Actuators as Predictive Maintenance Tools for MOVs*, Rotork Controls, Inc., Rochester, NY, 2000.


“Valve Sizing Technical Bulletin”, document MS-06-84-E, revision 3, Swagelok Company, MI, 2002.

Chapter 28

Variable-speed motor controls

An alternative to control valves for adjusting fluid flow is to regulate the speed of the machine(s) motivating fluid to flow. In the case of liquid flow control, this would take the form of variable-speed pumps. In the case of gas flow control, it would mean varying the rotational speed of compressors or blowers.

Flow control by machine speed control makes a lot of sense for some process applications. It is certainly more energy-efficient¹ to vary the speed of the machine pushing fluid to control flow, as opposed to letting the machine run at full speed all the time and adjusting flow rate by throttling the machine's discharge (outlet) or recycling fluid back to the machine's suction (inlet). The fact that the system has one less component in it (no control valve) also reduces capital investment and potentially increases system reliability:



Modern power electronics provide the means to electronically control the speed of almost any type and size of electric motor, using a device called a motor *drive*. DC motor drives vary voltage and current to the armature and field windings of the motor. In general, DC motor speed is directly proportional to armature voltage, and inversely proportional to field current. AC motor drives

¹Regulating fluid flow by using a throttling valve along with a constant-speed pump is analogous to regulating an automobile's speed by applying varying force to the brake pedal while holding the accelerator pedal at its full-power position!

vary the *frequency*² of the power applied to the motor's stator windings, because frequency is what establishes the speed of the stator's rotating magnetic field which the rotor follows.

DC motors were once considered clearly superior to AC motors in variable-speed applications where high starting torque (torque generated at zero speed) was needed. The advent of sophisticated variable-frequency drive (VFD) electronics, however, greatly expanded the useful operating speed range of AC induction motors to the point where one can do almost any task³ with an AC motor that used to be possible only with a DC motor. This is highly advantageous, because AC induction motors are much simpler and more reliable machines than DC motors. DC motors use commutators and brushes to conduct electrical power to their rotating armatures, both of these components being subject to wear. AC induction motors convey power to their rotors by electromagnetic induction, not by direct contact, and so neither commutators nor brushes are necessary. In fact, the only "wearing" component in an AC induction motor are the bearings holding the shaft, which of course are common to *all* rotating machines and therefore not a liability peculiar to AC induction motors.

Further advantages of electric motor speed control, whether DC or AC, include the ability to "soft-start" the machine instead of always accelerating rapidly from a full stop to full speed. Soft-starting electric motors greatly reduces the wear on machines, increasing their service life. In applications such as conveyor belt control, robotic machine motion control, and electric vehicle propulsion, variable-speed motor technology makes perfect sense as a control mechanism because the prime mover device is already (in most cases) an electric motor, with precise speed control of that motor yielding many practical benefits. In some applications, *regenerative braking* may be possible: where the motor is used as an electrical generator to slow down the machine on command. Regenerative braking transfers kinetic energy within the machine back to the power grid where it may be productively used in other processes, saving energy and reducing wear on any mechanical (friction) brakes already installed in the machine.

With all these advantages inherent to variable-speed pumps, fans, and compressors (as opposed to using dissipative control valves), one might wonder, "Why would anyone *ever* use a control valve to regulate flow? Why not control *all* fluid flows using variable-speed pumping machines?" Several good answers exist to this question:

- Variable-speed machines often cannot increase or decrease fluid flow rates as rapidly as control valves
- Control valves have the ability to positively halt flow; a stopped pump or blower will not necessarily prevent flow from going through
- Some process applications *must* contain a dissipative element in order for the system to function (e.g. let-down valves in closed-loop refrigeration systems)
- Split-ranging may be difficult or impossible to achieve with multiple machine speed control
- Limited options for fail-safe status
- In many cases, there is no machine dedicated to a particular flow path (e.g. a pressure release valve, or a valve controlling water flow from a dam) for us to control the speed of


²AC drives also vary the amount of voltage applied to the motor along with frequency, but this of secondary importance to the varying of frequency to control speed.

³This includes using an AC induction motor as a *servo* for precise positioning control!

28.1 DC motor speed control


DC electric motors generate torque by a reaction between two magnetic fields: one field established by stationary “field” windings (coils), and the other by windings in the rotating armature. Some DC motors lack field windings, substituting large permanent magnets in their place so that the stationary magnetic field is constant for all operating conditions.

In any case, the operating principle of a DC electric motor is that current passed through the armature creates a magnetic field that tries to align with the stationary magnetic field. This causes the armature to rotate:



However, a set of segmented copper strips called a *commutator* breaks electrical contact with the now-aligned coil and energizes another coil (or in the simple example shown above, it re-energizes the same loop of wire in the opposite direction) to create another out-of-alignment magnetic field that continues to rotate the armature. Electrical contact between the rotating commutator segments and the stationary power source is made through carbon *brushes*. These brushes wear over time (as does the commutator itself), and must be periodically replaced.

Most industrial DC motors are built with multiple armature coils, not just one as shown in the simplified illustration above. A photograph of a large (1250 horsepower) DC motor used to propel a ferry ship is shown here, with the field and armature poles clearly seen (appearing much like spokes in a wheel):



A close-up of one brush assembly on this large motor shows both the carbon brush, the brush's spring-loaded holder, and the myriad of commutator bars the brush makes contact with as the armature rotates:



DC motors exhibit the following relationships between mechanical and electrical quantities:


Torque:

- Torque is directly proportional to armature magnetic field strength, which in turn is directly proportional to current through the armature windings
- Torque is also directly proportional to the stationary pole magnetic field strength, which in turn is directly proportional to current through the field windings (in a motor with non-permanent field magnets)

Speed:

- Speed is limited by the counter-EMF generated by the armature as it spins through the stationary magnetic field. This counter-EMF is directly proportional to armature speed, and also directly proportional to stationary pole magnetic field strength (which is directly proportional to field winding current in a motor that is not permanent-magnet)
- Thus, speed is directly proportional to armature voltage
- Speed is also inversely proportional to stationary magnetic field strength, which is directly proportional to current through the field windings (in a motor with non-permanent field magnets)


A very simple method for controlling the speed and torque characteristics of a wound-field (non-permanent magnet) DC motor is to control the amount of current through the field winding:




Decreasing the field control resistor's resistance allows more current through the field winding, strengthening its magnetic field. This will have two effects on the motor's operation: first, the motor will generate more torque than it did before (for the same amount of armature current) because there is now a stronger magnetic field for the armature to react against; second, the motor's speed will decrease because more counter-EMF will be generated by the spinning armature for the same rotational speed, and this counter-EMF naturally attempts to equalize with the applied DC source voltage. Conversely, we may increase a DC motor's speed (and reduce its torque output) by increasing the field control resistor's resistance, weakening the stationary magnetic field through which the armature spins.

Regulating field current may alter the balance between speed and torque, but it does little to control total motor *power*. In order to control the power output of a DC motor, we must also regulate armature voltage and current. Variable resistors may also be used for this task, but this is generally frowned upon in modern times because of the wasted power.

A better solution is to have an electronic power control circuit very rapidly switch transistors on and off, switching power to the motor armature. This is called *pulse-width modulation*, or *PWM*.




The *duty cycle* (on time versus on+off time) of the pulse waveform will determine the fraction of total power delivered to the motor:




Such an electronic power-control circuit is generally referred to as a *drive*. Thus, a *variable-speed drive* or *VSD* is a high-power circuit used to control the speed of a DC motor. Motor drives may be manually set to run a motor at a set speed, or accept an electronic control signal to vary the motor speed in the same manner an electronic signal commands a control valve to move. When equipped with remote control signaling, a motor drive functions just like any other final control element: following the command of a process controller in order to stabilize some process variable at setpoint.

An older technology for pulsing power to a DC motor is to use a *controlled rectifier* circuit, using SCRs instead of regular rectifying diodes to convert AC to DC. Since the main power source of most industrial DC motors is AC anyway, and that AC must be converted into DC at some point in the system, it makes sense to integrate control right at the point of rectification:



Controlled rectifier circuits work on the principle of varying the “trigger” pulse times relative to the AC waveform pulses. The earlier the AC cycle each SCR is triggered on, the longer it will be on to pass current to the motor. The “phase control” circuitry manages all this pulse timing and generation.


A DC motor drive that simply varied power to the motor according to a control signal would be crude and difficult to apply to the control of most processes. What is ideally desired from a variable-speed drive is precise command over the motor's *speed*. For this reason, most VSDs are designed to receive feedback from a tachometer mechanically connected to the motor shaft, so the VSD "knows" how fast the motor is turning. The tachometer is typically a small DC generator, producing a DC voltage directly proportional to its shaft speed (0 to 10 volts is a common scale). With this information, the VSD may throttle electrical power to the motor as necessary to achieve whatever speed is being commanded by the control signal. Having a speed-control feedback loop built into the drive makes the VSD a "slave controller" in a cascade control system, the drive receiving a speed setpoint signal from whatever process controller is sending an output signal to it:



A photograph of the tachogenerators (dual, for redundancy) mechanically coupled to that large 1250 horsepower ferry ship propulsion motor appears here:



The SCRs switching power to this motor may be seen here, connected via twisted-pair wires to control boards issuing “firing” pulses to each SCR at the appropriate times:



The integrity of the tachogenerator feedback signal to the VSD is extremely important for safety reasons. If the tachogenerator becomes disconnected – whether mechanically or electrically (it doesn’t matter) – from the drive, the drive will “think” the motor is not turning. In its capacity as a speed *controller*, the drive will then send full power to the DC motor in an attempt to get it up to speed. Thus, loss of tachogenerator feedback causes the motor to immediately “run away” to full speed. This is undesirable at best, and likely dangerous in the case of motors as large as the one powering this ship.

As with all forms of electric power control based on pulse durations and duty cycles, there is a lot of electrical “noise” broadcast by VSD circuits. Square-edged pulse waveforms created by the rapid on-and-off switching of the semiconductor power devices are equivalent to infinite series


of high-frequency sine waves⁴, some of which may be of high enough frequency to self-propagate through space as electromagnetic waves. This *radio-frequency interference* or *RFI* may be quite severe given the high power levels of industrial motor drive circuits. For this reason, it is *imperative* that neither the motor power conductors nor the conductors feeding AC power to the drive circuit be routed anywhere near small-signal or control wiring, because the induced noise *will* wreak havoc with whatever systems utilize those low-level signals.

RFI noise on the AC power conductors may be reduced by routing the AC power through low-pass filter circuits called *line reactors* placed near the drive. These line reactors, consisting of ferrous metal core inductors wired in series with the drive, block high-frequency noise from propagating back to the rest of the AC power distribution wiring where it may influence other electronic equipment. However, there is little that may be done about the RFI noise between the drive and the motor other than to shield the conductors inside of well-grounded metallic conduit and/or to use grounded-shield power cables.

⁴This equivalence was mathematically proven by Jean Baptiste Joseph Fourier (1768-1830), and is known as a *Fourier series*.

28.2 AC motor speed control

AC induction motors are based on the principle of a *rotating magnetic field* produced by a set of stationary windings (called *stator* windings) energized by AC power of different phases. The effect is not unlike a series of blinking “chaser” light bulbs which appear to “move” in one direction due to the blinking sequence. If sets of wire coils (windings) are energized in a like manner – each coil reaching its peak field strength at a different time from its adjacent neighbor – the effect will be a magnetic field that “appears” to move in one direction. If these windings are oriented around the circumference of a circle, the moving magnetic field rotates about the center of the circle, as illustrated by this sequence of images (read left-to-right, top-to-bottom, as if you were reading words in a sentence):



To view a flip-book animation of this same sequence, turn to Appendix [A.2](#) beginning on page [2840](#).

Any magnetized object placed in the center of this circle will attempt to spin at the same rotational speed as the rotating magnetic field. *Synchronous* AC motors use this principle, where a magnetized rotor precisely follows the magnetic field’s speed.

Any electrically conductive object placed in the center of the circle will experience *induction* as the magnetic field direction changes around the conductor. This will induce electric currents within the conductive object, which in turn will react against the rotating magnetic field in such a way that the object will be “dragged along” by the field, always lagging a bit in speed. *Induction* AC motors use this principle, where a non-magnetized (but electrically conductive) rotor rotates at a speed slightly less⁵ than the synchronous speed of the rotating magnetic field.

⁵The difference between the synchronous speed and the rotor’s actual speed is called the motor’s *slip speed*.

The rotational speed of this magnetic field is directly proportional to the frequency of the AC power, and inversely proportional to the number of poles in the stator:

$$S = \frac{120f}{n}$$

Where,

S = Synchronous speed of rotating magnetic field, in revolutions per minute (RPM)

f = Frequency, in cycles per second (Hz)

n = Total number of stator poles per phase (the simplest possible AC induction motor design will have two poles)

The relationship between synchronous speed, frequency, and pole number may be understood by analogy: the speed at which the lights in a “chaser” light array appear to move is a function of the blinking frequency and the number of light bulbs per unit length. If the number of light bulbs in such an array is doubled by placing additional bulbs between the existing bulbs (so as to maintain the same array length), the apparent speed will be cut in half: with less distance between each pair of bulbs, it takes more cycles (more “blinks”) for the sequence to travel the entire length of the array. Likewise, an AC stator with more poles in its circumference will require more cycles of AC power for the rotating magnetic field to complete one revolution.

A *synchronous* AC motor will spin at the exact same speed as the rotating magnetic field: a practical example is a 4-pole synchronous motor spinning at 1800 RPM with an applied power frequency of 60 Hz. An *induction* AC motor will spin at slightly less than the speed of the magnetic field: a practical example is a 4-pole induction motor spinning at 1720 RPM with an applied power frequency of 60 Hz (i.e. 80 RPM “slip” speed). Induction motors are simpler both in construction and operation, making them the most popular of the two types of AC electric motors in industry.

While the number of poles in the motor’s stator is a quantity fixed⁶ at the time of the motor’s manufacture, the frequency of power we apply may be adjusted with the proper electronic circuitry. A high-power circuit designed to produce varying frequencies for an AC motor to run on is called a *variable-frequency drive*, or *VFD*.


Variable-frequency motor drives are incredibly useful devices, as they allow what would normally be a fixed-speed electric motor to provide useful power over a wide range of speeds. The benefits of variable-speed operation include reduced power consumption (only spinning the motor as fast as it needs to move, and no faster), reduced vibration (less speed = reduced vibrational forces), and the ability to ramp the motor’s speed up and down for reduced wear and tear on mechanical components resulting from acceleration forces.

Another feature common to most VFDs is the ability to actively *brake* the load. This is when the drive causes the motor to actively apply a negative torque to the load to slow it down. Some VFDs even provide means to recover the kinetic energy of the load during the braking process, resulting in further energy savings.

⁶Multi-speed motors do exist, with selectable pole configurations. An example of this is an electric motor with extra sets of stator windings, which may be connected to form a 4-pole configuration for high speed, and an 8-pole configuration for low speed. If the normal full-load “high” speed for this motor is 1740 RPM, the normal full-load “low” speed will be approximately half that, or 870 RPM. Given a fixed line frequency, this motor will only have these two speeds to choose from.

Variable-frequency AC motor drives consist of electronic components to convert the constant-frequency AC input power into variable-frequency (and variable-voltage) AC output power for the motor to run on. This usually takes place in three distinct sections. The *rectifier* section uses diodes to convert line AC power into DC. The *filter* “smooths” the rectified DC power so it has little ripple voltage. Lastly, the *inverter* section re-converts the filtered DC power back into AC, only this time at whatever levels of frequency and voltage is desired to run the motor at different speeds.


A simplified schematic diagram for a VFD is shown here, with a rectifier section on the left (to convert AC input power into DC), a filter capacitor to “smooth” the rectified DC power, and a transistor “bridge” to switch DC into AC at whatever frequency is desired to power the motor⁷. The transistor control circuitry has been omitted from this diagram for the sake of simplicity:



As with DC motor drives (VSDs), the power transistors in an AC drive (VFD) switch on and off very rapidly with a varying duty cycle. Unlike DC drives, however, the duty cycle of an AC drive’s power transistors must vary rapidly in order to synthesize an AC waveform from the DC “bus” voltage following the rectifier. A DC drive circuit’s PWM duty cycle controls motor power, and so it will remain at a constant value when the desired motor power is constant. Not so for an AC motor drive circuit: its duty cycle must vary from zero to maximum and back to zero repeatedly in order to create an AC waveform for the motor to run on.

⁷Note the reverse-connected diodes across the source and drain terminals of each power transistor. These diodes serve to protect the transistors against damage from reverse voltage drop, but they also permit the motor to “back feed” power to the DC bus (acting as a *generator*) when the motor’s speed exceeds that of the rotating magnetic field, which may happen when the drive commands the motor to slow down. This leads to interesting possibilities, such as *regenerative braking*, with the addition of some more components.

The equivalence between a rapidly-varied pulse-width modulation (PWM) waveform and a sine wave is shown in the following illustration:



This concept of rapid PWM transistor switching allows the drive to “carve” any arbitrary waveform out of the filtered DC voltage it receives from the rectifier. Virtually any frequency may be synthesized (up to a maximum limited by the frequency of the PWM pulsing), and any voltage (up to a maximum peak established by the DC bus voltage), giving the VFD the ability to power an induction motor over a wide range of speeds.

While frequency control is the key to synchronous and induction AC motor speed control, it is generally not enough on its own. While the speed of an AC motor is a direct function of frequency (controlling how fast the rotating magnetic field rotates around the circumference of the stator), torque is a function of stator current. Since the stator windings are inductors by nature, their reactance varies with frequency as described by the formula $X_L = 2\pi fL$. Thus, as frequency is increased, winding reactance increases right along with it. This increase in reactance would result in decreased stator current if the VFD’s output voltage remained constant. This undesirable scenario would result in torque loss at high speeds, and excessive torque (as well as excessive stator heat!) at low speeds. For this reason, the AC voltage output by a VFD is made to vary⁸ in proportion to the applied frequency, so that the stator current will remain within good operating limits throughout the speed range of the VFD. This correspondence is called the *voltage-to-frequency ratio*, abbreviated “V/F” ratio or “V/Hz” ratio.

To give an example of a VFD programmed with a constant V/F ratio, if the output line voltage to the motor is 480 volts RMS at full speed (60 Hz), then the output line voltage should be 240 volts RMS at half-speed (30 Hz), and 120 volts RMS at quarter-speed (15 Hz).

⁸The VFD achieves variable output voltage using the same technique used to create variable output frequency: rapid pulse-width-modulation of the DC bus voltage through the output transistors. When lower output voltage is necessary, the duty cycle of the pulses are reduced throughout the cycle (i.e. transistors are turned on for shorter periods of time) to generate a lower average voltage of the synthesized sine wave.

Variable-frequency motor drives are manufactured for industrial motor control in a wide range of sizes and horsepower capabilities. Some VFDs are small enough to hold in your hand, while others are large enough to require a freight train for transport. The following photograph shows a pair of moderately-sized Allen-Bradley VFDs (about 100 horsepower each, standing about 4 feet high), used to control pumps at a wastewater treatment plant:



Variable-frequency AC motor drives do not require motor speed feedback the way variable-speed DC motor drives do. The reason for this is quite simple: the controlled variable in an AC drive is the frequency of power sent to the motor, and rotating-magnetic-field AC motors are *frequency-controlled* machines by their very nature. For example, a 4-pole AC induction motor powered by 60 Hz has a base speed of 1728 RPM (assuming 4% slip). If a VFD sends 30 Hz AC power to this motor, its speed will be approximately half its base-speed value, or 864 RPM. There is really no need for speed-sensing feedback in an AC drive, because the motor's real speed will always be limited by the drive's output frequency. To control frequency *is* to control motor speed for AC synchronous and induction motors, so no tachogenerator feedback is necessary for an AC drive to “know” approximately⁹ how fast the motor is turning. The non-necessity of speed feedback for AC drives eliminates a potential safety hazard common to DC drives: the possibility of a “runaway” event where the drive loses its speed feedback signal and sends full power to the motor.

As with DC motor drives, there is a lot of electrical “noise” broadcast by VFD circuits. Square-edged pulse waveforms created by the rapid on-and-off switching of the power transistors are equivalent to infinite series of high-frequency sine waves¹⁰, some of which may be of high enough frequency to self-propagate through space as electromagnetic waves. This *radio-frequency interference* or *RFI* may be quite severe given the high power levels of industrial motor drive circuits. For this reason, it is *imperative* that neither the motor power conductors nor the conductors feeding

⁹For more precise control of AC motor speed (especially at low speeds where slip speed becomes a greater percentage of actual speed), speed sensors may indeed be necessary.

¹⁰This equivalence was mathematically proven by Jean Baptiste Joseph Fourier (1768-1830), and is known as a *Fourier series*.

AC power to the drive circuit be routed anywhere near small-signal or control wiring, because the induced noise *will* wreak havoc with whatever systems utilize those low-level signals.

RFI noise on the AC power conductors may be mitigated by routing the AC power through *filter* circuits placed near the drive. The filter circuits block high-frequency noise from propagating back to the rest of the AC power distribution wiring where it may influence other electronic equipment. However, there is little that may be done about the RFI noise between the drive and the motor other than to shield the conductors in well-grounded metallic conduit.

28.3 AC motor braking

There are several different methods useful for causing an AC induction motor to *brake*, or slow down:

- DC injection
- Dynamic braking
- Regenerative braking
- Plugging

DC injection uses the technique of energizing the stator windings with low-current DC instead of high-current AC as is the case when the motor runs. *Dynamic braking* works the motor as a generator, dissipating energy through a resistive load. *Regenerative braking* also works the motor as a generator, but instead of wasting energy in the form of resistive heating, a regenerating motor drive pumps that energy back into the power supply grid where it may be used by other loads. Lastly, *plugging* works by applying reverse power to the motor, and is the most aggressive means of bringing any motor to a halt.

All electronic motor braking techniques enjoy the advantage of mechanical simplicity. If the motor itself can be used as a brake, then a separate mechanical brake may not be needed. This simplifies the machinery of a system and potentially reduces maintenance costs.

A significant disadvantage of electronic braking techniques is that they all depend on the proper function of the motor drive, and in some cases the AC line power as well. If a VFD's braking ability depends on the presence of AC line power, and that line power suddenly is lost, the VFD will have no braking capacity at all! This means a large motor might suddenly have no ability to brake in the event of a power outage or a tripped circuit breaker, which could be a serious safety issue in some applications. In such cases, one must ensure the presence of other (alternative) braking methods to function in the event of line power failure.

28.3.1 DC injection braking


If a spinning AC induction motor's stator coils are energized with DC rather than AC, the rotor will find itself spinning inside a stationary magnetic field. This causes currents to be induced in the rotor bars, which in turn causes a braking force to develop in the rotor in accordance with Lenz's Law. The effect is exactly opposite of what happens when a motor is energized from a stand-still: there, currents are induced in the rotor bars because the rotor is stationary and the stator field is rotating. This method of braking is quite effective, with only small amounts of direct current through the stator winding being necessary to cause a large braking torque.

The braking torque produced by DC injection varies directly with the magnitude of the DC injection current, and also directly with the speed of the rotor. This means the braking force created by DC injection tends to diminish as the motor slows down to a stop.

When any motor acts as a brake, the kinetic energy of the motor and the mechanism it attaches to must go somewhere. This is a basic tenet of physics, codified as the *Law of Energy Conservation*: energy cannot be created or destroyed, only altered in form. When DC injection is used to brake a motor, the braking energy is dissipated in the form of heat by means of the induced currents circulating through the rotor bars and shorting rings. This is something one must be careful to consider when choosing DC injection as a braking method: can the rotor safely dissipate the heat when needed? Repeated braking cycles, especially with little time between cycles, may overheat the rotor and cause damage to the motor.

Modern solid-state AC motor drives easily provide DC injection for braking. All they need to do is energize their output transistors in such a way that one or more of the stator windings sees a constant voltage polarity instead of an alternating polarity as is the case when the motor is running. The following diagram shows the power flow into the motor during DC injection:

Example showing a VFD injecting DC to the motor



The intensity of the DC injection current may be varied by altering the pulse-width duty cycle of the transistors used to switch the braking current.

28.3.2 Dynamic braking


If a powered AC induction motor spins at a speed *faster* than its rotating magnetic field, it acts as a generator: supplying power back to the voltage source, transferring kinetic energy from the spinning rotor and machinery back into electrical power. This makes for an interesting experiment: take an internal combustion engine, steam turbine, water turbine, or some other mechanical prime mover and mechanically *force* a powered induction motor to spin faster than its synchronous speed (i.e. force it to achieve a *negative* slip speed). If a power meter is connected between this motor and the AC line power grid, the meter will register negative power (i.e. power flowing from the motor to the grid, rather than from the grid to the motor).

This principle holds true for an induction motor powered by a VFD as well: if the rotor is spun faster than the speed of the rotating magnetic field produced by the VFD, it will act as a generator, sending back more power to the VFD than it receives from the VFD. Since the magnetic field's rotational speed is variable – thanks to the VFD's ability to synthesize virtually any desired frequency – it means an induction motor may be made to operate as a generator at almost any speed we desire.

When acting as an electrical generator, an induction motor requires an input of mechanical energy. That is, it will require mechanical *effort* to keep the rotor spinning faster than synchronous speed, since the motor naturally “wants” to spin at synchronous speed or slower. This means a generating motor acts as a brake, attempting to slow down whatever is keeping it spinning faster than synchronous speed. This braking effect is in direct proportion to how much the generated energy is used or dissipated by an electrical load. If we build a VFD to dissipate this energy in a controlled manner, the motor will have the ability to act as a *dynamic* brake.

In a VFD circuit, the “reverse” power flow received from the motor takes the form of currents traveling through the reverse-protection diodes placed in parallel with the output transistors. This in turn causes the DC bus filter capacitor to charge, resulting in a raised DC bus voltage:


Generating currents through reverse-protection diodes



Without a place for this energy to dissipate, however, there will be little braking effort, and the capacitor will be quickly destroyed by the excessive DC bus voltage. Therefore, in order for dynamic braking to work, the VFD must be equipped with a *braking resistor* to dissipate the received energy. A special transistor rapidly switched on and off to regulate DC bus voltage ensures the capacitor will not be harmed, and that the braking is effective.

This next schematic diagram shows how a braking resistor and its accompanying transistor could be added to the simple VFD circuit. Once again, the switching circuitry used to turn the braking transistor rapidly on and off has been omitted for simplicity:

Braking resistor provides a means of energy dissipation




The braking transistor switches on in direct proportion to the DC bus voltage. The higher the DC bus voltage, the greater the duty cycle (on time versus total time) of the braking transistor. Thus, the transistor functions as a *shunt voltage regulator*, placing a controlled load on the DC bus in direct proportion to its degree of over-voltage. This transistor never turns on when the DC bus voltage is within normal (motoring) operating range. It only turns on to clamp DC bus voltage to reasonable levels when the motor spins faster than synchronous speed.

With this braking circuit in place, the only action a VFD must take to dynamically brake an AC induction motor is simply slow down the applied AC frequency to the motor until that frequency is less than the equivalent rotor speed (i.e. create a condition of negative slip speed).

As with DC injection braking, the braking torque created by dynamic braking is a function of magnetic field strength and rotor speed. More precisely, it is a function of the Volts/Hz ratio applied by the VFD to the motor, and the magnitude of the negative slip speed. Braking torque is primarily limited by the braking resistor's power rating and also the power rating of the VFD. Since the kinetic energy dissipation occurs outside the motor, there is little rotor heating as is the case with DC injection braking.

28.3.3 Regenerative braking

Regenerative braking takes the concept of dynamic braking one step further, in converting the DC bus over-voltage into usable AC power to be placed back on the AC line for other AC devices to use. Rather than regulate DC bus voltage via a shunt resistor switched on and off by a special transistor, a regenerative drive manages the same task by augmenting the bridge rectifier diode array with a set of six more power transistors, then switching those transistors on and off synchronously with the line voltage (the AC power source). This line-synchronized switching takes the DC bus voltage and “inverts” it to AC so that the drive may send real power back into the AC power system from whence it originated:




Rectifier circuits equipped with a set of line-synchronized power transistors are often referred to as an *active front end* to the motor drive. The term “active” refers to the transistors (diodes are “passive” devices), and the term “front end” simply refers to the bridge being at the incoming (front) side of the VFD power circuit. In such a drive, the front end’s transistors are sequenced as needed to clamp the DC bus voltage to reasonable maximum levels, just like the braking transistor is pulsed in a drive with dynamic braking to shunt-regulate DC bus voltage. If DC bus voltage in a regenerating drive rises too high, the active front end transistors will pulse for longer periods of time (i.e. with greater duty cycles) to apply more of that braking energy to the AC power grid.

Regenerative braking enjoys the unique advantage of putting the kinetic energy lost through braking back into productive use. No other method of motor braking does this. The cost of doing this, of course, is increased component count and complexity in the motor drive itself, leading to a more expensive and (potentially) fault-prone VFD. However, in applications where the recovered energy is significant, the cost savings of regenerative braking will rapidly offset the additional capital expense of the regenerative drive.

A simpler and cheaper way to enjoy the benefits of regenerative braking without adding a lot of complexity to the VFD circuitry is to take multiple VFDs and simply connect their DC bus circuits in parallel. If one of the drives slows down its motor, the raised DC bus voltage will be available at the other motor drives to help them drive their motors.

The following schematic diagram shows two interconnected VFD circuits, with the upper drive braking and the lower drive motoring (driving):



The major disadvantage to regeneratively braking in this fashion is that the braking energy is only recoverable by the other motor(s) with their DC busses paralleled, and only at the exact same time one or more of those motors are braking. This is not as convenient or practical as AC line regenerative braking, where a virtually unlimited number of loads exist on the grid to absorb the braking energy at any time. However, for certain applications¹¹ it may be practical, and in those applications the installed cost of the VFDs will be less than a comparable installation with AC line regeneration.

As with dynamic braking, motor heating is reduced (compared to DC injection braking) because the kinetic energy is dissipated elsewhere.

¹¹One such application is machine motion control, where one part of the machine always needs to slow down while another part is accelerating. Another application is coupling the drive motors of two conveyor belts together, where one conveyor always lifts the load uphill and the other conveyor always lowers the load downhill.

28.3.4 Plugging

Plugging is the most powerful method of braking an electric motor, consisting of actively applying power to the motor in the opposite direction of its rotation. This is analogous to reversing the engine thrust of a power boat or an airplane in order to quickly bring it to a halt. For a VFD, this means a reversal of phase rotation while carefully applying power to the AC induction motor.

Like DC injection braking, plugging requires power be applied to the motor in order to make it stop, and it also results in all the kinetic energy being dissipated in the rotor. The advantage held by plugging over DC injection braking is that the braking torque may be maintained and precisely controlled all the way to zero speed.

28.4 Motor drive features

Modern DC and AC motor drives provide features useful when using electric motors as final control elements. Some common features seen in both VSDs and VFDs are listed here:

- Speed limiting
- Torque limiting
- Torque profile curves (used to regulate the amount of torque available at different motor speeds)
- Acceleration (speed rate-of-change) limiting
- Deceleration (speed rate-of-change) limiting
- DC injection braking (applying DC to a motor to turn it into an electromagnetic brake)
- Dynamic braking (turning the motor into an electromagnetic brake¹²)
- Regenerative braking (turning the motor into a generator to recover kinetic energy)
- Plugging (applying reverse-direction power to a motor to *quickly* stop it)
- Overcurrent monitoring and automatic shut-down
- Overvoltage monitoring and automatic shut-down
- PWM frequency adjustment (may be helpful in reducing electromagnetic interference with some equipment)

Not only are some of these limiting parameters useful in extending the life of the motor, but they may also help extend the operating life of the mechanical equipment powered by the motor. It is certainly advantageous, for example, to have torque limiting on a conveyor belt motor, so that the motor does not apply full rated torque (i.e. stretching force) to the belt during start-up.


If a motor drive is equipped with digital network communication capability (e.g. Modbus), it is usually possible for a host system such as a PLC or DCS to update these control parameters as the motor is running.

In order for a VSD or VFD to properly and safely control an electric motor, that drive must be programmed with the motor's nameplate data (voltage rating, current rating, maximum speed, etc.). Failure to properly configure an electronic motor drive with these "base" parameters may even result in damage to the motor, for example if the drive is configured to output more current than the motor is rated for! As such, it is recommended that you *first* program these parameters into a motor drive before setting any other drive parameters.

¹²This is accomplished in very different ways for DC versus AC motors. To dynamically brake a DC motor, the field winding must be kept energized while a high-power load resistor is connected to the armature. As the motor turns, the armature will push current through the resistor, generating a braking torque as it does. One way to dynamically brake an AC motor is to inject a small DC current through the stator windings, causing large braking currents to be induced in the rotor. Another way is to regeneratively brake into a resistive load.

28.5 Use of line reactors

Regulating the electric power sent to an electric motor is a task performed by high-speed switching transistors inside a motor drive, modulating the pulse-width of a high-frequency square wave to the motor. The high-speed switching happening inside of a motor “drive” circuit results in the drive drawing current from the AC power source as high-frequency pulses rather than as sinusoidal waves. These current pulses tend to distort the voltage of the AC power source so that other devices powered by the same AC source will “see” high-frequency noise on the power lines. This is true for DC and AC motor drives alike:



As French mathematician and physicist Jean Baptiste Joseph Fourier (1768-1830) mathematically proved centuries ago, any repeating waveform – no matter how strange the shape may be – is equivalent to a series of sine and cosine waves at integer multiples (“harmonics”) of some fundamental frequency. Thus, the normal sine-wave AC power supplied to an operating motor drive unit will be tainted by harmonic frequencies in addition to the fundamental frequency of 60 Hz¹³.

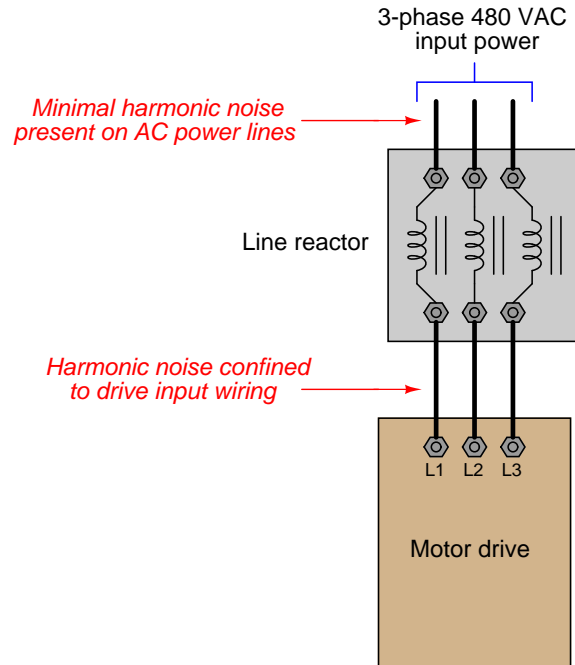
Such high-frequency noise may be very troublesome to nearby electronic devices and to other electrical components connected to the same AC power system. Power transformers will suffer increased core heating from harmonic currents. System capacitances and inductances may *resonate*

¹³In Europe, the fundamental power line frequency is 50 Hz rather than 60 Hz. Also noteworthy is the fact that since the distortion caused by motor drives is typically symmetrical above and below the center-line of the AC waveform, the only significant harmonics will be odd and not even. In a 60 Hz system, the odd harmonics will include 180 Hz (3rd), 300 Hz (5th), 420 Hz (7th), and higher. For a 50 Hz system, the corresponding harmonic frequencies are 150 Hz, 250 Hz, 350 Hz, etc.

at these harmonic frequencies causing high currents and voltages to form. So-called *triplen harmonics*¹⁴ are especially troublesome in three-phase power circuits, where they tend to add in the neutral conductors of Wye-connected system components and circulate through the phase elements in Delta-connected system components. In some industrial installations, the magnitude of triplen harmonic currents in a 4-wire Wye system have been so great that the neutral conductor actually overheated from excessive current, even though the three line conductors were well within their rated load current capacities!

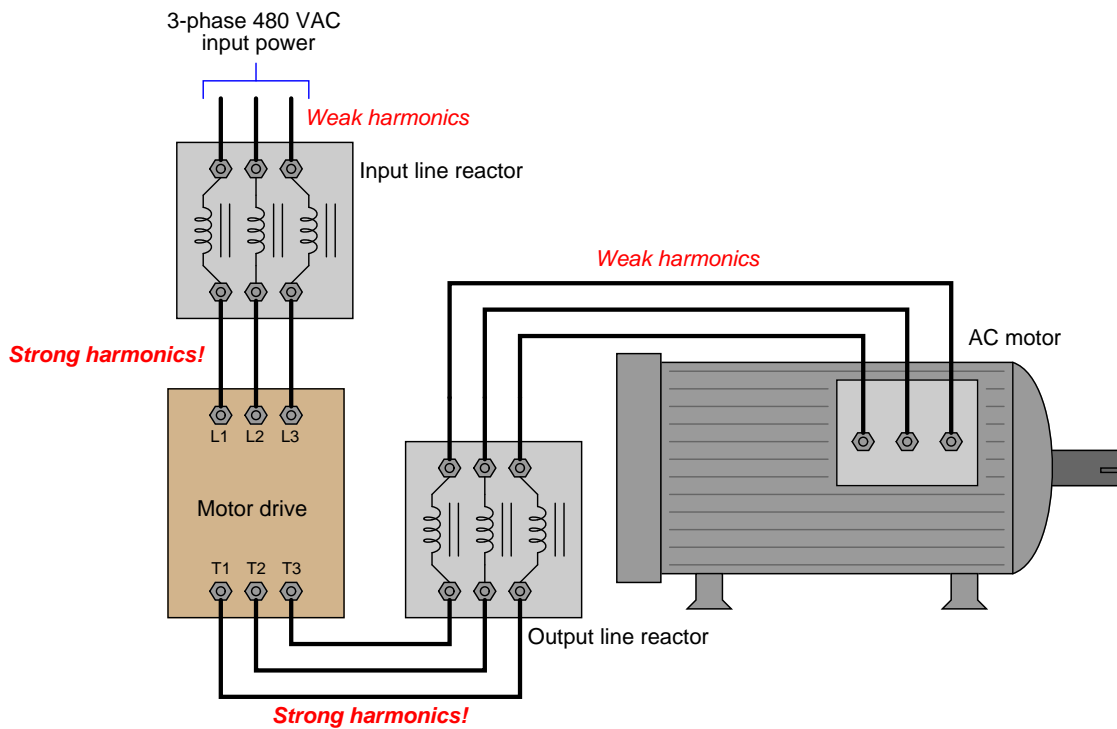
¹⁴Harmonic voltages and currents whose frequencies are multiples of three of the fundamental (e.g. 3rd, 6th, 9th, 12th, 15th harmonics). The reason these particular harmonics are noteworthy in three-phase systems is due to their relative phase shifts. Whereas the fundamental phase shift angle between different phase elements of a three-phase electrical system is 120° , the phase shift between triplen harmonics is zero. Thus, triplen harmonics are *directly additive* in three-phase systems.

One method useful to combat these effects is to filter harmonic frequencies from the rest of the AC power system, preventing the subsequent “corruption” of the AC power source by the motor drive’s pulsing currents. The most direct way to filter harmonic frequencies is to use an electrical component acting as a low-pass filter – a simple *inductor* connected in series with the motor drive. For three-phase-powered motor drives, this takes the form of three inductor elements, commonly referred to in industry as *reactors*:



Line reactors work by presenting a greater series impedance to high-frequency harmonic currents than to low-frequency fundamental currents, following the inductive reactance formula $X_L = 2\pi fL$. The greater the frequency (f) of current, the greater the inductive reactance (X_L) and therefore the greater the attenuation of that current through that conductor. As one might expect, line reactors cannot *prevent* harmonic distortion in the AC power system, but they do a great deal to mitigate the ill effects of harmonics produced by a motor drive.

Line reactors may also be used on the *output* of an AC motor drive to filter harmonics from the motor itself. Like transformers, AC induction motors suffer greater core losses when exposed to harmonic currents, causing the motor to heat up more than it would if powered by AC power of one pure frequency:



The presence of strong harmonic distortion on the motor drive's input wiring means those conductors should be kept short as possible to minimize electromagnetic interference with nearby electrical and electronic components.

Not only do output line reactors help reduce heating effects in the AC motors powered by variable-frequency drives, the reactors also reduce the severity of fault currents resulting from short-circuit transistor failures in the motor drive, as well as minimize the ill effects of reflected signals in the conductors stretching between the output line reactor and the motor itself¹⁵. With such benefits arguing for the installation of line reactors in variable-speed motor control circuits, the only reason for their non-installation is added expense, and/or insufficient space inside the enclosure with the motor drive.

¹⁵As you may recall, any sufficiently long set of conductors will act as a *transmission line* for high-frequency pulse signals. An unterminated (or poorly-terminated) transmission line will *reflect* pulse signals reaching its ends. In the case of a motor drive circuit, these reflected pulses may constructively interfere to produce nodes of high voltage or high current, causing premature wiring failure. Output line reactors help minimize these effects by filtering out high-frequency pulse signals from reaching the long motor power conductors.

28.6 Metering pumps

A very common method for directly controlling low flow rates of fluids is to use a device known as a *metering pump*. A “metering pump” is a pump mechanism, motor, and drive electronics contained in a monolithic package. Simply supply 120 VAC power and a control signal to a metering pump, and it is ready to use.

Metering pumps are commonly used in water treatment processes to inject small quantities of treatment chemicals (e.g. coagulants, disinfectants, acid or caustic liquids for pH neutralization, corrosion-control chemicals) into the water flowstream, as is the Milton-Roy unit shown in this photograph:



In the case of water treatment chemical injection, the flow rates of each chemical must be proportioned to the flow rate of the water being treated. This is why a simple “manual” set flow rate is insufficient for the task. Each chemical injection pump’s flow must be automatically adjustable, so that a control system is able to modulate the injection of each chemical according to the needs of the process, without human operator intervention.

A standard 4-20 mA DC control signal adjusts the output flow of the metering pump, from 0% flow (4 mA) to 100% full flow (20 mA). Adjustment knobs on the front of the pump establish the maximum flow rate at a control signal value of 100% (i.e. the controlled flow “range” of the pump):



While some metering pumps use rotary motor and pump mechanisms, many use a “plunger” style mechanism operated by a solenoid at variable intervals. Thus, the latter type of metering pump does not provide continuous flow control, but rather a flow consisting of discrete pulses distributed over a period of time. The “plunger” metering pumps are quite simple and reliable, and are entirely appropriate if non-continuous flow is permissible for the process. The Milton-Roy pump shown in these photographs is of that design: a plunger injects pulses of liquid into the process line, the frequency of that plunger’s action determined by the 4-20 mA control signal.

28.7 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Conservation of energy:** energy cannot be created or destroyed, only converted between different forms. Relevant to motor braking techniques in variable-frequency drives (VFDs) – the kinetic energy taken from the rotating machine during braking must go somewhere.
- **Newton's Second Law of motion:** $F = ma$, describing how the acceleration of an object is directly proportional to the amount of applied (resultant) force and inversely proportional to its mass. Relevant to the calculation of motor current required to accelerate a machine part.
- **Lenz's Law:** any magnetic field arising from electromagnetic induction opposes the inducing field. Relevant to the operation of all induction AC motors, as well as to DC injection, dynamic, and regenerative braking of AC motors.
- **Rotating magnetic field:** this is necessary to cause an AC induction motor to spin in a particular direction, and is generated by polyphase field poles (i.e. multiple magnetic fields that are out-of-phase with each other). All AC induction motors require such a polyphase magnetic field to start up in a particular direction, although a single-phase magnetic field is sufficient to maintain rotation once started.

References


Irwin, J. David, *The Industrial Electronics Handbook*, CRC Press, Boca Raton, FL, 1997.

Chapter 29

Closed-loop control

Instrumentation is the science of automated measurement and control. Applications of this science abound in modern research, industry, and everyday living. From automobile engine control systems to home thermostats to aircraft autopilots to the manufacture of pharmaceutical drugs, automation surrounds us. This chapter explains some of the fundamental principles of automatic process control.


Transmitters are represented in process diagrams by small circles with identifying letters inside, in this case, “TT,” which stands for **T**emperature **T**ransmitter:



The signal output by the transmitter (represented by the “PV” dashed line), representing the heated fluid’s exiting temperature, is called the *process variable*. Like a variable in a mathematical equation that represents some story-problem quantity, this signal represents the measured quantity we wish to control in the process.


In order to exert control over the process variable, we must have some way of altering fluid flow through the heat exchanger, either of the process fluid, the steam, or both. Generally, it makes more sense to alter the flow of the heating medium (the steam), and let the process fluid flow rate be dictated by the demands of the larger process. If this heat exchanger were part of an oil refinery unit, for example, it would be far better to throttle steam flow to control oil temperature rather than to throttle the oil flow itself, since altering the oil’s flow will also affect other process variables upstream and downstream of the exchanger. Ideally, the heat exchanger temperature control system would provide consistent temperature of the exiting oil, for any given incoming oil temperature and flow-rate of oil through it.

One convenient way to throttle steam flow into the heat exchanger is to use a control valve (labeled “TV” because it is a **T**emperature **V**alve). In general terms, a control valve is known as a *final control element*. Other types of final control elements exist (servo motors, variable-flow pumps, and other mechanical devices used to vary some physical quantity at will), but valves are the most common, and probably the simplest to understand. With a final control element in place, the steam flow becomes known as the *manipulated variable*, because it is the quantity we will manipulate in order to gain control over the process variable:



Valves come in a wide variety of sizes and styles. Some valves are hand-operated: that is, they have a “wheel” or other form of manual control that may be moved to “pinch off” or “open up” the flow passage through the pipe. Other valves come equipped with signal receivers and positioner devices, which move the valve mechanism to various positions at the command of a signal (usually an electrical signal, like the type output by transmitter instruments). This feature allows for remote control, so a human operator or computer device may exert control over the manipulated variable from a distance. In the previous illustration, the steam control valve is equipped with such an electrical signal input, represented by the “control signal” dashed line.

This brings us to the final component of the heat exchanger temperature control system: the *controller*. This is a device designed to interpret the transmitter's process variable signal and decide how far open the control valve needs to be in order to maintain that process variable at the desired value.




Here, the circle with the letters “TC” in the center represents the controller. Those letters stand for **T**emperature **C**ontroller, since the process variable being controlled is the process fluid’s *temperature*. Usually, the controller consists of a computer making automatic decisions to open and close the valve as necessary to stabilize the process variable at some predetermined *setpoint*.

Note that the controller’s circle has a solid line going through the center of it, while the transmitter and control valve circles are open. An open circle represents a field-mounted device according to the ISA standard for instrumentation symbols, and a single solid line through the middle of a circle tells us the device is located on the front of a control panel in a main control room location. So, even though the diagram might appear as though these three instruments are located close to one another, they in fact may be quite far apart. Both the transmitter and the valve must be located near the heat exchanger (out in the “field” area rather than inside a building), but the controller may be located a long distance away where human operators can adjust the setpoint from inside a safe and secure control room.

These elements comprise the essentials of a *feedback control system*: the *process* (the system

to be controlled), the *process variable* (the specific quantity to be measured and controlled), the *transmitter* (the device used to measure the process variable and output a corresponding signal), the *controller* (the device that decides what to do to bring the process variable as close to setpoint as possible), the *final control element* (the device that directly exerts control over the process), and the *manipulated variable* (the quantity to be directly altered to effect control over the process variable).

Feedback control may be viewed as a sort of information “loop,” from the transmitter (measuring the process variable), to the controller, to the final control element, and through the process itself, back to the transmitter. Ideally, a process control “loop” not only holds the process variable at a steady level (the setpoint), but also maintains control over the process variable given changes in setpoint, and even changes in other variables of the process:



Specifically, the type of feedback we are employing here to control the process is *negative* or *degenerative* feedback. The term “negative” refers to the direction of action the control system takes in response to any measured change in the process variable. If something happens to drive the process variable up, the control system will automatically respond in such a way as to bring the process variable back down where it belongs. If the process variable happens to sag below setpoint, the control system will automatically act to drive the process variable back up to setpoint. Whatever the process variable does in relation to setpoint, the control system takes the opposite (inverse, or negative) action in an attempt to stabilize it at setpoint.

For example, if the unheated process fluid flow rate were to suddenly increase, the heat exchanger outlet temperature would fall due to the physics of heat transfer, but once this drop was detected by the transmitter and reported to the controller, the controller would automatically call for additional steam flow to compensate for the temperature drop, thus bringing the process variable back in agreement with the setpoint. Ideally, a well-designed and well-tuned control loop will sense and


compensate for *any* change in the process or in the setpoint, the end result being a process variable value that always holds steady at the setpoint value.

The unheated fluid flow rate is an example of an uncontrolled, or *wild*, variable because our control system here has no ability to influence it. This flow is also referred to as a *load* because it “loads” or affects the process variable we are trying to stabilize. Loads are present in nearly every controlled system, and indeed are the primary factor necessitating a control system at all. Referring back to our heat exchanger process again, we could adequately control the operating temperature of it with just a manually-set steam control valve if only none of the other factors (steam temperature, fluid flow rate, incoming fluid temperature, etc.) ever changed!

Many types of processes lend themselves to feedback control. Consider an aircraft autopilot system, keeping an airplane on a steady course heading despite the effects of loads such as side-winds: reading the plane’s heading (process variable) from an electronic compass and using the rudder as a final control element to change the plane’s “yaw.” An automobile’s “cruise control” is another example of a feedback control system, with the process variable being the car’s velocity, and the final control element being the engine’s throttle. The purpose of a cruise control is to maintain constant driving speed despite the influence of loads such as hills, head-winds, tail-winds, and road roughness. Steam boilers with automatic pressure controls, electrical generators with automatic voltage and frequency controls, and water pumping systems with automatic flow controls are further examples of how feedback may be used to maintain control over certain process variables.

Modern technology makes it possible to control nearly anything that may be measured in an industrial process. This extends beyond the pale of simple pressure, level, temperature, and flow variables to include even certain chemical properties.

In municipal water and wastewater treatment systems, for example, numerous chemical quantities must be measured and controlled automatically to ensure maximum health and minimum environmental impact. Take for instance the chlorination of treated wastewater, before it leaves the wastewater treatment facility into a large body of water such as a river, bay, or ocean. Chlorine is added to the water to kill any residual bacteria so they do not consume oxygen in the body of water they are released to. Too little chlorine added, and not enough bacteria are killed, resulting in a high *biological oxygen demand* or *BOD* in the water which will asphyxiate the fish swimming in it. Too much chlorine added, and the chlorine itself poses a hazard to marine life. Thus, the chlorine content must be carefully controlled at a particular setpoint, and the control system must take aggressive action if the dissolved chlorine concentration strays too low or too high:




Now that we have seen the basic elements of a feedback control system, we will concentrate on the *algorithms* used in the controller to maintain a process variable at setpoint. For the scope of this topic, an “algorithm” is a mathematical relationship between the process variable and setpoint inputs of a controller, and the output (manipulated variable). Control algorithms determine *how* the manipulated variable quantity is deduced from PV and SP inputs, and range from the elementary to the very complex. In the most common form of control algorithm, the so-called “PID” algorithm, calculus is used to determine the proper final control element action for any combination of input signals.

29.2 Diagnosing feedback control problems

Negative feedback systems, in general, tend to cause much confusion for those first learning their fundamental principles and behaviors. The closed-cycle “loop” formed by the interaction of sensing element, controller, final control element, and process means essentially that *everything affects everything else*. This is especially problematic when the feedback control system in question contains a fault and must be diagnosed. For example, if an operator happens to notice that the process variable (as indicated by a manual measurement or by some trusted indicating instrument) is not holding to setpoint, it could be the result of a fault in *any* portion of the system (sensor, controller, FCE, or even the process itself).

Recall that every feedback control loop consists of four basic elements: an element that *senses* the process variable (e.g. primary sensing element, transmitter), an element that *decides* what how to regulate this process variable (e.g. a PID controller), an element that *influences* the process variable (e.g. a control valve, motor drive, or some other final control device), and finally the process itself which *reacts* to the final control device’s actions:



One of the basic diagnostic strategies for any instrumentation system is to assess whether the *input value(s)* and *output value(s)* correspond for each instrument. We may apply this same strategy to each of the four elements of a feedback control “loop” to identify where the problem might exist. If you encounter one of these four system portions whose output does not correspond with its input, you know that portion of the system is faulted.

You can check each element of your feedback control loop by comparing its input with its output to see if each element is doing what it should. I recommend beginning with the controller (the decision-making element) because typically those values are the most easily monitored:


- **Decision-making:** Carefully examine the controller faceplate, looking at the values of PV, SP, and Output. Is the controller taking appropriate action to force PV equal to SP? In other words, is the Output signal at a value you would expect if the controller were functioning properly to regulate the process variable at setpoint? If so, then the controller's action and tuning are most likely not at fault. If not, then the problem definitely lies with the controller.
- **Sensing:** Compare the controller's displayed value for PV with the actual process variable value as indicated by local gauges, by feel, or by any other means of detection. If there is good correspondence between the controller's PV display and the real process variable, then there probably isn't anything wrong with the measurement portion of the control loop (e.g. transmitter, impulse lines, PV signal wiring, analog input of controller, etc.). If the displayed PV disagrees with the actual process variable value, then something is definitely wrong here.
- **Influencing:** Compare the controller's displayed value for Output with the actual status of the final control element. If there is good correspondence between the controller's Output display and the FCE's status, then there probably isn't anything wrong with the output portion of the control loop (e.g. FCE, output signal wiring, analog output of controller, etc.). If the controller Output value differs from the FCE's state, then something is definitely wrong here.
- **Reacting:** Compare the process variable value with the final control element's state. Is the process doing what you would expect it to? If so, the problem is most likely not within the process (e.g. manual valves, relief valves, pumps, compressors, motors, and other process equipment). If, however, the process is not reacting the way you would expect it to given the final control element's state, then something is definitely awry with the process itself.

29.3 On/off control

Once while working as an instrument technician in an aluminum foundry, a mechanic asked me what it was that I did. I began to explain my job, which was essentially to calibrate, maintain, troubleshoot, document, and modify (as needed) all automatic control systems in the facility. The mechanic seemed puzzled as I explained the task of “tuning” loop controllers, especially those controllers used to maintain the temperature of large, gas-fired industrial furnaces holding many tons of molten metal. “Why does a controller have to be ‘tuned’?” he asked. “All a controller does is turn the burner on when the metal’s too cold, and turn it off when it becomes too hot!”

In its most basic form, the mechanic’s assessment of the control system was correct: to turn the burner on when the process variable (molten metal temperature) drops below setpoint, and turn it off when it rises above setpoint. However, the actual algorithm is much more complex than that, finely adjusting the burner intensity according to the amount of *error* between PV and SP, the amount of time the error has accumulated, and the rate-of-change of the error over time. In his casual observation of the furnace controllers, though, he had noticed nothing more than the full-on/full-off action of the controller.

The technical term for a control algorithm that merely checks for the process variable exceeding or falling below setpoint is *on/off control*. In colloquial terms, it is known as *bang-bang* control, since the manipulated variable output of the controller rapidly switches between fully “on” and fully “off” with no intermediate state. Control systems this crude usually provide very imprecise control of the process variable. Consider our example of the shell-and-tube heat exchanger, if we were to implement simple on/off control¹:



As you can see, the degree of control is rather poor. The process variable “cycles” between the upper and lower setpoints (USP and LSP) without ever stabilizing at the setpoint, because that

¹To be precise, this form of on/off control is known as *differential gap* because there are two setpoints with a gap in between. While on/off control is possible with a single setpoint (FCE on when below setpoint and off when above), it is usually not practical due to the frequent cycling of the final control element.


would require the steam valve to be position somewhere *between* fully closed and fully open.

This simple control algorithm may be adequate for temperature control in a house, but not for a sensitive chemical process! Can you imagine what it would be like if an automobile's cruise control system relied on this algorithm? Not only is the lack of precision a problem, but the frequent cycling of the final control element may contribute to premature failure due to mechanical wear. In the heat exchanger scenario, thermal cycling (hot-cold-hot-cold) will cause metal fatigue in the tubes, resulting in a shortened service life. Furthermore, every excursion of the process variable above setpoint is wasted energy, because the process fluid is being heated to a greater temperature than what is necessary.

Clearly, the only practical answer to this dilemma is a control algorithm able to *proportion* the final control element rather than just operate it at zero or full effect (the control valve fully closed or fully open). This, in its simplest form, is called *proportional control*.

29.4 Proportional-only control

Imagine a liquid-level control system for a vessel, where the position of a level-sensing float directly sets the stem position of a control valve. As the liquid level rises, the valve opens up proportionally:



Despite its crude mechanical nature, this *proportional* control system would in fact help regulate the level of liquid inside the process vessel. If an operator wished to change the “setpoint” value of this level control system, he or she would have to adjust the coupling between the float and valve stems for more or less distance between the two. Increasing this distance (lengthening the connection) would effectively raise the level setpoint, while decreasing this distance (shortening the connection) would lower the setpoint.

We may generalize the proportional action of this mechanism to describe *any* form of controller where the output is a direct function of process variable (PV) and setpoint (SP):

$$m = K_p e + b$$

Where,

m = Controller output

e = Error (difference between PV and SP)

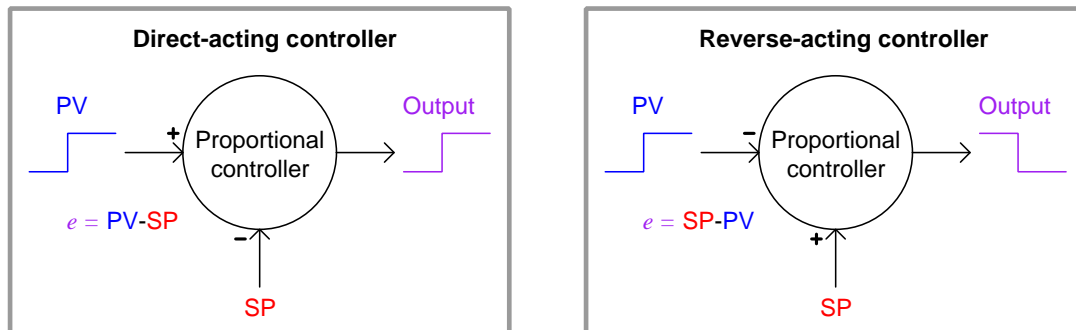
K_p = Proportional gain

b = Bias

A new term introduced with this formula is e , the “error” or difference between process variable and setpoint. Error may be calculated as $SP - PV$ or as $PV - SP$, depending on whether or not the controller must produce an *increasing* output signal in response to an increase in the process variable (“direct” acting), or output a *decreasing* signal in response to an increase in the process variable (“reverse” acting):


$$m = K_p(PV - SP) + b \quad (\text{Direct-acting proportional controller})$$

$$m = K_p(SP - PV) + b \quad (\text{Reverse-acting proportional controller})$$



The optional “+” and “-” symbols clarify the effect each input has on the controller output: a “-” symbol representing an *inverting* effect and a “+” symbol representing a *noninverting* effect. When we say that a controller is “direct-acting” or “reverse-acting” we are referring to its reaction to the PV signal, therefore the output signal from a “direct-acting” controller goes in the same direction as the PV signal and the output from a “reverse-acting” controller goes in the opposite direction of its PV signal. It is important to note, however, that the response to a change in setpoint (SP) will yield the *opposite* response as does a change in process variable (PV): a rising SP will drive the output of a direct-acting controller *down* while a rising SP drives the output of a reverse-acting controller *up*. “+” and “-” symbols explicitly show the effect both inputs have on the controller output, helping to avoid confusion when analyzing the effects of PV changes versus the effects of SP changes.

The direction of action required of the controller is determined by the nature of the process, transmitter, and final control element. In the case of the crude mechanical level controller, the action needs to be *direct* so that a greater liquid level will result in a further-open control valve to drain the vessel faster. In the case of the automated heat exchanger shown earlier, we are assuming that an increasing output signal sent to the control valve results in increased steam flow, and consequently higher temperature, so our controller will need to be reverse-acting (i.e. an increase in measured temperature results in a decrease in output signal; error calculated as $SP - PV$):



After the error has been calculated, the controller then multiplies the error signal by a constant value called the *gain*, which is programmed into the controller. The resulting figure, plus a “bias” quantity, becomes the output signal sent to the valve to proportion it. The “gain” value is exactly what it seems to be for anyone familiar with electronic amplifier circuits: a ratio of output to input. In this case, the gain of a proportional controller is the ratio of output signal change to input signal change, or how *aggressive* the controller reacts to changes in input (PV or SP).

To give a numerical example, a loop controller set to have a gain of 4 will change its output signal by 40% if it sees an input change of 10%: the ratio of output change to input change will be 4:1. Whether the input change comes in the form of a setpoint adjustment, a drift in the process variable, or some combination of the two does not matter to the magnitude of the output change.


The bias value of a proportional controller is simply the value of its output whenever process

variable happens to be equal to setpoint (i.e. a condition of zero *error*). Without a bias term in the proportional control formula, the valve would always return to a fully shut (0%) condition if ever the process variable reached the setpoint value. The bias term allows the final control element to achieve a non-zero state at setpoint.


If the controller could be configured for infinite gain, its response would duplicate on/off control. That is, *any* amount of error will result in the output signal becoming “saturated” at either 0% or 100%, and the final control element will simply turn on fully when the process variable drops below setpoint and turn off fully when the process variable rises above setpoint. Conversely, if the controller is set for zero gain, it will become completely unresponsive to changes in either process variable *or* setpoint: the valve will hold its position at the bias point no matter what happens to the process.

Obviously, then, we must set the gain somewhere between infinity and zero in order for this algorithm to function any better than on/off control. Just how much gain a controller needs to have depends on the process and all the other instruments in the control loop.

If the gain is set too high, there will be oscillations as the PV converges on a new setpoint value:




If the gain is set too low, the process response will be stable under steady-state conditions but relatively slow to respond to changes in setpoint, as shown in the following trend recording:



A characteristic deficiency of proportional control action, exacerbated with low controller gain values, is a phenomenon known as *proportional-only offset* where the PV never fully reaches SP. A full explanation of proportional-only offset is too lengthy for this discussion and will be presented in a subsequent section of the book, but may be summarized here simply by drawing attention to the proportional controller equation which tells us the output always returns to the bias value when PV reaches SP (i.e. $m = b$ when $PV = SP$). If anything changes in the process to require a different output value than the bias (b) to stabilize the PV, an error between PV and SP *must* develop to drive the controller output to that necessary output value. This means it is only by chance that the PV will settle precisely at the SP value – most of the time, the PV will deviate from SP in order to generate an output value sufficient to stabilize the PV and prevent it from drifting. This persistent error, or offset, worsens as the controller gain is reduced. Increasing controller gain causes this offset to decrease, but at the expense of oscillations.

With proportional-only control, the choice of gain values is really a compromise between excessive oscillations and excessive offset. A well-tuned proportional controller response is shown here:



An unnecessarily confusing aspect of proportional control is the existence of two completely different ways to express controller proportionality. In the proportional-only equation shown earlier, the degree of proportional action was specified by the constant K_p , called *gain*. However, there is another way to express the sensitivity of proportional action, and that is to state the percentage of error change necessary to make the output (m) change by 100%. Mathematically, this is the inverse of gain, and it is called *proportional band* (PB):

$$K_p = \frac{1}{\text{PB}} \quad \text{PB} = \frac{1}{K_p}$$

Gain is always specified as a unitless value², whereas proportional band is always specified as a percentage. For example, a gain value of 2.5 is equivalent to a proportional band value of 40%, because the error input to this controller must change by 40% in order to make the output change a full 100%.

²In electronics, the unit of *decibels* is commonly used to express gains. Thankfully, the world of process control was spared the introduction of decibels as a unit of measurement for controller gain. The last thing we need is a *third* way to express the degree of proportional action in a controller!

Due to the existence of these two completely opposite conventions for specifying proportional action, you may see the proportional term of the control equation written differently depending on whether the author assumes the use of gain or the use of proportional band:

$$K_p = \text{gain} \quad \text{PB} = \text{proportional band}$$



$$K_p e \quad \frac{1}{\text{PB}} e$$

Many modern digital electronic controllers allow the user to conveniently select the unit they wish to use for proportional action. However, even with this ability, anyone tasked with adjusting a controller's "tuning" values may be required to translate between gain and proportional band, especially if certain values are documented in a way that does not match the unit configured for the controller.

When you communicate the proportional action setting of a process controller, you should always be careful to specify either "gain" or "proportional band" to avoid ambiguity. *Never* simply say something like, "The proportional setting is twenty," for this could mean either:

- Proportional band = 20%; Gain = 5 . . . or . . .
- Gain = 20; Proportional band = 5%


As you can see here, the real-life difference in controller response to an input disturbance (wave) depending on whether it has a proportional band of 20% or a gain of 20 is quite dramatic:



29.5 Proportional-only offset

A fundamental limitation of proportional control has to do with its response to changes in setpoint and changes in process *load*. A “load” in a controlled process is any variable not controlled by the loop controller which nevertheless affects the process variable the controller is trying to regulate. In other words, a “load” is any factor the loop controller must compensate for while maintaining the process variable at setpoint.

In our hypothetical heat exchanger system, the temperature of the incoming process fluid is an example of a load:



If the incoming fluid temperature were to suddenly decrease, the immediate effect this would have on the process would be to decrease the outlet temperature (which is the temperature we are trying to maintain at a steady value). It should make intuitive sense that a colder incoming fluid will require more heat input to raise it to the same outlet temperature as before. If the heat input remains the same (at least in the immediate future), this colder incoming flow must make the outlet flow colder than it was before. Thus, incoming feed temperature has an impact on the outlet temperature whether we like it or not, and the control system must compensate for these unforeseen and uncontrolled changes. This is precisely the definition of a “load”: a burden³ on the control

³One could argue that the presence of loads actually *justifies* a control system, for if there were no loads, there

system.

Of course, it is the job of the controller to counteract any tendency for the outlet temperature to stray from setpoint, but as we shall soon see this cannot be perfectly achieved with proportional control alone.

Let us perform a “thought experiment” to demonstrate this phenomenon of proportional-only offset. Imagine the controller has been controlling outlet temperature exactly at setpoint ($PV = SP$), and then suddenly the inlet feed temperature drops and remains colder than before. Recall that the equation for a reverse-acting proportional controller is as follows:

$$m = K_p(SP - PV) + b$$

Where,

m = Controller output
 K_p = Proportional gain
 SP = Setpoint
 PV = Process variable
 b = Bias

The introduction of colder feed fluid to the heat exchanger makes the outlet temperature (PV) begin to fall. As the PV falls, the controller calculates a positive error ($SP - PV$). This positive error, when multiplied by the controller’s gain value, drives the output to a greater value. This opens up the steam valve, adding more heat to the exchanger.

As more heat is added, the rate of temperature drop slows down. The further the PV drops, the more the steam valve opens, until enough additional heat is being added to the heat exchanger to maintain a constant outlet temperature. However, this new stable PV value will be less than it was prior to the introduction of colder feed (i.e. less than the SP). In fact, the controller’s automatic action can *never* return the PV to its original (SP) value so long as the feed remains colder than before. The reason for this is that a greater flow of steam is necessary to balance a colder feed coming in, and the only way a proportional controller is ever going to automatically drive the steam valve to this greater-flow position is if an error develops between PV and SP . Thus, an *offset* inevitably develops between PV and SP due to the load (colder feed).

We may prove the inevitability of this offset another way: imagine somehow that the PV did actually return to the SP value despite the colder feed fluid (remaining colder). If this happened, the steam valve would also return to its former throttling position where it was before the feed temperature dropped. However, we know that this former position will not allow enough steam through to the exchanger to overcome the colder feed – if it did, the PV never would have decreased to begin with! A further-open valve is precisely what we need to stabilize the PV given this colder feed, yet the only way the proportional-only controller can achieve this is if the PV actually falls below SP .

To summarize: the only way a proportional-only controller can automatically generate a new output value (m) is if the PV deviates from SP . Therefore, load changes (requiring new output values to compensate) force the PV to deviate from SP .

would be nothing to compensate for, and therefore no need for an automatic control system at all! In the total absence of loads, a manually-set final control element would be enough to hold most process variables at setpoint.

Another “thought experiment” may be helpful to illustrate the phenomenon of proportional-only offset. Imagine building your own cruise control system for your automobile based on the proportional-only equation: the engine’s throttle position is a function of the difference between PV (road speed) and SP (the desired “target” speed). Let us further suppose that you carefully adjust the bias value of your cruise control system to achieve PV = SP on level ground at a speed of 70 miles per hour (70% on a 0 to 100 MPH speedometer scale), with the throttle at a position of 40%, and a gain (K_p) of 2:

$$m = K_p(SP - PV) + b$$

$$40\% = 2(70 - 70) + 40\%$$

Imagine now that after cruising precisely at setpoint (70% = 70 MPH), the road begins to incline uphill for several miles. This, obviously, is a load on the cruise control system. With the cruise control disengaged, the automobile would slow down because the same throttle position (40%) sufficient to maintain setpoint (70 MPH) on level ground is not enough power to maintain that same setpoint on an incline.

With the cruise control engaged, the engine throttle will automatically open further as speed drops. At a speed of 69 MPH, the throttle opens up to 42%. At a speed of 68 MPH, the throttle opens up to 44%. Every drop in speed of 1 MPH results in a 2% further-open throttle to send more power to the wheels.

Suppose the demands of this particular inclined road require a 50% throttle position for this automobile to maintain a constant speed. In order for your proportional-only cruise control system to deliver this necessary 50% throttle position, the speed will have to “droop” by 5 MPH below setpoint:

$$m = K_p(SP - PV) + b$$


$$50\% = 2(70 - 65) + 40\%$$

There is simply no other way for your proportional-only controller to automatically achieve the requisite 50% throttle position aside from letting the speed sag below setpoint by 5% (5 MPH). Given this fact, the only way the proportional-only cruise control will ever return the speed to setpoint (70 MPH) is if and when the load conditions change to allow for a lesser throttle position of 40%. So long as the load demands a different throttle position than the bias value, the speed *must* deviate from the setpoint value of 70 MPH.

This necessary error developing between PV and SP is called *proportional-only offset*, sometimes called *droop*. The amount of droop depends on how severe the load change is, and how aggressive the controller responds (i.e. how much gain it has). The term “droop” is very misleading, as it is possible for the error to develop the other way (i.e. the PV might rise above SP due to a load change!). Imagine the opposite load-change scenario in our steam heat exchanger process, where the incoming feed temperature suddenly *rises* instead of falls. If the controller was controlling exactly at setpoint before this upset, the final result will be an outlet temperature that settles at some point *above* setpoint, enough so the controller is able to pinch the steam valve far enough closed to stop any further rise in temperature.

Proportional-only offset also occurs as a result of setpoint changes. We could easily imagine the same sort of effect following an operator's increase of setpoint for the temperature controller on the heat exchanger. After increasing the setpoint, the controller immediately increases the output signal, sending more steam to the heat exchanger. As temperature rises, though, the proportional algorithm causes the output signal to decrease. When the rate of heat energy input by the steam equals the rate of heat energy carried away from the heat exchanger by the heated fluid (a condition of *energy balance*), the temperature stops rising. This new equilibrium temperature will not be at setpoint, assuming the temperature was holding at setpoint prior to the human operator's setpoint increase. The new equilibrium temperature indeed *cannot* ever achieve any setpoint value higher than the one it did in the past, for if the error ever returned to zero ($PV = SP$), the steam valve would return to its old position, which we know would be insufficient to raise the temperature of the heated fluid to a new value.

An example of proportional-only control in the context of electronic power supply circuits is the following opamp voltage regulator, used to stabilize voltage to a load with power supplied by an unregulated voltage source:



In this circuit, a zener diode establishes a “reference” voltage (which may be thought of as a “setpoint” for the controlling opamp to follow). The operational amplifier acts as the proportional-only controller, sensing voltage at the load (PV), and sending a driving output voltage to the base of the power transistor to keep load voltage constant despite changes in the supply voltage or changes in load current (both “loads” in the process-control sense of the word, since they tend to influence voltage at the load circuit without being under the control of the opamp).

If everything functions properly in this voltage regulator circuit, the load's voltage will be stable over a wide range of supply voltages and load currents. However, the load voltage cannot ever *precisely* equal the reference voltage established by the zener diode, even if the operational amplifier (the “controller”) is without defect. The reason for this incapacity to perfectly maintain “setpoint” is the simple fact that in order for the opamp to generate any output signal at all, there *absolutely must be* a differential voltage between the two input terminals for the amplifier to amplify. Operational amplifiers (ideally) generate an output voltage equal to the enormously high gain value (A_V) multiplied by the difference in input voltages (in this case, $V_{ref} - V_{load}$). If V_{load} (the “process

variable”) were to ever achieve equality with V_{ref} (the “setpoint”), the operational amplifier would experience absolutely no differential input voltage to amplify, and its output signal driving the power transistor would fall to zero. Therefore, there must always exist some *offset* between V_{load} and V_{ref} (between process variable and setpoint) in order to give the amplifier some input voltage to amplify.

The amount of offset is ridiculously small in such a circuit, owing to the enormous gain of the operational amplifier. If we take the opamp’s transfer function to be $V_{out} = A_V(V_{(+)} - V_{(-)})$, then we may set up an equation predicting the load voltage as a function of reference voltage (assuming a constant 0.7 volt drop between the base and emitter terminals of the transistor):

$$V_{out} = A_V(V_{(+)} - V_{(-)})$$

$$V_{out} = A_V(V_{ref} - V_{load})$$

$$V_{load} + 0.7 = A_V(V_{ref} - V_{load})$$

$$V_{load} + 0.7 = A_V V_{ref} - A_V V_{load}$$

$$V_{load} + A_V V_{load} = A_V V_{ref} - 0.7$$

$$(A_V + 1)V_{load} = A_V V_{ref} - 0.7$$

$$V_{load} = \frac{A_V V_{ref} - 0.7}{A_V + 1}$$

If, for example, our zener diode produced a reference voltage of 5.00000 volts and the operational amplifier had an open-loop voltage gain of 250000, the load voltage would settle at a theoretical value of 4.9999772 volts: just barely below the reference voltage value. If the opamp’s open-loop voltage gain were much less – say only 100 – the load voltage would only be 4.94356 volts. This still is quite close to the reference voltage, but definitely not as close as it would be with a greater opamp gain!

Clearly, then, we can minimize proportional-only offset by increasing the gain of the process controller gain (i.e. decreasing its proportional band). This makes the controller more “aggressive” so it will move the control valve further for any given change in PV or SP. Thus, not as much error needs to develop between PV and SP to move the valve to any new position it needs to go. However, too much controller gain makes the control system unstable: at best it will exhibit residual oscillations after setpoint and load changes, and at worst it will oscillate out of control altogether. Extremely high gains work well to minimize offset in operational amplifier circuits, only because time delays are negligible between output and input. In applications where large physical processes are being controlled (e.g. furnace temperatures, tank levels, gas pressures, etc.) rather than voltages across small electronic loads, such high controller gains would be met with debilitating oscillations.


If we are limited in how much gain we can program in to the controller, how do we minimize this offset? One way is for a human operator to periodically place the controller in manual mode and move the control valve just a little bit more so the PV once again reaches SP, then place the controller back

into automatic mode. In essence this technique adjusts the “Bias” term of the controller equation. The disadvantage of this technique is rather obvious: it requires human intervention. What is the point of having an automation system requiring periodic human intervention to maintain setpoint?

A more sophisticated method for eliminating proportional-only offset is to add a different control action to the controller: one that takes action based on the amount of error between PV and SP and the amount of time that error has existed. We call this control mode *integral*, or *reset*.

29.6 Integral (reset) control

Imagine a liquid-level control system for a vessel, where the position of a level-sensing float sets the position of a potentiometer, which then sets the *speed* of a motor-actuated control valve. If the liquid level is above setpoint, the valve continually opens up; if below setpoint, the valve continually closes off:



Unlike the *proportional* control system where valve position was a direct function of float position, this control system sets the *speed* of the motor-driven valve according to the float position. The further away from setpoint the liquid level is, the *faster* the valve moves open or closed. In fact, the only time the valve will ever halt its motion is when the liquid level is precisely at setpoint; otherwise, the control valve will be in constant motion.

This control system does its job in a very different manner than the all-mechanical float-based proportional control system illustrated previously. Both systems are capable of regulating liquid level inside the vessel, but they take very different approaches to doing so. One of the most significant differences in control behavior is how the proportional system would inevitably suffer from *offset* (a persistent error between PV and SP), whereas this control system actively works at all times to eliminate offset. The motor-driven control valve literally does not rest until all error has been eliminated!

Instead of characterizing this control system as *proportional*, we call it *integral*⁴ in honor of the calculus principle (“integration”) whereby small quantities are accumulated over some span to form a total. Don’t let the word “calculus” scare you! You are probably already familiar with the concept of numerical integration even though you may have never heard of the term before.

Calculus is a form of mathematics dealing with *changing* variables, and how rates of change relate between different variables. When we “integrate” a variable with respect to time, what we are doing is *accumulating* that variable’s value as time progresses. Perhaps the simplest example of this is a vehicle odometer, accumulating the total distance traveled by the vehicle over a certain time period. This stands in contrast to a speedometer, indicating the rate of distance traveled *per* unit of time.

Imagine a car moving along at exactly 30 miles per hour. How far will this vehicle travel after 1 hour of driving this speed? Obviously, it will travel 30 miles. Now, how far will this vehicle travel if it continues for another 2 hours at the exact same speed? Obviously, it will travel 60 more miles, for a total distance of 90 miles since it began moving. If the car’s speed is a constant, calculating total distance traveled is a simple matter of multiplying that speed by the travel time.

The odometer mechanism that keeps track of the mileage traveled by the car may be thought of as *integrating* the speed of the car with respect to time. In essence, it is multiplying speed times time continuously to keep a running total of how far the car has gone. When the car is traveling at a high speed, the odometer “integrates” at a faster rate. When the car is traveling slowly, the odometer “integrates” slowly.

If the car travels in reverse, the odometer will decrement (count down) rather than increment (count up) because it sees a negative quantity for speed⁵. The rate at which the odometer decrements depends on how fast the car travels in reverse. When the car is stopped (zero speed), the odometer holds its reading and neither increments nor decrements.

Now let us return to the context of an automated process to see how this calculus principle works inside a process controller. Integration is provided either by a pneumatic mechanism, an electronic opamp circuit, or by a microprocessor executing a digital integration algorithm. The variable being integrated is *error* (the difference between PV and SP) over time. Thus the integral mode of the controller ramps the output either up or down over time in response to the amount of error existing between PV and SP, and the sign of that error. We saw this “ramping” action in the behavior of the liquid level control system using a motor-driven control valve commanded by a float-positioned potentiometer: the valve stem continuously moves so long as the liquid level deviates from setpoint. The reason for this ramping action is to increase or decrease the output *as far as it is necessary* in order to completely eliminate any error and force the process variable to precisely equal setpoint. Unlike proportional action, which simply moves the output an amount proportional to any change in PV or SP, integral control action never stops moving the output until all error is eliminated.

⁴An older term for this mode of control is *floating*, which I happen to think is particularly descriptive. With a “floating” controller, the final control element continually “floats” to whatever value it must in order to completely eliminate offset.

⁵At least the old-fashioned mechanical odometers would. Modern cars use a pulse detector on the driveshaft which cannot tell the difference between forward and reverse, and therefore their odometers always increment. Shades of the movie *Ferris Bueller’s Day Off*.

If proportional action is defined by the error telling the output how *far* to move, integral action is defined by the error telling the output how *fast* to move. One might think of integral as being how “impatient” the controller is, with integral action constantly ramping the output as far as it needs to go in order to eliminate error. Once the error is zero (PV = SP), of course, the integral action stops ramping, leaving the controller output (valve position) at its last value just like a stopped car’s odometer holds a constant value.

If we add an integral term to the controller equation, we get something that looks like this⁶:

$$m = K_p e + \frac{1}{\tau_i} \int e dt + b$$

Where,

- m = Controller output
- e = Error (difference between PV and SP)
- K_p = Proportional gain
- τ_i = Integral time constant (minutes)
- t = Time
- b = Bias

The most confusing portion of this equation for those new to calculus is the part that says “ $\int e dt$ ”. The integration symbol (looks like an elongated letter “S”) tells us the controller will accumulate (“sum”) multiple products of error (e) over tiny slices of time (dt). Quite literally, the controller multiplies error by time (for very short segments of time, dt) and continuously adds up all those products to contribute to the output signal which then drives the control valve (or other final control element). The integral time constant (τ_i) is a value set by the technician or engineer configuring the controller, proportioning this cumulative action to make it more or less aggressive over time.

To see how this works in a practical sense, let’s imagine how a proportional + integral controller would respond to the scenario of a heat exchanger whose inlet temperature suddenly dropped. As we saw with proportional-only control, an inevitable offset occurs between PV and SP with changes in load, because an error *must* develop if the controller is to generate the different output signal value necessary to halt further change in PV. We called this effect *proportional-only offset*.

Once this error develops, though, integral action begins to work. Over time, a larger and larger quantity accumulates in the integral mechanism (or register) of the controller due to the persistent error between PV and SP. That accumulated value adds to the controller’s output, driving the steam control valve further and further open. This, of course, adds heat at a faster rate to the heat exchanger, which causes the outlet temperature to rise. As the temperature re-approaches setpoint, the error becomes smaller and thus the integral action proceeds at a slower rate (like a car’s odometer incrementing at a slower rate as the car’s speed decreases). So long as the PV is below SP (the outlet temperature is still too cool), the controller will continue to integrate upwards, driving the control valve further and further open. Only when the PV rises to exactly meet SP does

⁶The equation for a proportional + integral controller is often written without the bias term (b), because the presence of integral action makes it unnecessary. In fact, if we let the integral term completely replace the bias term, we may consider the integral term to be a self-*resetting* bias. This, in fact, is the meaning of the word “reset” in the context of PID controller action: the “reset” term of the controller acts to eliminate offset by continuously adjusting (resetting) the bias as necessary.

integral action finally rest, holding the valve at a steady position. Integral action tirelessly works to eliminate any offset between PV and SP, thus neatly eliminating the offset problem experienced with proportional-only control action.

As with proportional action, there are (unfortunately) two completely opposite ways to specify the degree of integral action offered by a controller. One way is to specify integral action in terms of *minutes* or *minutes per repeat*. A large value of “minutes” for a controller’s integral action means a less aggressive integral action over time, just as a large value for proportional band means a less aggressive proportional action. The other way to specify integral action is the inverse: how many *repeats per minute*, equivalent to specifying proportional action in terms of gain (large value means aggressive action). For this reason, you will sometimes see the integral term of a PID equation written differently:

$$\begin{array}{ll} \tau_i = \text{minutes per repeat} & K_i = \text{repeats per minute} \\ \frac{1}{\tau_i} \int e \, dt & K_i \int e \, dt \end{array}$$

Many modern digital electronic controllers allow the user to select the unit they wish to use for integral action, just as they allow a choice between specifying proportional action as gain or as proportional band.

Integral is a highly effective mode of process control. In fact, some processes respond so well to integral controller action that it is possible to operate the control loop on integral action alone, without proportional. Typically, though, process controllers implement some form of proportional plus integral (“PI”) control.

Just as too much proportional gain will cause a process control system to oscillate, too much integral action (i.e. an integral time constant that is too short) will also cause oscillation. If the integration happens at too fast a rate, the controller’s output will “saturate” either high or low before the process variable can make it back to setpoint. Once this happens, the only condition that will “unwind” the accumulated integral quantity is for an error to develop of the opposite sign, and remain that way long enough for a canceling quantity to accumulate. Thus, the PV must cross over the SP, guaranteeing at least another half-cycle of oscillation.

A similar problem called *reset windup* (or *integral windup*) happens when external conditions make it impossible for the controller to achieve setpoint. Imagine what would happen in the heat exchanger system if the steam boiler suddenly stopped producing steam. As outlet temperature dropped, the controller’s proportional action would open up the control valve in a futile effort to raise temperature. If and when steam service is restored, proportional action would just move the valve back to its original position as the process variable returned to its original value (before the boiler died). This is how a proportional-only controller would respond to a steam “outage”: nice and predictably. If the controller had integral action, however, a much worse condition would result. All the time spent with the outlet temperature below setpoint causes the controller’s integral term to “wind up” in a futile attempt to admit more steam to the heat exchanger. This accumulated quantity can only be un-done by the process variable rising above setpoint for an equal error-time product⁷, which means when the steam supply resumes, the temperature will rise well above setpoint

⁷Since integration is fundamentally a process of multiplication followed by addition, the units of measurement are always the product (multiplication) of the function’s variables. In the case of reset (integral) control, we are multiplying

until the integral action finally “unwinds” and brings the control valve back to a same position again.

Various techniques exist to manage integral windup. Controllers may be built with limits to restrict how far the integral term can accumulate under adverse conditions. In some controllers, integral action may be turned off completely if the error exceeds a certain value. The surest fix for integral windup is human operator intervention, by placing the controller in manual mode. This typically resets the integral accumulator to a value of zero and loads a new value into the bias term of the equation to set the valve position wherever the operator decides. Operators usually wait until the process variable has returned at or near setpoint before releasing the controller into automatic mode again.

While it might appear that operator intervention is again a problem to be avoided (as it was in the case of having to correct for proportional-only offset), it is noteworthy to consider that the conditions leading to integral windup usually occur only during shut-down conditions. It is customary for human operators to run the process manually anyway during a shutdown, and so the switch to manual mode is something they would do anyway and the potential problem of windup often never manifests itself.

Integral control action has the unfortunate tendency to create loop oscillations (“cycling”) if the final control element exhibits hysteresis, such as the case with a “sticky” control valve. Imagine for a moment our steam-heated heat exchanger system where the steam control valve possesses excessive packing friction and therefore refuses to move until the applied air pressure changes far enough to overcome that friction, at which point the valve “jumps” to a new position and then “sticks” in that new position. If the valve happens to stick at a stem position resulting in the product temperature settling slightly below setpoint, the controller’s integral action will continually increase the output signal going to the valve in an effort to correct this error (as it should). However, when that output signal has risen far enough to overcome valve friction and move the stem further open, it is very likely the stem will once again “stick” but this time do so at a position making the product temperature settle *above* setpoint. The controller’s integral action will then ramp downward in an effort to correct this new error, but due to the valve’s friction making precise positioning impossible, the controller can never achieve setpoint and therefore it cyclically “hunts” above and below setpoint.

The best solution to this “reset cycling” phenomenon, of course, is to correct the hysteresis in the final control element. Eliminating friction in the control valve will permit precise positioning and allow the controller’s integral action to achieve setpoint as designed. Since it is practically impossible to eliminate *all* friction from a control valve, however, other solutions to this problem exist. One of them is to program the controller to stop integrating whenever the error is less than some pre-configured value (sometimes referred to as the “integral deadband” or “reset deadband” of the controller). By activating reset control action only for significant error values, the controller ignores small errors rather than “compulsively” trying to correct for any detected error no matter how small.

controller error (the difference between PV and SP, usually expressed in percent) by time (usually expressed in minutes or seconds). Therefore the result will be an “error-time” product. In order for an integral controller to self-recover following windup, the error must switch signs and the error-time product accumulate to a sufficient value to cancel out the error-time product accumulated during the windup period.

29.7 Derivative (rate) control

The final element of PID control is the “D” term, which stands for *derivative*. This is a calculus concept like integral, except most people consider it easier to understand. Simply put, derivative is the expression of a variable’s *rate-of-change* with respect to another variable. Finding the derivative of a function (differentiation) is the inverse operation of integration. With integration, we calculated accumulated value of some variable’s product with time. With derivative, we calculate the ratio of a variable’s change per unit of time. Whereas integration is fundamentally a multiplicative operation (products), differentiation always involves division (ratios).

A controller with derivative (or *rate*) action looks at how fast the process variable changes per unit of time, and takes action proportional to that rate of change. In contrast to integral (reset) action which represents the “impatience” of the controller, derivative (rate) action represents the “caution” of the controller.

If the process variable starts to change at a high rate of speed, the job of derivative action is to move the final control element in such a direction as to counteract this rapid change, and thereby moderate the speed at which the process variable changes. In simple terms, derivative action works to limit how fast the error can change.

What this will do is make the controller “cautious” with regard to rapid changes in process variable. If the process variable is headed toward the setpoint value at a rapid rate, the derivative term of the equation will diminish the output signal, thus tempering the controller’s response and slowing the process variable’s approach toward setpoint. This is analogous to a truck driver preemptively applying the brakes to slow the approach to an intersection, knowing that the heavy truck doesn’t “stop on a dime.” The heavier the truck’s load, the sooner a cautious driver will apply the brakes, to avoid “overshoot” beyond the stop sign and into the intersection. For this reason, derivative control action is also called *pre-act* in addition to being called *rate*, because it acts “ahead of time” to avoid overshoot.

If we modify the controller equation to incorporate differentiation, it will look something like this:

$$m = K_p e + \frac{1}{\tau_i} \int e dt + \tau_d \frac{de}{dt} + b$$

Where,

m = Controller output

e = Error (difference between PV and SP)

K_p = Proportional gain

τ_i = Integral time constant (minutes)

τ_d = Derivative time constant (minutes)

t = Time

b = Bias

The $\frac{de}{dt}$ term of the equation expresses the rate of change of error (e) over time (t). The lower-case letter “d” symbols represent the calculus concept of *differentials* which may be thought of in this context as very tiny increments of the following variables. In other words, $\frac{de}{dt}$ refers to the ratio of a very small change in error (de) over a very small increment of time (dt). On a graph, this is interpreted as the slope of a curve at a specific point (slope being defined as *rise over run*).

It is also possible to build a controller with proportional and derivative actions, but lacking integral action. These are most commonly used in applications prone to wind-up⁸, and where the elimination of offset is not critical:

$$m = K_p e + \tau_d \frac{de}{dt} + b$$

Many PID controllers offer the option of calculating derivative response based on rates of change for the process variable (PV) only, rather than the error (PV – SP or SP – PV). This avoids huge “spikes” in the output of the controller if ever a human operator makes a sudden change in setpoint⁹. The mathematical expression for such a controller would look like this:

$$m = K_p e + \frac{1}{\tau_i} \int e dt + \tau_d \frac{dPV}{dt} + b$$

Even when derivative control action is calculated on PV alone (rather than on error), it is still useful for controlling processes dominated by large lag times. The presence of derivative control action in a PID controller generally means the proportional (P) and integral (I) terms may be adjusted more aggressively than before, since derivative (D) will act to limit overshoot. In other words, the judicious presence of derivative action in a PID controller lets us “get away” with using a bit more P and I action than we ordinarily could, resulting in faster approach to setpoint with minimal overshoot.

It should be mentioned that derivative mode should be used with caution. Since it acts on rates of change, derivative action will “go crazy” if it sees substantial noise in the PV signal. Even small amounts of noise possess extremely large rates of change (defined as percent PV change per minute of time) owing to the relatively high frequency of noise compared to the timescale of physical process changes.

Ziegler and Nichols, the engineers who wrote the ground-breaking paper entitled “Optimum Settings for Automatic Controllers” had these words to say regarding “pre-act” control (page 762 of the November 1942 *Transactions of the A.S.M.E.*):

The latest control effect made its appearance under the trade name “Pre-Act.” On some control applications, the addition of pre-act response made such a remarkable improvement that it appeared to be in embodiment of mythical “anticipatory” controllers. On other applications it appeared to be worse than useless. Only the difficulty of predicting the usefulness and adjustment of this response has kept it from being more widely used.

⁸An example of such an application is where the output of a loop controller may be “de-selected” or otherwise “over-ridden” by some other control function. This sort of control strategy is often used in energy-conserving controls, where multiple controllers monitoring different process variables selectively command a single FCE.

⁹It should not be assumed that such spikes are always undesirable. In processes characterized by long lag times, such a response may be quite helpful in overcoming that lag for the purpose of rapidly achieving new setpoint values. Slave (secondary) controllers in cascaded systems – where the controller receives its setpoint signal from the output of another (primary, or master) controller – may similarly benefit from derivative action calculated on error instead of just PV. As usual, the specific needs of the application dictate the ideal controller configuration.

29.8 Summary of PID control terms

PID control can be a confusing concept to understand. Here, a brief summary of each term within PID (P, I, and D) is presented for your learning benefit.

29.8.1 Proportional control mode (P)

Proportional – sometimes called *gain* or *sensitivity* – is a control action reproducing changes in input as changes in output. Proportional controller action responds to present changes in input by generating immediate and commensurate changes in output. When you think of “proportional action” (P), think *prompt*: this control action works immediately (never too soon or too late) to match changes in the input signal.

Mathematically defined, proportional action is the ratio of output change to input change. This may be expressed as a quotient of differences, or as a derivative (a rate of change, using calculus notation):

$$\text{Gain value} = \frac{\Delta\text{Output}}{\Delta\text{Input}}$$

$$\text{Gain value} = \frac{d\text{Output}}{d\text{Input}} = \frac{dm}{de}$$

For example, if the PV input of a proportional-only process controller with a gain of 2 suddenly changes (“steps”) by 5 percent, and the output will immediately jump by 10 percent ($\Delta\text{Output} = \text{Gain} \times \Delta\text{Input}$). The direction of this output jump in relation to the direction of the input jump depends on whether the controller is configured for direct or reverse action.

A legacy term used to express this same concept is *proportional band*: the mathematical reciprocal of gain. “Proportional band” is defined as the amount of input change necessary to evoke full-scale (100%) output change in a proportional controller. Incidentally, it is always expressed as a percentage, never as fraction or as a per unit value:

$$\text{Proportional Band value} = \frac{\Delta\text{Input}}{\Delta\text{Output}}$$

$$\text{Proportional Band value} = \frac{d\text{Input}}{d\text{Output}} = \frac{de}{dm}$$

Using the same example of a proportional controller exhibiting an output “step” of 10% in response to a PV “step” of 5%, the proportional band would be 50%: the reciprocal of its gain ($\frac{1}{2} = 50\%$). Another way of saying this is that a 50% input “step” would be required to change the output of this controller by a full 100%, since its gain is set to a value of 2.

29.8.2 Integral control mode (I)

Integral – sometimes called *reset* or *floating control* – is a control action causing the output signal to change over time at a rate proportional to the amount of error (the difference between PV and SP values). Integral controller action responds to error accumulated over time, ramping the output signal as far as it needs to go to completely eliminate error. If proportional (P) action tells the output how *far* to move when an error appears, integral (I) action tells the output how *fast* to move when an error appears. If proportional (P) action acts on the *present*, integral (I) action acts on the *past*. Thus, how far the output signal gets driven by integral action depends on the *history* of the error over time: how much error existed, and for how long. When you think of “integral action” (I), think *impatience*: this control action drives the output further and further the longer PV fails to match SP.

Mathematically defined, integral action is the ratio of output *velocity* to input error:

$$\text{Integral value (repeats per minute)} = \frac{\text{Output velocity}}{\text{Input error}}$$

$$\text{Integral value (repeats per minute)} = \frac{\frac{dm}{dt}}{e}$$

An alternate way to express integral action is to use the reciprocal unit of “minutes per repeat.” If we define integral action in these terms, the defining equations must be reciprocated:

$$\text{Integral time constant (minutes per repeat)} = \tau_i = \frac{\text{Input error}}{\text{Output velocity}}$$

$$\text{Integral time constant (minutes per repeat)} = \tau_i = \frac{e}{\frac{dm}{dt}}$$

For example, if an error of 5% appears between PV and SP on an integral-only process controller with an integral value of 3 repeats per minute (i.e. an integral time constant of 0.333 minutes per repeat), the output will begin ramping at a rate of 15% per minute ($\frac{dm}{dt} = \text{Integral_value} \times e$, or $\frac{dm}{dt} = \frac{e}{\tau_i}$). In most PI and PID controllers, integral response is also multiplied by proportional gain, so the same conditions applied to a PI controller that happened to also have a gain of 2 would result in an output ramping rate of 30% per minute ($\frac{dm}{dt} = \text{Gain_value} \times \text{Integral_value} \times e$, or $\frac{dm}{dt} = \text{Gain_value} \times \frac{e}{\tau_i}$). The direction of this ramping in relation to the direction (sign) of the error depends on whether the controller is configured for direct or reverse action.

29.8.3 Derivative control mode (D)

Derivative – sometimes called *rate* or *pre-act* – is a control action causing the output signal to be offset by an amount proportional to the rate at which the input is changing. Derivative controller action responds to how quickly the input changes over time, biasing the output signal commensurate with that rate of input change. If proportional (P) action tells the output how *far* to move when an error appears, derivative (D) action tells the output how far to move when the input *ramps*. If proportional (P) action acts on the *present* and integral (I) action acts on the *past*, derivative (D) action acts on the *future*: it effectively “anticipates” overshoot by tempering the output response according to how fast the process variable is rising or falling. When you think of “derivative action” (D), think *discretion*: this control action is cautious and prudent, working against change.

Mathematically defined, derivative action is the ratio of output offset to input *velocity*:

$$\text{Derivative time constant (minutes)} = \tau_d = \frac{\text{Output offset}}{\text{Input velocity}}$$

$$\text{Derivative time constant (minutes)} = \tau_d = \frac{\Delta\text{Output}}{\frac{de}{dt}}$$

For example, if the PV signal begins to ramp at a rate of 5% per minute on a process controller with a derivative time constant of 4 minutes, the output will immediately become offset by 20% ($\Delta\text{Output} = \text{Derivative_value} \times \frac{de}{dt}$). In most PD and PID controllers, derivative response is also multiplied by proportional gain, so the same conditions applied to a PD controller that happened to also have a gain of 2 would result in an immediate offset of 40% ($\Delta\text{Output} = \text{Gain_value} \times \text{Derivative_value} \times \frac{de}{dt}$). The direction (sign) of this offset in relation to the direction of the input ramping depends on whether the controller is configured for direct or reverse action.


29.9 P, I, and D responses graphed

A very helpful method for understanding the operation of proportional, integral, and derivative control terms is to analyze their respective responses to the same input conditions over time. This section is divided into subsections showing P, I, and D responses for several different input conditions, in the form of graphs. In each graph, the controller is assumed to be *direct-acting* (i.e. an increase in process variable results in an increase in output).

It should be noted that these graphic illustrations are all qualitative, not quantitative. There is too little information given in each case to plot exact responses. The illustrations of P, I, and D actions focus only on the *shapes* of the responses, not their exact numerical values.


In order to *quantitatively* predict PID controller responses, one would have to know the values of all PID settings, as well as the original starting value of the output before an input change occurred and a time index of when the change(s) occurred.

29.9.1 Responses to a single step-change




Proportional action directly mimics the shape of the input change (a step). Integral action ramps at a rate proportional to the magnitude of the input step. Since the input step holds a constant value, the integral action ramps at a constant rate (a constant *slope*). Derivative action interprets the step as an *infinite* rate of change, and so generates a “spike¹⁰” driving the output to saturation.

When combined into one PID output, the three actions produce this response:



¹⁰This is the meaning of the vertical-pointing arrowheads shown on the trend graph: momentary saturation of the output all the way up to 100%.

29.9.2 Responses to a momentary step-and-return




Proportional action directly mimics the shape of the input change (an up-and-down step). Integral action ramps at a rate proportional to the magnitude of the input step, for as long as the PV is unequal to the SP. Once $PV = SP$ again, integral action stops ramping and simply holds the last value¹¹. Derivative action interprets both steps as *infinite* rates of change, and so generates “spikes¹²” at the leading and at the trailing edges of the step. Note how the leading (rising) edge causes derivative action to saturate high, while the trailing (falling) edge causes it to saturate low.


¹¹This is a good example of how integral controller action represents the *history* of the $PV - SP$ error. The continued offset of integral action from its starting point “remembers” the area accumulated under the rectangular “step” between PV and SP. This offset will go away only if a *negative* error appears having the same percent-minute product (area) as the positive error step.

¹²This is the meaning of the vertical-pointing arrowheads shown on the trend graph: momentary saturation of the output all the way up to 100% (or down to 0%).

When combined into one PID output, the three actions produce this response:




29.9.3 Responses to two momentary steps-and-returns




Proportional action directly mimics the shape of all input changes. Integral action ramps at a rate proportional to the magnitude of the input step, for as long as the PV is unequal to the SP. Once $PV = SP$ again, integral action stops ramping and simply holds the last value. Derivative action interprets each step as an *infinite* rate of change, and so generates a “spike” at the leading and at the trailing edges of each step. Note how a leading (rising) edge causes derivative action to saturate high, while a trailing (falling) edge causes it to saturate low.

When combined into one PID output, the three actions produce this response:




29.9.4 Responses to a ramp-and-hold




Proportional action directly mimics the ramp-and-hold shape of the input. Integral action ramps slowly at first (when the error is small) but increases ramping rate as error increases. When error stabilizes, integral rate likewise stabilizes. Derivative action offsets the output according to the input's ramping rate.

When combined into one PID output, the three actions produce this response:




29.9.5 Responses to an up-and-down ramp




Proportional action directly mimics the up-and-down ramp shape of the input. Integral action ramps slowly at first (when the error is small) but increases ramping rate as error increases, then ramps slower as error decreases back to zero. Once PV = SP again, integral action stops ramping and simply holds the last value. Derivative action offsets the output according to the input's ramping rate: first positive then negative.

When combined into one PID output, the three actions produce this response:




29.9.6 Responses to a multi-slope ramp




Proportional action directly mimics the ramp shape of the input. Integral action ramps slowly at first (when the error is small) but increases ramping rate as error increases, then accelerates its increase as the PV ramps even steeper. Once PV = SP again, integral action stops ramping and simply holds the last value. Derivative action offsets the output according to the input's ramping rate: first positive, then more positive, then it spikes negative when the PV suddenly returns to SP.

When combined into one PID output, the three actions produce this response:




29.9.7 Responses to a multiple ramps and steps




Proportional action directly mimics the ramp-and-step shape of the input. Integral action ramps slowly at first (when the error is small) but increases ramping rate as error increases. Which each higher ramp-and-step in PV, integral action winds up at an ever-increasing rate. Since PV never equals SP again, integral action never stops ramping upward. Derivative action steps with each ramp of the PV.

When combined into one PID output, the three actions produce this response:



29.9.8 Responses to a sine wavelet



As always, proportional action directly mimics the shape of the input. The 90° phase shift seen in the integral and derivative responses, compared to the PV wavelet, is no accident or coincidence. The derivative of a sinusoidal function is *always* a cosine function, which is mathematically identical to a sine function with the angle advanced by 90° :

$$\frac{d}{dx}(\sin x) = \cos x = \sin(x + 90^\circ)$$


Conversely, the integral of a sine function is *always* a negative cosine function¹³, which is mathematically identical to a sine function with the angle retarded by 90° :

$$\int \sin x \, dx = -\cos x = \sin(x - 90^\circ)$$

In summary, the derivative operation always adds a positive (leading) phase shift to a sinusoidal input waveform, while the integral operation always adds a negative (lagging) phase shift to a sinusoidal input waveform.

¹³In this example, I have omitted the constant of integration (C) to keep things simple. The actual integral is as such: $\int \sin x \, dx = -\cos x + C = \sin(x - 90^\circ) + C$. This constant value is essential to explaining why the integral response does not immediately “step” like the derivative response does at the beginning of the PV sine wavelet.

When combined into one PID output, these particular integral and derivative actions mostly cancel, since they happen to be sinusoidal wavelets of equal amplitude and opposite phase. Thus, the only way that the final (PID) output differs from proportional-only action in this particular case is the “steps” caused by derivative action responding to the input’s sudden rise at the beginning and end of the wavelet:




If the I and D tuning parameters were such that the integral and derivative responses were *not* equal in amplitude, their effects would not completely cancel. Rather, the resultant of P, I, and D actions would be a sine wavelet having a phase shift somewhere between -90° and $+90^\circ$ exclusive, depending on the relative strengths of the P, I, and D actions.

The 90 degree phase shifts associated with the integral and derivative operations are useful to understand when tuning PID controllers. If one is familiar with these phase shift relationships, it is relatively easy to analyze the response of a PID controller to a sinusoidal input (such as when a process oscillates following a sudden load or setpoint change) to determine if the controller’s response is dominated by any one of the three actions. This may be helpful in “de-tuning” an over-tuned (overly aggressive) PID controller, if an excess of P, I, or D action may be identified from a phase comparison of PV and output waveforms.

29.9.9 Note to students regarding quantitative graphing


A common exercise for students learning the function of PID controllers is to practice graphing a controller's output given input (PV and SP) conditions, either qualitatively or quantitatively. This can be a frustrating experience for some students, as they struggle to accurately combine the effects of P, I, and/or D responses into a single output trend. Here, I will present a way to ease the pain.

Suppose for example you were tasked with graphing the response of a PD (proportional + derivative) controller to the following PV and SP inputs over time. You are told the controller has a gain of 1, a derivative time constant of 0.3 minutes, and is reverse-acting:




My first recommendation is to *qualitatively* sketch the individual P and D responses. Simply draw two different trends, each one right above or below the given PV/SP trends, showing the shapes of each response over time. You might even find it easier to do if you re-draw the original PV and SP trends on a piece of non-graph paper with the qualitative P and D trends also sketched on the same piece of non-graph paper. The purpose of the qualitative sketches is to separate the task of determining shapes from the task of determining numerical values, in order to simplify the process.

After sketching the separate P and D trends, label each one of the “features” (changes either up or down) in these qualitative trends. This will allow you to more easily combine the effects into one output trend later:



Now, you may qualitatively sketch an output trend combining each of these “features” into one graph. Be sure to label each ramp or step originating with the separate P or D trends, so you know where each “feature” of the combined output graph originates from:




Once the general shape of the output has been qualitatively determined, you may go back to the separate P and D trends to calculate numerical values for each of the labeled “features.”


Note that each of the PV ramps is 15% in height, over a time of 15 seconds (one-quarter of a minute). With a controller gain of 1, the proportional response to each of these ramps will also be a ramp that is 15% in height.

Taking our given derivative time constant of 0.3 minutes and multiplying that by the PV’s rate-of-change ($\frac{dPV}{dt}$) during each of its ramping periods (15% per one-quarter minute, or 60% per minute) yields a derivative response of 18% during each of the ramping periods. Thus, each derivative response “step” will be 18% in height.

Going back to the qualitative sketches of P and D actions, and to the combined (qualitative) output sketch, we may apply the calculated values of 15% for each proportional ramp and 18% for each derivative step to the labeled “features.” We may also label the starting value of the output trend as given in the original problem (35%), to calculate actual output values at different points in time. Calculating output values at specific points in the graph becomes as easy as cumulatively adding and subtracting the P and D “feature” values to the starting output value:



Now that we know the output values at all the critical points, we may quantitatively sketch the output trend on the original graph:



29.10 Different PID equations

For better or worse, there are no fewer than *three* different forms of PID equations implemented in modern PID controllers: the *parallel*, *ideal*, and *series*. Some controllers offer the choice of more than one equation, while others implement just one. It should be noted that more variations of PID equation exist than these three, but that these are the three major variations.

29.10.1 Parallel PID equation

The equation used to describe PID control so far in this chapter is the simplest form, sometimes called the *parallel* equation, because each action (P, I, and D) occurs in separate terms of the equation, with the combined effect being a simple sum:

$$m = K_p e + \frac{1}{\tau_i} \int e \, dt + \tau_d \frac{de}{dt} + b \quad \text{Parallel PID equation}$$

In the parallel equation, each action parameter (K_p , τ_i , τ_d) is independent of the others. At first, this may seem to be an advantage, for it means each adjustment made to the controller should only affect one aspect of its action. However, there are times when it is better to have the gain parameter affect all three control actions (P, I, and D)¹⁴.

We may show the independence of the three actions mathematically, by breaking the equation up into three different parts, each one describing its contribution to the output (Δm):

$$\Delta m = K_p \Delta e \quad \text{Proportional action}$$

$$\Delta m = \frac{1}{\tau_i} \int e \, dt \quad \text{Integral action}$$

$$\Delta m = \tau_d \frac{de}{dt} \quad \text{Derivative action}$$

As you can see, the three portions of this PID equation are completely separate, with each tuning parameter (K_p , τ_i , and τ_d) acting independently within its own term of the equation.

¹⁴An example of a case where it is better for gain (K_p) to influence all three control modes is when a technician re-ranges a transmitter to have a larger or smaller span than before, and must re-tune the controller to maintain the same loop gain as before. If the controller's PID equation takes the parallel form, the technician must adjust the P, I, and D tuning parameters proportionately. If the controller's PID equation uses K_p as a factor in all three modes, the technician need only adjust K_p to re-stabilize the loop.

29.10.2 Ideal PID equation

An alternate version of the PID equation designed such that the gain (K_p) affects all three actions is called the *Ideal* or *ISA* equation:

$$m = K_p \left(e + \frac{1}{\tau_i} \int e \, dt + \tau_d \frac{de}{dt} \right) + b \quad \text{Ideal or ISA PID equation}$$

Here, the gain constant (K_p) is distributed to all terms within the parentheses, equally affecting all three control actions. Increasing K_p in this style of PID controller makes the P, the I, *and* the D actions equally more aggressive.

We may show this mathematically, by breaking the “ideal” equation up into three different parts, each one describing its contribution to the output (Δm):

$$\Delta m = K_p \Delta e \quad \text{Proportional action}$$

$$\Delta m = \frac{K_p}{\tau_i} \int e \, dt \quad \text{Integral action}$$

$$\Delta m = K_p \tau_d \frac{de}{dt} \quad \text{Derivative action}$$

As you can see, all three portions of this PID equation are influenced by the gain (K_p) owing to algebraic distribution, but the integral and derivative tuning parameters (τ_i and τ_d) act independently within their own terms of the equation.

29.10.3 Series PID equation

A third version, with origins in the peculiarities of pneumatic controller mechanisms and analog electronic circuits, is called the *Series* or *Interacting* equation:

$$m = K_p \left[\left(\frac{\tau_d}{\tau_i} + 1 \right) e + \frac{1}{\tau_i} \int e \, dt + \tau_d \frac{de}{dt} \right] + b \quad \text{Series or Interacting PID equation}$$

Here, the gain constant (K_p) affects all three actions (P, I, and D) just as with the “ideal” equation. The difference, though, is the fact that both the integral and derivative constants have an effect on proportional action as well! That is to say, adjusting either τ_i or τ_d does not merely adjust those actions, but also influences the aggressiveness of proportional action¹⁵.

We may show this mathematically, by breaking the “series” equation up into three different parts, each one describing its contribution to the output (Δm):

$$\Delta m = K_p \left(\frac{\tau_d}{\tau_i} + 1 \right) \Delta e \quad \text{Proportional action}$$

$$\Delta m = \frac{K_p}{\tau_i} \int e \, dt \quad \text{Integral action}$$

$$\Delta m = K_p \tau_d \frac{de}{dt} \quad \text{Derivative action}$$

As you can see, all three portions of this PID equation are influenced by the gain (K_p) owing to algebraic distribution. However, the proportional term is also affected by the values of the integral and derivative tuning parameters (τ_i and τ_d). Therefore, adjusting τ_i affects both the I and P actions, adjusting τ_d affects both the D and P actions, and adjusting K_p affects all three actions.

This “interacting” equation is an artifact of certain pneumatic and electronic controller designs. Back when these were the dominant technologies, and PID controllers were modularly designed such that integral and derivative actions were separate hardware modules included in a controller at additional cost beyond proportional-only action, the easiest way to implement the integral and derivative actions was in a way that just happened to have an interactive effect on controller gain. In other words, this odd equation form was a sort of compromise made for the purpose of simplifying the physical design of the controller.

Interestingly enough, many digital PID controllers are programmed to implement the “interacting” PID equation even though it is no longer an artifact of controller hardware. The rationale for this programming is to have the digital controller behave identically to the legacy analog electronic or pneumatic controller it is replacing. This way, the proven tuning parameters of the old controller may be plugged into the new digital controller, yielding the same results. In essence, this is a form of “backward compatibility” between digital PID control and analog (electronic or pneumatic) PID control.

¹⁵This becomes especially apparent when using derivative action with low values of τ_i (aggressive integral action). The error-multiplying term $\frac{\tau_d}{\tau_i} + 1$ may become quite large if τ_i is small, even with modest τ_d values.


29.11 Pneumatic PID controllers

A *pneumatic* controller receives a process variable (PV) signal as a variable air pressure, compares that signal against a desired setpoint (SP) value, and then mechanically generates another air pressure signal as the output, driving a final control element.

Throughout this section I will make reference to a pneumatic controller mechanism of my own design. This mechanism does not directly correspond to any particular manufacturer or model of pneumatic controller, but shares characteristics common to many. This design is shown here for the purpose of illustrating the development of P, I, and D control actions in as simple a context as possible.

29.11.1 Proportional control action

Many pneumatic PID controllers use the *force-balance* principle. One or more input signals (in the form of pneumatic pressures) exert a force on a beam by acting through diaphragms, bellows, and/or bourdon tubes, which is then counter-acted by the force exerted on the same beam by an output air pressure acting through a diaphragm, bellows, or bourdon tube. The self-balancing mechanical system “tries” to keep the beam motionless through an exact balancing of forces, the beam’s position precisely detected by a nozzle/baffle mechanism:




The action of this particular controller is *direct*, since an increase in process variable signal (pressure) results in an increase in output signal (pressure). Increasing process variable (PV) pressure attempts to push the right-hand end of the beam up, causing the baffle to approach the nozzle. This blockage of the nozzle causes the nozzle’s pneumatic backpressure to increase, thus increasing the amount of force applied by the output feedback bellows on the left-hand end of the beam and returning the flapper (very nearly) to its original position. If we wished to reverse the controller’s action, all we would need to do is swap the pneumatic signal connections between the input bellows, so that the PV pressure was applied to the upper bellows and the SP pressure to the lower bellows.

Any factor influencing the ratio of input pressure(s) to output pressure may be exploited as a gain (proportional band) adjustment in this mechanism. Changing bellows area (either both the PV and SP bellows equally, or the output bellows by itself) would influence this ratio, as would a change in output bellows position (such that it pressed against the beam at some difference distance

from the fulcrum point). Moving the fulcrum left or right is also an option for gain control, and in fact is usually the most convenient to engineer.


In this illustration the fulcrum is shown moved to two different positions, to effect a change in gain:



Moving the fulcrum closer to the output bellows places that bellows at a mechanical disadvantage for generating torque (leverage) on the beam. This means any given change in input (PV or SP) force is more difficult for the output bellows to counterbalance. The output pressure, therefore, must change to a greater degree in order for this force-balance mechanism to achieve balance. A greater change in output pressure for a given change in input pressure is the definition of a gain *increase*.

Conversely, moving the fulcrum farther away from the output bellows increases that bellows' mechanical advantage. This additional leverage makes it easier for the output bellows to counteract changes in input force, resulting in less output pressure change required to balance any given input pressure change. A lesser change in output pressure for a given change in input pressure is characteristic of a gain *decrease*.

Some pneumatic controllers employ the *motion-balance* principle instead of the force-balance principle in their operation. In contrast to a force-balance system where opposing forces cancel each other to restrain motion of the mechanism, a motion-balance system freely moves as the signal pressures traverse their working ranges. A simple motion-balance proportional controller design appears here:




As the process variable signal increases, the right-hand end of the lever is forced up. This motion draws the lever away from the nozzle, resulting in decreased nozzle backpressure. The decreased backpressure causes the output bellows to collapse¹⁶, moving the left-hand end of the lever down and returning the nozzle/lever gap to (approximately) where it was before the PV signal change. This behavior identifies this controller as *reverse-acting*. If direct action were desired, all we would need to do is swap the process variable and setpoint input pressure connections.

Unlike the force-balance controller mechanism where the lever is maintained in an essentially stationary position by equal and opposite forces, the lever in this motion-balance system is free to tilt. In fact, tilting is precisely how a (nearly) constant nozzle gap is maintained: as one end of the lever moves (either up or down), the other end moves in the opposite direction to keep the nozzle/lever gap constant in the middle.

¹⁶Being a motion-balance mechanism, these bellows must act as spring elements in order to produce consistent pressure/motion behavior. Some pneumatic controllers employ coil springs inside the brass bellows assembly to provide the necessary “stiffness” and repeatability.

The gain of such a mechanism may be changed by moving the position of the nozzle along the lever's length. However, it must be understood that this position change will have the opposite effect on gain compared with the fulcrum position change described for the force-balance mechanism. Here in the motion-balance system, it is the relative *travel* of each bellows that matters for gain, not the relative *leverage* (torque):




With the nozzle positioned closer to the output bellows, that bellows need not stretch or collapse as much in order to maintain the nozzle gap constant even with a large motion at the input (right-hand) end of the lever. The output pressure in this case will change only slightly for large changes in PV or SP pressures: characteristic of a low gain.

Moving the nozzle closer to the input (PV and SP) bellows gives those bellows more influence over the nozzle/lever gap. The output bellows must expand and contract quite a bit more than the input bellows in order to maintain a constant nozzle gap for any motion at the input side. This requires a greater change in output pressure for a given change in input pressure: the definition of increased gain.

29.11.2 Automatic and manual modes

A more practical pneumatic proportional controller mechanism is shown in the next illustration, complete with setpoint and bias adjustments, and a manual control mode:



“Bumpless” transfer between automatic and manual modes is a very important feature for any loop controller because it allows human operators to change the mode of the controller without introducing an unnecessary disturbance to the process being controlled. Without provision for bumpless transfer, the output signal of the controller may suddenly change whenever the mode is switched between automatic and manual. This sudden signal change will cause the final control element to suddenly “step” to some new level of effect on the process.

In this particular pneumatic controller, bumpless auto/manual transfer is accomplished by the operator paying attention to the *balance indicator* revealing any air pressure difference between the output bellows and the output adjust pressure regulator. When in automatic mode, a switch to manual mode involves adjusting the output regulator until the balance indicator registers zero pressure difference, then switching the transfer valve to the “manual” position. The controller output is then at the direct command of the output adjust pressure regulator, and will not respond

to changes in either PV or SP. “Bumplessly” switching back to automatic mode requires that the setpoint pressure regulator be adjusted until the balance indicator once again registers zero pressure difference, then switching the transfer valve to the “auto” position. The controller output will once again respond to changes in PV and SP.


A photograph showing a Foxboro model 43AP pneumatic controller manual/auto transfer switch and balance indicator appears here:



The metal ball within the curved plastic tube indicates equal pressures between automatic and manual modes when centered in the tube. To achieve bumpless transfer between automatic and manual modes, one must never switch the auto/manual valve unless that ball is centered. To center the ball while in automatic mode, the manual output pressure must be adjusted to achieve balance with the automatic-mode output pressure. To center the ball while in manual mode, the automatic-mode output pressure must be adjusted to achieve balance with the manual-mode output pressure – a condition attained by adjusting the *setpoint* knob.

29.11.3 Derivative control action

Derivative (rate) control action is relatively easy to add to this pneumatic controller mechanism. All we need to do is place a restrictor valve between the nozzle tube and the output feedback bellows, causing the bellows to delay filling or emptying its air pressure over time:




If any sudden change occurs in PV or SP, the output pressure will saturate before the output bellows has the opportunity to equalize in pressure with the output signal tube. Thus, the output pressure “spikes” with any sudden “step change” in input: exactly what we would expect with derivative control action.

If either the PV or the SP ramps over time, the output signal will ramp in direct proportion (proportional action), but there will *also* be an added offset of pressure at the output signal in order to keep air flowing either in or out of the output bellows at a constant rate to generate the force necessary to balance the changing input signal. Thus, derivative action causes the output pressure to shift either up or down (depending on the direction of input change) more than it would with just proportional action alone in response to a ramping input: exactly what we would expect from a controller with both proportional and derivative control actions.

29.11.4 Integral control action

Adding integral action to our hypothetical pneumatic controller mechanism requires the placement of a second bellows (a “reset” bellows) opposite the output feedback bellows, and another restrictor valve to the mechanism¹⁷:




This second bellows takes air pressure from the output line and translates it into force that opposes the original feedback bellows. At first, this may seem counter-productive, for it nullifies the ability of this mechanism to continuously balance the force generated by the PV and SP bellows. Indeed, it would render the force-balance system completely ineffectual if this new “reset” bellows were allowed to inflate and deflate with no time lag. However, with a time lag provided by the

¹⁷Practical integral action also requires the elimination of the bias spring and adjustment, which formerly provided a constant downward force on the left-hand side of the beam to give the output signal the positive offset necessary to avoid saturation at 0 PSI. Not only is a bias adjustment completely unnecessary with the addition of integral action, but it would actually cause problems by making the integral action “think” an error existed between PV and SP when there was none.

restriction of the integral adjustment valve and the volume of the bellows (a sort of pneumatic “RC time constant”), the nullifying force of this bellows becomes delayed over time. As this bellows slowly fills (or empties) with pressurized air from the nozzle, the change in force on the beam causes the regular output bellows to have to “stay ahead” of the reset bellows action by constantly filling (or emptying) at some rate over time.


To better understand this integrating action, let us perform a “thought experiment” on a simplified version of the controller. The following mechanism has been stripped of all unnecessary complexity so that we may focus on just the proportional and integral actions. Here, we envision the PV and SP air pressure signals differing by 3 PSI, causing the force-balance mechanism to instantly respond with a 3 PSI output pressure to the feedback bellows (assuming a central fulcrum location, giving a controller gain of 1). The reset (integral) valve has been completely shut off at the start of this thought experiment:



With 0 PSI of air pressure in the reset bellows, it is as though the reset bellows does not exist at all. The mechanism is a simple proportional-only pneumatic controller.

Now, imagine opening up the reset valve just a little bit, so that the output air pressure of 3 PSI begins to slowly fill the reset bellows. As the reset bellows fills with pressurized air, it begins to push down on the left-hand end of the force beam. This forces the baffle closer to the nozzle, causing the output pressure to rise. The regular output bellows has no restrictor valve to impede its filling, and so it *immediately* applies more upward force on the beam with the rising output pressure. With this greater output pressure, the reset bellows has an even greater “final” pressure to achieve, and so its rate of filling continues.

The result of these two bellows' opposing forces (one instantaneous, one time-delayed) is that the lower bellows' pressure must always *lead 3 PSI ahead of the upper bellows' pressure* in order to maintain a pressure difference of 3 PSI necessary to balance force with the PV and SP bellows (whose pressures differ by 3 PSI). This creates a constant 3 PSI differential pressure across the reset restriction valve, resulting in a constant flow of air into the reset bellows at a rate determined by that pressure drop and the opening of the restrictor valve. Eventually this will cause the output pressure to saturate at maximum, but until then the practical importance of this rising pressure action is that the mechanism now exhibits *integral control response* to the constant error between PV and SP:



The greater the difference in pressures between PV and SP (i.e. the greater the *error*), the more pressure drop will develop across the reset restriction valve, causing the reset bellows to fill (or empty, depending on the sign of the error) with compressed air at a faster rate¹⁸, causing the output pressure to change at a faster rate. Thus, we see in this mechanism the defining nature of integral control action: that the magnitude of the error determines the *velocity* of the output signal (its rate of change over time, or $\frac{dm}{dt}$). The rate of integration may be finely adjusted by changing the opening of the restrictor valve, or adjusted in large steps by connecting *capacity tanks* to the reset bellows to greatly increase its effective volume.


¹⁸These restrictor valves are designed to encourage laminar air flow, making the relationship between volumetric flow rate and differential pressure drop *linear* rather than quadratic as it is for large control valves. Thus, a doubling of pressure drop across the restrictor valve results in a doubling of flow rate into (or out of) the reset bellows, and a consequent doubling of integration rate. This is precisely what we desire and expect from a controller with integral action.

29.11.5 Fisher MultiTrol

Front (left) and rear (right) photographs of a real pneumatic controller (a Fisher “MultiTrol” unit) appear here:



The mechanism is remarkably similar to the one used throughout the explanatory discussion, with the important distinction of being *motion-balance* instead of force balance. Proportional and integral control modes are implemented through the actions of four brass bellows pushing as opposing pairs at either end of a beam:





The nozzle may be seen facing down at the middle of the beam, with the center of the beam acting as a baffle. Setpoint control is achieved by moving the position of the nozzle up and down with respect to the beam. A setpoint dial (labeled “Increase Output Pressure”) turns a cam which moves the nozzle closer to or farther away from the beam. This being a motion-balance system, an offset in nozzle position equates to a biasing of the output signal, causing the controller to seek a new process variable value.


Instead of altering the position of a fulcrum to alter the gain (proportional band) of this controller, gain control is effected through the use of a “pressure divider” valve proportioning the amount of output air pressure sent to the feedback bellows. Integral rate control is implemented exactly the same way as in the hypothetical controller mechanism illustrated in the discussion: by adjusting a valve restricting air flow to and from the reset bellows. Both valves are actuated by rotary knobs with calibrated scales. The reset knob is actually calibrated in units of minutes per repeat, while the proportional band knob is labeled with a scale of arbitrary numbers:



Selection of direct versus reverse action is accomplished in the same way as selection between proportional and snap-action (on-off) control: by movable manifolds re-directing air pressure signals to different bellows in the mechanism. The direct/reverse manifold appears in the left-hand photograph (the letter “D” stands for *direct* action) while the proportional/snap manifold appears in the right-hand photograph (the letter “P” stands for *proportional* control):



Either setting is made by removing the screw holding the manifold plate to the controller body, rotating the plate one-quarter turn, and re-attaching. The following photograph shows one of the manifold plates removed and turned upside-down for inspection of the air passages:





The two quarter-circumference slots seen in the manifold plate connect adjacent air ports together. Rotating the plate 90 degrees connects the four air ports together as two different pairs.

29.11.6 Foxboro model 43AP

The Fisher MultiTrol pneumatic controller is a very simple device, intended for field-mounting near the pneumatic transmitter and control valve to form a control loop for non-precision applications. A more sophisticated field-mounted pneumatic controller is the Foxboro model 43AP, sporting actual PV and SP indicating pointers, plus more precise tuning controls. The following photographs show one of these controllers, with the access door closed (left) and open (right):



At the heart of this controller is a motion-balance “pneumatic control unit” mechanism. A dial for setting proportional band (and direct/reverse action) appears on the front of the mechanism:




Note the simple way in which direct and reverse actions are described on this dial: either *increasing measurement, decreasing output* (reverse action) or *increasing measurement, increasing output* (direct action).

29.11.7 Foxboro model 130

Foxboro also manufactured panel-mounted pneumatic controllers, the model 130 series, for larger-scale applications where multiple controllers needed to be located in one compact space. A bank of four Foxboro model 130 pneumatic controllers appears in the next photograph:





Each controller may be partially removed (slid out) from its slot in the rack, the P, I, and D settings adjustable on the left side panel with a screwdriver:



With the side panel removed, the entire mechanism is open to viewing:



The heart of the model 130 controller is a four-bellows force-balance mechanism, identical in principle to the hypothetical force-balance PID controller mechanism used throughout the explanatory discussion. Instead of the four bellows acting against a straight beam, however, these bellows push against a circular disk:




A nozzle (shown in the next photograph) detects if the disk is out of position (unbalanced), sending a back-pressure signal to an amplifying relay which then drives the feedback bellows:




The disk rocks along an axis established by a movable bar. As this bar is rotated at different angles relative to the face of the disk, the fulcrum shifts with respect to the four bellows, providing a simple and effective gain adjustment:

Gain is equal to 1 (PB = 100%)




Gain exceeds 1 (PB < 100%)



If the moment arms (lever lengths) between the input (PV and SP) bellows and the feedback bellows are equal, both sets of bellows will have equal leverage, and the gain will be one (a proportional band setting of 100%). However, if the fulcrum bar is rotated to give the input bellows more leverage and the feedback bellows less leverage, the feedback bellows will have to “work harder” (exert more force) to counteract any imbalance of force created by the input (PV and SP) bellows, thus creating a greater gain: more output pressure for the same amount of input pressure.

The fourth (lower-left) bellows acting on the disk provides an optional reset (integral) function. Its moment arm (lever length) of course is always equal to that of the feedback bellows, just as the PV and SP bellows’ moment arm lengths are always equal, being positioned opposite the fulcrum line.

Selection between direct and reverse action works on the exact same principle as in the Fisher MultiTrol controller – by connecting four air ports in one of two paired configurations. A selector (movable with a hex wrench) turns an air signal port “switch” on the bottom of the four-bellows unit, effectively switching the PV and SP bellows:




An interesting characteristic of most pneumatic controllers is modularity of function: it is possible to order a pneumatic controller that is proportional-only (P), proportional plus integral (P+I), or full PID. Since each control mode requires additional components to implement, a P-only pneumatic controller costs less than a P+I pneumatic controller, which in turn costs less than a full PID pneumatic controller. This explains the relative scarcity of full PID pneumatic controllers in industry: why pay for additional functionality if less will suffice for the task at hand?

29.11.8 External reset (integral) feedback

Some pneumatic controllers come equipped with an option for *external reset*: a feature useful in control systems to avoid integral windup if and when the process stops responding to changes in controller output. Instead of receiving a pneumatic signal directly from the output line of the controller, the reset bellows receives its signal through another pneumatic line, connected to a location in the control system where the final *effect* of the output signal (m) is seen. If for some reason the final control element cannot achieve the state called for by the controller, the controller will sense this through the external reset signal, and will cease integration to avoid “wind-up.”

In the following illustration¹⁹, the external reset signal comes from a pneumatic *position transmitter* (ZT) mounted to the sliding stem of the control valve, sending back a 3-15 PSI signal representing valve stem position:



If something happens to the control valve causing it to freeze position when the controller commands it to move – suppose the stem encounters a mechanical “stop” limiting travel, or a piece of solid material jams the valve trim so it cannot close further – the pneumatic pressure signal sent from the position transmitter to the controller’s reset bellows will similarly freeze. After the

¹⁹In case you are wondering, this controller happens to be *reverse-acting* instead of direct. This is of no consequence to the feature of external reset.

pneumatic lag caused by the reset restrictor valve and bellows passes, the reset bellows force will remain fixed. This halts the controller's integral action, which was formerly based on a "race" between the output feedback bellows and the reset bellows, causing the feedback bellows to "lead" the reset bellows pressure by an amount proportional to the error between PV and SP. This "race" caused the output pressure to wind either up or down depending on the sign of the error. Now that the reset bellows pressure is frozen due to the control valve stem position being frozen, however, the "race" comes to an end and the controller exhibits only proportional action. Thus, the dreaded effect of integral windup – where the integral action of a controller continues to act even though the change in output is of no effect on the process – is averted.

29.12 Analog electronic PID controllers

Although analog electronic process controllers are considered a newer technology than pneumatic process controllers, they are actually "more obsolete" than pneumatic controllers. Panel-mounted (inside a control room environment) analog electronic controllers were a great improvement over panel-mounted pneumatic controllers when they were first introduced to industry, but they were superseded by digital controller technology later on. Field-mounted pneumatic controllers were either replaced by panel-mounted electronic controllers (either analog or digital) or left alone. Applications still exist for field-mounted pneumatic controllers, even now at the beginning of the 21st century, but very few applications exist for analog electronic controllers in any location.

Analog electronic controllers enjoy two inherent advantages over digital electronic controllers: greater reliability²⁰ and faster response. However, these advantages have been diminishing as digital control technology has advanced. Today's digital electronic technology is far more reliable than the digital technology available during the heyday of analog electronic controllers. Now that digital controls have achieved very high levels of reliability, the first advantage of analog control is largely academic²¹, leaving only the second advantage for practical consideration. The advantage of faster speed may be fruitful in applications such as motion control, but for most industrial processes even the slowest digital controller is fast enough²². Furthermore, the numerous advantages offered by digital technology (data recording, networking capability, self-diagnostics, flexible configuration, function blocks for implementing different control strategies) severely weaken the relative importance of reliability and speed.

Most analog electronic PID controllers utilize *operational amplifiers* in their designs. It is relatively easy to construct circuits performing amplification (gain), integration, differentiation, summation, and other useful control functions with just a few opamps, resistors, and capacitors.

²⁰The reason for this is the low component count compared to a comparable digital control circuit. For any given technology, a simpler device will tend to be more reliable than a complex device if only due to there being fewer components to fail. This also suggests a third advantage of analog controllers over digital controllers, and that is the possibility of easily designing and constructing your own for some custom application such as a hobby project. A digital controller is not outside the reach of a serious hobbyist to design and build, but it is definitely more challenging due to the requirement of programming expertise in addition to electronic hardware expertise.

²¹It is noteworthy that analog control systems are completely immune from "cyber-attacks" (malicious attempts to foil the integrity of a control system by remote access), due to the simple fact that their algorithms are fixed by physical laws and properties of electronic components rather than by code which may be edited. This new threat constitutes an inherent weakness of digital technology, and has spurred some thinkers in the field to reconsider analog controls for the most critical applications.

²²The real problem with digital controller speed is that the time delay between successive "scans" translates into dead time for the control loop. Dead time is the single greatest impediment to feedback control.

29.12.1 Proportional control action

The basic proportional-only control algorithm follows this formula:

$$m = K_p e + b$$

Where,

m = Controller output

e = Error (difference between PV and SP)

K_p = Proportional gain

b = Bias


The “error” variable (e) is the mathematical difference between process variable and setpoint. If the controller is direct-acting, $e = PV - SP$. If the controller is reverse-acting, $e = SP - PV$. Thus,

$$m = K_p(PV - SP) + b \quad \text{Direct-acting}$$

$$m = K_p(SP - PV) + b \quad \text{Reverse-acting}$$


Mathematical operations such as subtraction, multiplication by a constant, and addition are quite easy to perform using analog electronic (operational amplifier) circuitry. Prior to the advent of reliable digital electronics for industrial applications, it was natural to use analog electronic circuitry to perform proportional control for process control loops.

For example, the subtraction function necessary to calculate error (e) from process variable and setpoint signals may be performed with a three-amplifier “subtractor” circuit:




This particular subtractor circuit calculates error for a reverse-acting controller. As the PV signal increases, the error signal decreases (becomes more negative). It could be modified for direct action simply by swapping the two inputs: SP on top and PV on bottom such that the Output becomes $PV - SP$.

Gain is really nothing more than multiplication by a constant, in this case the constant being K_p . A very simple one-amplifier analog circuit for performing this multiplication is the *inverting*²³ *amplifier* circuit:



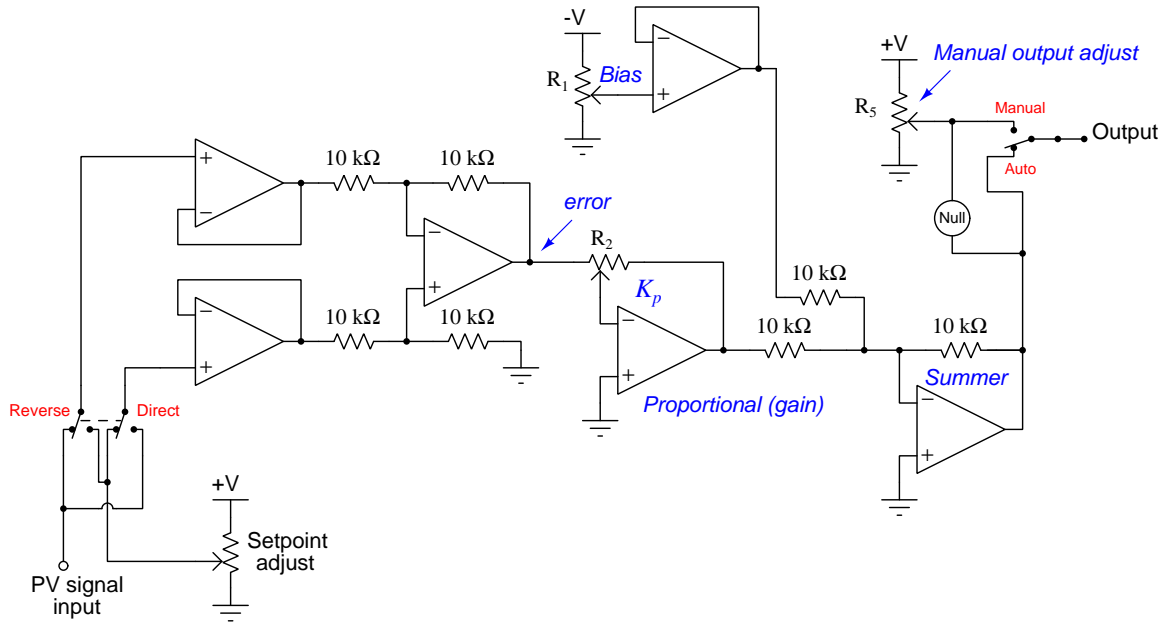
With the potentiometer's wiper in mid-position, the voltage gain of this circuit will be 1 (with an inverted polarity which we shall ignore for now). Moving the wiper toward the left-hand side of the potentiometer increases the circuit's gain past unity, while moving the wiper toward the right-hand side of the potentiometer decreases the gain toward zero.

In order to add the bias (b) term in the proportional control equation, we need an analog circuit capable of summing two voltage signals. This need is nicely met in the *inverting summer* circuit, shown here:



²³This circuit configuration is called "inverting" because the mathematical sign of the output is always opposite that of the input. This sign inversion is not an intentional circuit feature, but rather a consequence of the input signal facing the opamp's inverting input. Non-inverting multiplier circuits also exist, but are more complicated when built to achieve multiplication factors less than one.

Combining all these analog functions together into one circuit, and adding a few extra features such as direct/reverse action selection, bias adjustment, and manual control with a null voltmeter to facilitate bumpless mode transfer, gives us this complete analog electronic proportional controller:



29.12.2 Derivative and integral control actions

Differentiating and integrating live voltage signals with respect to time is quite simple using operational amplifier circuits. Instead of using all resistors in the negative feedback network, we may implement these calculus functions by using a combination of *capacitors* and resistors, exploiting the capacitor's natural derivative relationship between voltage and current:

$$I = C \frac{dV}{dt}$$

Where,


I = Current through the capacitor (amperes)

C = Capacitance of capacitor (farads)

V = Voltage across the capacitor (volts)

$\frac{dV}{dt}$ = Rate-of-change of voltage across the capacitor (volts per second)

If we build an operational amplifier with a resistor providing negative feedback current through a capacitor, we create a *differentiator* circuit where the output voltage is proportional to the rate-of-change of the input voltage:



Since the inverting input of the operational amplifier is held to ground potential by feedback (a “virtual ground”), the capacitor experiences the full input voltage of signal A . So, as A varies over time, the current through that capacitor will directly represent the signal A ’s rate of change over time ($I = C \frac{dA}{dt}$). This current passes through the feedback resistor, creating a voltage drop at the output of the amplifier directly proportional to signal A ’s rate of change over time. Thus, the output voltage of this circuit reflects the input voltage’s instantaneous rate of change, albeit with an inverted polarity. The mathematical term RC is the *time constant* of this circuit. For a differentiator circuit such as this, we typically symbolize its time constant as τ_d (the “derivative” time constant).

For example, if the input voltage to this differentiator circuit were to ramp at a constant rate of +4.3 volts per second (rising) with a resistor value of 10 k Ω and a capacitor value of 33 μ F (i.e. $\tau_d = 0.33$ seconds), the output voltage would be a constant -1.419 volts:

$$V_{out} = -RC \frac{dV_{in}}{dt}$$


$$V_{out} = -(10000 \Omega)(33 \times 10^{-6} \text{ F}) \left(\frac{4.3 \text{ V}}{\text{s}} \right)$$

$$V_{out} = -(0.33 \text{ s}) \left(\frac{4.3 \text{ V}}{\text{s}} \right)$$

$$V_{out} = -1.419 \text{ V}$$

Recall that the purpose of derivative action in a PID controller is to react to sudden changes in either the error (e) or the process variable (PV). This circuit fulfills that function, by generating an output proportional to the input voltage’s rate of change.

If we simply swap²⁴ the locations of the resistor and capacitor in the feedback network of this operational amplifier circuit, we create an *integrator* circuit where the output voltage rate-of-change is proportional to the input voltage:



This integrator circuit provides the exact inverse function of the differentiator. Rather than a changing input signal generating an output signal proportional to the input's rate of change, an input signal in this circuit controls the rate at which the output signal changes.

The way it works is by acting as a *current source*, pumping current into the capacitor at a value determined by the input voltage and the resistor value. Just as in the previous (differentiator) circuit where the inverting terminal of the amplifier was a “virtual ground” point, the input voltage in this circuit is impressed across the resistor R . This creates a current which must go through capacitor C on its way either to or from the amplifier's output terminal. As we have seen in the capacitor's equation ($I = C \frac{dV}{dt}$), a current forced through a capacitor causes the capacitor's voltage to change over time. This changing voltage becomes the output signal of the integrator circuit. As in the case of the differentiator circuit, the mathematical term RC is the *time constant* of this circuit as well. Being an integrator, we customarily represent this “integral” time constant as τ_i .

Any amount of change in output voltage (ΔV_{out}) occurring between some initial time (t_0) and a finishing time (t_f) may be calculated by the following integral:

$$\Delta V_{out} = -\frac{1}{RC} \int_{t_0}^{t_f} V_{in} dt$$

If we wish to know the absolute output voltage at the end of that time interval, all we need to do is add the circuit's initial output voltage (V_0 , i.e. the voltage stored in the capacitor at the initial time t_0) to the calculated change:

$$V_{out} = -\frac{1}{RC} \int_{t_0}^{t_f} V_{in} dt + V_0$$

²⁴This inversion of function caused by the swapping of input and feedback components in an operational amplifier circuit points to a fundamental principle of negative feedback networks: namely, that placing a mathematical element within the feedback loop causes the amplifier to exhibit the inverse of that element's intrinsic function. This is why voltage dividers placed within the feedback loop cause an opamp to have a multiplicative gain (division \rightarrow multiplication). A circuit element exhibiting a logarithmic response, when placed within a negative feedback loop, will cause the amplifier to exhibit an exponential response (logarithm \rightarrow exponent). Here, an element having a time-differentiating response, when placed inside the feedback loop, causes the amplifier to time-integrate (differentiation \rightarrow integration). Since the opamp's output voltage must assume any value possible to maintain (nearly) zero differential voltage at the input terminals, placing a mathematical function in the feedback loop forces the output to assume the inverse of that function in order to “cancel out” its effects and achieve balance at the input terminals.

For example, if we were to input a constant DC voltage of +1.7 volts to this circuit with a resistor value of 81 k Ω and a capacitor value of 47 μ F (i.e. $\tau_i = 3.807$ seconds), the output voltage would ramp at a constant rate of -0.447 volts per second²⁵. If the output voltage were to begin at -3.0 volts and be allowed to ramp for exactly 12 seconds at this rate, it would reach a value of -8.359 volts at the conclusion of that time interval:

$$V_{out} = -\frac{1}{RC} \int_{t_0}^{t_f} V_{in} dt + V_0$$

$$V_{out} = -\left(\frac{1}{(81000 \Omega)(47 \times 10^{-6} \text{ F})}\right) \left(\int_0^{12} 1.7 \text{ V } dt\right) - 3 \text{ V}$$

$$V_{out} = -\left(\frac{1}{3.807 \text{ s}}\right) (20.4 \text{ V} \cdot \text{s}) - 3 \text{ V}$$

$$V_{out} = -5.359 \text{ V} - 3 \text{ V}$$

$$V_{out} = -8.359 \text{ V}$$


If, after ramping for some amount of time, the input voltage of this integrator circuit is brought to zero, the integrating action will cease. The circuit's output will simply hold at its last value until another non-zero input signal voltage appears.

Recall that the purpose of integral action in a PID controller is to eliminate offset between process variable and setpoint by calculating the error-time product (how far PV deviates from SP, and for how long). This circuit will fulfill that function if the input voltage is the error signal, and the output voltage contributes to the output signal of the controller.

²⁵If this is not apparent, imagine a scenario where the +1.7 volt input existed for precisely one second's worth of time. However much the output voltage ramps in that amount of time must therefore be its rate of change in volts per second (assuming a linear ramp). Since we know the area accumulated under a constant value of 1.7 (high) over a time of 1 second (wide) must be 1.7 volt-seconds, and τ_i is equal to 3.807 seconds, the integrator circuit's output voltage must ramp 0.447 volts during that interval of time. If the input voltage is positive and we know this is an inverting opamp circuit, the direction of the output voltage's ramping must be negative, thus a ramping rate of *-0.447 volts per second*.

29.12.3 Full-PID circuit design

The following schematic diagram shows a full PID controller implemented using eight operational amplifiers, designed to input and output voltage signals representing PV, SP, and Output²⁶:



It is somewhat stunning to realize that such a controller, fully capable of controlling many industrial process types, may be constructed using only two integrated circuit “chips” (two “quad” operational amplifiers) and a handful of passive electronic components. The only significant engineering challenge in this simple circuit design is achieving slow enough time constants (in the

²⁶The two input terminals shown, $Input_{(+)}$ and $Input_{(-)}$ are used as PV and SP signal inputs, the correlation of each depending on whether one desires direct or reverse controller action.


range of minutes rather than seconds) in the integrator and differentiator functions using non-polarized capacitors²⁷.

This controller implements the so-called *ideal* PID algorithm, with the proportional (gain) value distributing to the integral and derivative terms:

$$m = K_p \left(e + \frac{1}{\tau_i} \int e dt + \tau_d \frac{de}{dt} \right) \quad \text{Ideal PID equation}$$


We may determine this from the schematic diagram by noting that the I and D functions each receive their input signals from the output of the proportional amplifier (the one with the R_{prop} potentiometer). Adjusting R_{prop} affects not only the controller's proportional gain, but also the sensitivity of τ_i and τ_d .

An actual implementation of this PID controller in printed circuit board form appears here:



²⁷This particular design has integral and derivative time value limits of 10 seconds, maximum. These relatively “quick” tuning values are the result of having to use non-polarized capacitors in the integrator and differentiator stages. The practical limits of cost and size restrict the maximum value of on-board capacitance to around 10 μF each.

It is possible to construct an analog PID controller with fewer components. An example is shown here:



As you can see, a *single* operational amplifier does all the work of calculating proportional, integral, *and* derivative responses. The first two amplifiers do nothing but buffer the input signals and calculate error (PV – SP, or SP – PV, depending on the direction of action).

One of the consequences of consolidating all three control terms in a single amplifier is that those control terms interact with each other. The mathematical expression of this control action is shown here, called the *series* or *interacting* PID equation:

$$m = K_p \left[\left(\frac{\tau_d}{\tau_i} + 1 \right) e + \frac{1}{\tau_i} \int e \, dt + \tau_d \frac{de}{dt} \right] \quad \text{Series or Interacting PID equation}$$

Not only does a change in gain (K_p) alter the relative responses of integral and derivative in the series equation (as it also does in the ideal equation), but changes in either integral or derivative time constants also have an effect on proportional response! This is especially noticeable when the integral time constant is set to some very small value, which is typically the case on fast-responding, self-regulating processes such as liquid flow or liquid pressure control.

It should be apparent that an analog controller implementing the series equation is simpler in construction than one implementing either the parallel or ideal PID equation. This also happens to be true for pneumatic PID controller mechanisms: the simplest analog controller designs all implement the series PID equation²⁸.

²⁸An interesting example of engineering tradition is found in electronic PID controller designs. While it is not too terribly difficult to build an analog electronic controller implementing either the parallel or ideal PID equation (just a few more parts are needed), it is quite challenging to do the same in a pneumatic mechanism. When analog electronic controllers were first introduced to industry, they were often destined to replace old pneumatic controllers. In order to ease the transition from pneumatic to electronic control, manufacturers built their new electronic controllers to behave exactly the same as the old pneumatic controllers they would be replacing. The same legacy followed the advent of digital electronic controllers: many digital controllers were programmed to behave in the same manner as the old pneumatic controllers, for the sake of operational familiarity, not because it was easier to design a digital controller that way.

29.12.4 Single-loop analog controllers

One popular analog electronic controller was the Foxboro model 62H, shown in the following photographs. Like the model 130 pneumatic controller, this electronic controller was designed to fit into a rack next to several other controllers. Tuning parameters were adjustable by moving potentiometer knobs under a side-panel accessible by partially removing the controller from its rack:



The Fisher corporation manufactured a series of analog electronic controllers called the AC², which were similar in construction to the Foxboro model 62H, but very narrow in width so that many could be fit into a compact panel space. Here we see a pair of Fisher AC² controllers mounted side-by-side in the same rack, used to control liquid level in a pulping process:



Like the pneumatic panel-mounted controllers preceding, and digital panel-mount controllers to follow, the tuning parameters for a panel-mounted analog electronic controller were typically

accessed on the controller's side. The controller could be slid partially out of the panel to reveal the P, I, and D adjustment knobs (as well as direct/reverse action switches and other configuration controls).

Indicators on the front of an analog electronic controller served to display the process variable (PV), setpoint (SP), and manipulated variable (MV, or output) for operator information. Many analog electronic controllers did not have separate meter indications for PV and SP, but rather used a single meter movement to display the *error signal*, or difference between PV and SP. On the Foxboro model 62H, a hand-adjustable knob provided both indication and control over SP, while a small edge-reading meter movement displayed the error. A negative meter indication showed that the PV was below setpoint, and a positive meter indication showed that the PV was above setpoint.

The Fisher AC² analog electronic controller used the same basic technique, cleverly applied in such a way that the PV was displayed in real engineering units. The setpoint adjustment was a large wheel, mounted so the edge faced the operator. Along the circumference of this wheel was a scale showing the process variable range, from the LRV at one extreme of the wheel's travel to the URV at the other extreme of the wheel's travel. The actual setpoint value was the middle of the wheel from the operator's view of the wheel edge. A single meter movement needle traced an arc along the circumference of the wheel along this same viewable range. If the error was zero ($PV = SP$), the needle would be positioned in the middle of this viewing range, pointed at the same value along the scale as the setpoint. If the error was positive, the needle would rise up to point to a larger (higher) value on the scale, and if the error was negative the needle would point to a smaller (lower) value on the scale. For any fixed value of PV, this error needle would therefore move in exact step with the wheel as it was rotated by the operator's hand. Thus, a single adjustment and a single meter movement displayed both SP and PV in very clear and unambiguous form.

Taylor manufactured a line of analog panel-mounted controllers that worked much the same way, with the SP adjustment being a graduated tape reeled to and fro by the SP adjustment knob. The middle of the viewable section of tape (as seen through a plastic window) was the setpoint value, and a single meter movement needle pointed to the PV value as a function of error. If the error happened to be zero ($PV = SP$), the needle would point to the middle of this viewable section of tape, which was the SP value.

Another popular panel-mounted analog electronic controller was the Moore Syncro, which featured plug-in modules for implementing different control algorithms (different PID equations, nonlinear signal conditioning, etc.). These plug-in function modules were a hardware precursor to the software "function blocks" appearing in later generations of digital controllers: a simple way of organizing controller functionality so that technicians unfamiliar with computer programming could easily configure a controller to do different types of control functions. Later models of the Syncro featured fluorescent bargraph displays of PV and SP for easy viewing in low-light conditions.

Analog single-loop controllers are largely a thing of the past, with the exception of some low-cost or specialty applications. An example of the former is shown here, a simple analog temperature controller small enough to fit in the palm of my hand:



This particular controller happened to be part of a sulfur dioxide analyzer system, controlling the internal temperature of a gas regulator panel to prevent vapors in the sample stream from condensing in low spots of the tubing and regulator system. The accuracy of such a temperature control application was not critical – if temperature was regulated to ± 5 degrees Fahrenheit it would be more than adequate. This is an application where an analog controller makes perfect sense: it is very compact, simple, extremely reliable, and inexpensive. None of the features associated with digital PID controllers (programmability, networking, precision) would have any merit in this application.

29.12.5 Multi-loop analog control systems

In contrast to single-loop analog controllers, *multi-loop* systems control dozens or even hundreds of process loops at a time. Prior to the advent of reliable digital technology, the only electronic process control systems capable of handling the numerous loops within large industrial installations such as power generating plants, oil refineries, and chemical processing facilities were analog systems, and several manufacturers produced multi-loop analog systems just for these large-scale control applications.

One of the most technologically advanced analog electronic products manufactured for industrial control applications was the Foxboro SPEC 200 system²⁹. Although the SPEC 200 system used panel-mounted indicators, recorders, and other interface components resembling panel-mounted control systems, the actual control functions were implemented in a separate equipment rack which Foxboro called a *nest*³⁰. Printed circuit boards plugged into each “nest” provided all the control functions (PID controllers, alarm units, integrators, signal selectors, etc.) necessary, with analog signal wires connecting the various functions together with panel-mounted displays and with field instruments to form a working system.

Analog field instrument signals (4-20 mA, or in some cases 10-50 mA) were all converted to a 0-10 VDC range for signal processing within the SPEC 200 nest. Operational amplifiers (mostly the model LM301) formed the “building blocks” of the control functions, with a ± 15 VDC power supply providing DC power for everything to operate.

²⁹Although the SPEC 200 system – like most analog electronic control systems – is considered “mature” (Foxboro officially declared the SPEC 200 and SPEC 200 Micro systems as such in March 2007), working installations may still be found at the time of this writing (2010). A report published by the Electric Power Research Institute (see References at the end of this chapter) in 2001 documents a SPEC 200 analog control system installed in a nuclear power plant in the United States as recently as 1992, and another as recently as 2001 in a Korean nuclear power plant.

³⁰Foxboro provided the option of a self-contained, panel-mounted SPEC 200 controller unit with all electronics contained in a single module, but the split architecture of the display/nest areas was preferred for large installations where many dozens of loops (especially cascade, feedforward, ratio, and other multi-component control strategies) would be serviced by the same system.


The following photographs show a model 2AX+A4 proportional-integral (P+I) controller card for a SPEC 200 system inserted into a metal frame (called a “module” by Foxboro). This module was designed to fit into a slot in a SPEC 200 “nest” where it would reside alongside many other similar cards, each card performing its own control function:



Tuning and alarm adjustments may be seen in the right-hand photograph. This particular controller is set to a proportional band value of approximately 170%, and an integral time constant of just over 0.01 minutes per repeat. A two-position rotary switch near the bottom of the card selected either reverse (“Dec”) or direct (“Inc”) control action.


The array of copper pins at the top of the module form the male half of a cable connector, providing connection between the control card and the front-panel instrument accessible to operations personnel. Since the tuning controls appear on the face of this controller card (making it a “card tuned” controller), they were not accessible to operators but rather only to the technical personnel with access to the nest area. Other versions of controller cards (“control station tuned”) had blank places where the P and I potentiometer adjustments appear on this model, with tuning adjustments provided on the panel-mounted instrument displays for easier access to operators.

The set of ten screw terminals at the bottom of the module provided connection points for the input and output voltage signals. The following list gives the general descriptions of each terminal pair, with the descriptions for this particular P + I controller written in *italic type*:




- Terminals (1+) and (1−): Input signal #1 (*Process variable input*)
- Terminals (2+) and (2−): Output signal #1 (*Manipulated variable output*)
- Terminals (3+) and (3−): Input #2, Output #4, or Option #1 (*Remote setpoint*)
- Terminals (4+) and (4−): Input #3, Output #3, or Option #2 (*Optional alarm*)
- Terminals (5+) and (5−): Input #4, Output #2, or Option #3 (*Optional 24 VAC*)

A photograph of the printed circuit board (card) removed from the metal module clearly shows the analog electronic components:



The following photograph shows a set of functioning Foxboro SPEC 200 controller modules residing in a “nest,” used to control a flow loop and a pressure loop in a wood pulping process. Both active controller modules are P + I units. An unused PID module resides just to the left of the flow controller module:



Foxboro engineers went to great lengths in their design process to maximize reliability of the SPEC 200 system, already an inherently reliable technology by virtue of its simple, analog nature. As a result, the reliability of SPEC 200 control systems is the stuff of legend³¹.

29.13 Digital PID controllers

The vast majority of PID controllers in service today are digital in nature. Microprocessors executing PID algorithms provide many advantages over any form of analog PID control (pneumatic or electronic), not the least of which being the ability to network with personal computer workstations and other controllers over wired or wireless (radio) networks.

³¹I once encountered an engineer who joked that the number “200” in “SPEC 200” represented the number of years the system was designed to continuously operate. At another facility, I encountered instrument technicians who were a bit afraid of a SPEC 200 system running a section of their plant: the system had *never suffered a failure of any kind* since it was installed decades ago, and as a result no one in the shop had any experience troubleshooting it. As it turns out, the entire facility was eventually shut down and sold, with the SPEC 200 nest running faithfully until the day its power was turned off! The functioning SPEC 200 controllers shown in the photograph were in continuous use at British Columbia Institute of Technology at the time of the photograph, taken in December of 2014.

29.13.1 Stand-alone digital controllers


If the internal components of a panel-mounted pneumatic or analog electronic controller (such as the Foxboro models 130 or 62, respectively) were completely removed and replaced by all-digital electronic componentry, the result would be a *stand-alone digital PID controller*. From the outside, such a digital controller looks very similar its technological ancestors, but its capabilities are far greater.

An example of a popular panel-mounted digital controller is the Siemens model 353 (formerly the Moore Products model 353):




This particular controller, like many high-end digital controllers and larger digital control systems, is programmed in a function block language. Each function block in the controller is a software subroutine performing a specific function on input signals, generating at least one output signal. Each function block has a set of configuration parameters telling it how to behave. For example, the PID function block in a digital controller would have parameters specifying direct or reverse action, gain (K_p), integral time constant (τ_i), derivative time constant (τ_d), output limits, etc.

Even the “stock” configuration for simple, single-loop PID control is a collection of function blocks linked together:



The beauty of function block programming is that the same blocks may be easily re-linked to implement custom control strategies. Take for instance the following function block program written for a Siemens model 353 controller to provide a pulse-width-modulation (PWM, or time-proportioned) output signal instead of the customary 4-20 mA DC analog output signal. The application is for an electric oven temperature control system, where the oven’s heating element could only be turned on and off fully rather than continuously varied:



In order to specify links between function blocks, each of the used lettered block inputs is mapped to the output channel of another block. In the case of the time-proportioned function block program, for example, the “P” (process variable) input of the PID function block is set to get its signal from the “01” output channel of the AIN1 (analog input 1) function block. The “TV” (tracking value) input of the SETPT (setpoint) function block is also set to the “01” output channel of the AIN1 function block, so that the setpoint value generator has access to the process variable value in order to implement setpoint tracking. Any function block output may drive an unlimited number of function block inputs (fan-out), but each function block input may receive a signal from *only one* function block output. This is a rule followed within all function block languages to prevent multiple block output signals from conflicting (attempting to insert different signal values into the same input).

In the Siemens controllers, function block programming may be done by entering configuration data using the front-panel keypad, or by using graphical software running on a personal computer networked with the controller.

For applications not requiring so much capability, and/or requiring a smaller form factor, other panel-mounted digital controllers exist. The Honeywell model UDC3000 is a popular example of a 1/4 DIN (96 mm × 96 mm) size digital controller:




Even smaller panel-mounted controllers are produced by a wide array of manufacturers for applications requiring minimum functionality: 1/8 DIN (96 mm × 48 mm), 1/16 DIN (48 mm × 48 mm), and even 1/32 DIN (48 mm × 24 mm) sizes are available.

One of the advantageous capabilities of modern stand-alone controllers is the ability to exchange data over digital networks. This provides operations and maintenance personnel alike the ability to remotely monitor and even control (adjust setpoints, switch modes, change tuning parameters, etc.) the process controller from a computer workstation. The Siemens model 353 controller (with appropriate options) has the ability to digitally network over Ethernet, a very common and robust digital network standard. The following photographs show three such controllers connected to a network through a common 4-port Ethernet “hub” device:




Special software (in this case, Siemens *Procidia*) running on a computer workstation connected to the same Ethernet network acquires data from and sends data to the networked controllers. Screenshots of this software show typical displays allowing complete control over the function of the process controllers:



29.13.2 Direct digital control (DDC)

A microprocessor operating at sufficient clock speed is able to execute more than one PID control algorithm for a process loop, by “time-sharing” its calculating power: devoting slices of time to the evaluation of each PID equation in rapid succession. This not only makes multiple-loop digital control possible for a single microprocessor, but also makes it very attractive given the microprocessor’s natural ability to manage data archival, transfer, and networking. A single computer is able to execute PID control for multiple loops, and also make that loop control data accessible between loops (for purposes of cascade, ratio, feedforward, and other control strategies) and accessible on networks for human operators and technicians to easily access.

Such *direct digital* control (DDC) has been applied with great success to the problem of building automation, where temperature and humidity controls for large structures benefit from large-scale data integration. The following photograph shows a Siemens APOGEE building automation system with multiple I/O (input/output) cards providing interface between analog instrument signals and the microprocessor’s digital functions:




A close-up view of an APOGEE processor shows the device handling all mathematical calculations for the PID control:



Other than a few LEDs, there is no visual indication in this panel of what the system is doing at any particular time. Operators, engineers, and technicians alike must use software running on a networked personal computer to access data in this control system.

An example of the HMI (Human-Machine Interface) software one might see used in conjunction with a DDC controller is shown here, also from a Siemens APOGEE building control system:



This particular screenshot shows monitored and controlled variables for a heat exchanger (“heat wheel”) used to exchange heat between outgoing and incoming air for the building.


A smaller-scale example of a DDC system is the Delta model DSC-1280 controller, an example shown in the following photograph:



This system does not have plug-in I/O cards like the Siemens APOGEE, but instead is monolithic in design, with all inputs and outputs part of one large “motherboard” PCB. The model DSC-1280 controller has 12 input channels and 8 output channels (hence the model number “1280”). An Ethernet cable (RJ-45 plug) is seen in the upper-left corner of this unit, through which a remotely-located personal computer communicates with the DDC using a high-level protocol called BACnet. In many ways, BACnet is similar to Modbus, residing at layer 7 of the OSI Reference Model (the so-called *Application Layer*), unconcerned with the details of data communication at the Physical or Data Link layers. This means, like Modbus, BACnet commands may be sent and received over a variety of lower-level network standards, with Ethernet being the preferred³² option at the time of this writing.

³²Thanks to the explosion of network growth accompanying personal computers in the workplace, Ethernet is ubiquitous. The relatively high speed and low cost of Ethernet communications equipment makes it an attractive network standard over which a great many high-level industrial protocols communicate.


Another example of a small-scale DDC is this Distech model ECP-410 unit:




This controller uses *LonWorks* as its communication protocol rather than BACnet. Like BACnet, LonWorks is an upper-layer protocol and may be transported over Ethernet as well as simpler serial communication formats. The model ECP-410 DDC controller has a two-wire connection for its LonWorks network.

Programming of DDC controllers ranges from text-based languages (similar to BASIC) to function-block programming. The Delta DSC-1280 is an example of a controller programmed in text, while the Distech ECP-410 supports function blocks in addition to text-based programming.

The application-specific nature of DDC (environmental controls for the interior of large buildings and other facilities) lends itself to controller units pre-programmed to perform well-defined tasks rather than general-purpose controllers designed to be user-programmable for any task. An example of such a “fixed” program controller is the Distech model EC-RTU-L, designed to control a “rooftop unit” for air handling:



A more common application of industrial DDC is the use of programmable logic controllers (PLCs) to control multiple loops. PLCs were originally invented for on/off (discrete) process control functions, but have subsequently grown in speed and capability to execute analog PID control functions as well. This next photograph shows an Allen-Bradley (Rockwell) ControlLogix PLC used to control the operation of a gas turbine engine. The PLC may be seen in the upper-left corner of the enclosure, with the rest of the enclosure devoted to terminal blocks and accessory components:



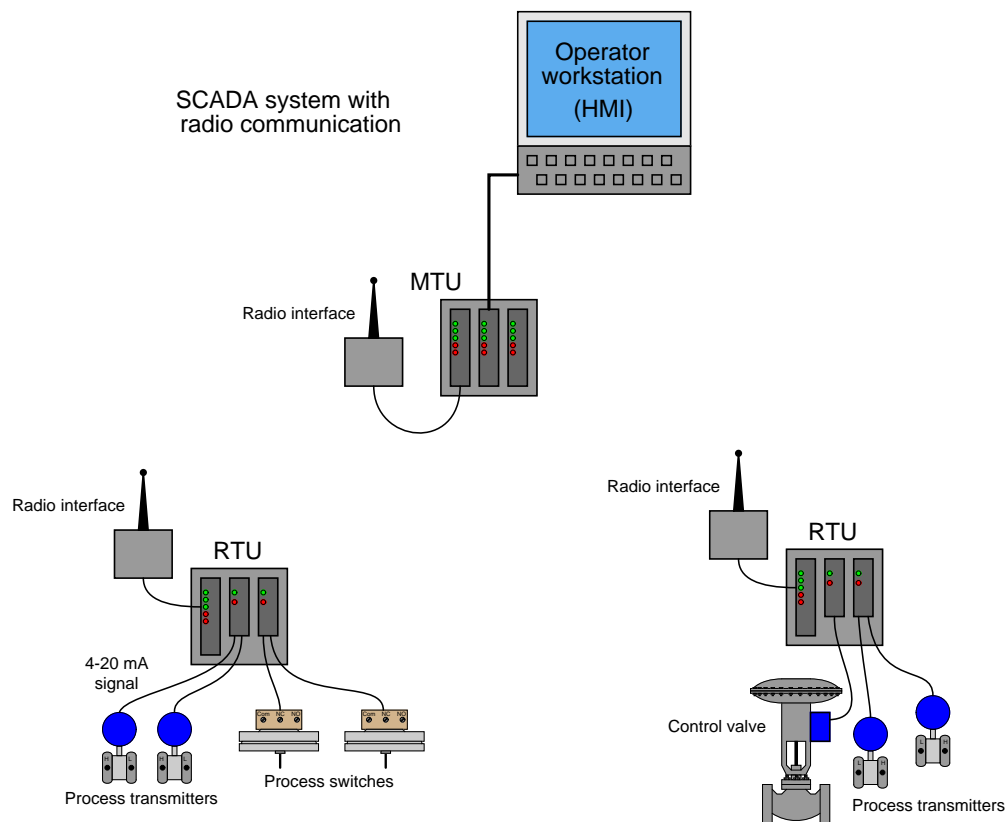
A strong advantage of using PLCs for analog loop control is the ability to easily integrate discrete controls with the analog controls. It is quite easy, for example, to coordinate the sequential start-up and shut-down functions necessary for intermittent operation with the analog PID controls necessary for continuous operation, all within one programmable logic controller. It should be noted, however, that many early PLC implementations of PID algorithms were crude at best, lacking the finesse of stand-alone PID controllers. Even some modern PLC analog functions are mediocre³³ at the time of this writing (2008).

³³An aspect common to many PLC implementations of PID control is the use of the “parallel” PID algorithm instead of the superior “ISA” or “non-interacting” algorithm. The choice of algorithm may have a profound effect on tuning, and on tuning procedures, especially when tuning parameters must be re-adjusted to accommodate changes in transmitter range.

29.13.3 SCADA and telemetry systems

A similar control system architecture to Direct Digital Control (DDC) – assigning a single microprocessor to the task of managing multiple control functions, with digital communication between the microprocessor units – is used for the management of systems which are by their very nature spread over wide geographical regions. Such systems are generally referred to as *SCADA*, which is an acronym standing for *Supervisory Control And Data Acquisition*.

The typical SCADA system consists of multiple *Remote Terminal Unit* (RTU) devices connected to process transmitters and final control elements, implementing basic control functions such as motor start/stop and PID loop control. These RTU devices communicate digitally to a *Master Terminal Unit* (MTU) device at a central location where human operators may monitor the process and issue commands.



A photograph of an RTU “rack” operating at a large electric power substation is shown here:



Some RTU hardware, such as the substation monitoring system shown above, is custom-manufactured for the application. Other RTU hardware is more general in purpose, intended for the monitoring and control of natural gas and oil production wells, but applicable to other applications as well.

The Fisher ROC 800 – shown in the photograph below – is an example of an RTU designed to operate with a minimum of electrical power, so that a single solar panel and battery will be sufficient for year-round operation in remote environments. The particular unit shown is installed in a natural gas metering station, where it monitors gas pressure, temperature, and flow rate, and also controls the injection of an “odorizing” compound into the gas to give it a bad smell:




Standard programmable logic controllers (PLCs) are ideal candidates for use as RTU devices. Modern PLCs have all the I/O, networking, and control algorithm capability necessary to function as remote terminal units. Commercially available Human-Machine Interface (HMI) software allowing personal computers to display PLC variable values potentially turns every PC into a Master Terminal Unit (MTU) where operators can view process variables, change setpoints, and issue other commands for controlling the process.

A photograph of such HMI software used to monitor a SCADA system for a set of natural gas compressors is shown here:



Another photograph of a similar system used to monitor and control drinking water reservoirs for a city is shown here:



A concept closely related to SCADA is that of *telemetry*, the word literally meaning “distance measuring” (i.e. measuring something over a distance). The acronym SCADA, by containing the word “control,” implies two-way communication (measurement and control) between the master location and the remote location. In applications where the flow of information is strictly one-way (simplex) from the remote location to the master location, “telemetry” is a more apt description.

Telemetry systems find wide application in scientific research. Seismographs, river and stream flowmeters, weather stations, and other remotely-located measurement instruments connected (usually by radio links) to some centralized data collection center are all examples of telemetry. Any industrial measurement (-only) application spanning a large distance could likewise be classified as a telemetry system, although you will sometimes find the term “SCADA” applied even when the communication is simplex in nature.


29.13.4 Distributed Control Systems (DCS)

A radically new concept appeared in the world of industrial control in the mid-1970's: the notion of *distributed* digital control. Direct digital control during that era³⁴ suffered a substantial problem: the potential for catastrophic failure if the single digital computer executing *multiple* PID control functions were to ever halt. Digital control brings many advantages, but it isn't worth the risk if the entire operation will shut down (or catastrophically fail!) following a hardware or software failure within that one computer.

Distributed control directly addressed this concern by having multiple control computers – each one responsible for only a handful of PID loops – distributed throughout the facility and networked together to share information with each other and with operator display consoles. With individual process control “nodes” scattered throughout the campus, each one dedicated to controlling just a few loops, there would be less concentration of liability as there would be with a single-computer DDC system. Such distribution of computing hardware also shortened the analog signal wiring, because now the hundreds or thousands of analog field instrument cables only had to reach as far as the distributed nodes, not all the way to a centralized control room. Only the networking cable had to reach that far, representing a drastic reduction in wiring needs. Furthermore, distributed control introduced the concept of *redundancy* to industrial control systems: where digital signal acquisition and processing hardware units were equipped with “spare” units designed to automatically take over all critical functions in the event of a primary failure.

³⁴Modern DDC systems of the type used for building automation (heating, cooling, security, etc.) almost always consist of networked control nodes, each node tasked with monitoring and control of a limited area. The same may be said for modern PLC technology, which not only exhibits advanced networking capability (fieldbus I/O networks, Ethernet, Modbus, wireless communications), but is often also capable of redundancy in both processing and I/O. As technology becomes more sophisticated, the distinction between a DDC (or a networked PLC system) and a DCS becomes more ambiguous.

The following illustration shows a typical distributed control system (DCS) architecture:




Each “rack” contains a microprocessor to implement all necessary control functions, with individual I/O (input/output) “cards” for converting analog field instrument signals into digital format, and vice-versa. Redundant processors, redundant network cables, and even redundant I/O cards address the possibility of component failure. DCS processors are usually programmed to perform routine self-checks³⁵ on redundant system components to ensure availability of the spare components in the event of a failure.

If there ever was a total failure in one of the “control racks” where the redundancy proved insufficient for the fault(s), the only PID loops faulted will be those resident in that rack, not any of the other loops throughout the system. Likewise, if ever the network cables become severed or otherwise faulted, only the information flow between those two points will suffer; the rest of the system will continue to communicate data normally. Thus, one of the “hallmark” features of a DCS is its tolerance to serious faults: even in the event of severe hardware or software faults, the impact to process control is minimized by design.

³⁵An example of such a self-check is scheduled switching of the networks: if the system has been operating on network cable “A” for the past four hours, it might switch to cable “B” for the next four hours, then back again after another four hours to continually ensure both cables are functioning properly.

One of the very first distributed control systems in the world was the Honeywell TDC2000 system³⁶, introduced in 1975. By today's standards, the technology was crude³⁷, but the concept was revolutionary.

Each rack (called a “box” by Honeywell) consisted of an aluminum frame holding several large printed circuit boards with card-edge connectors. A “basic controller” box appears in the left-hand photograph. The right-hand photograph shows the termination board where the field wiring (4-20 mA) connections were made. A thick cable connected each termination board to its respective controller box:



Controller redundancy in the TDC2000 DCS took the form of a “spare” controller box serving as a backup for up to eight other controller boxes. Thick cables routed all analog signals to this spare controller, so that it would have access to them in the event it needed to take over for a failed controller. The spare controller would become active on the event of *any* fault in any of the other controllers, including failures in the I/O cards. Thus, this redundancy system provided for processor failures as well as I/O failures. All TDC2000 controllers communicated digitally by means of a dual coaxial cable network known as the “Data Hiway.” The dual cables provided redundancy in network communications.

³⁶To be fair, the Yokogawa Electric Corporation of Japan introduced their CENTUM distributed control system the same year as Honeywell. Unfortunately, while I have personal experience maintaining and using the Honeywell TDC2000 system, I have zero personal experience with the Yokogawa CENTUM system, and neither have I been able to obtain technical documentation for the original incarnation of this DCS (Yokogawa's latest DCS offering goes by the same name). Consequently, I can do little in this chapter but mention its existence, despite the fact that it deserves just as much recognition as the Honeywell TDC2000 system.

³⁷Just to give some perspective, the original TDC2000 system used whole-board processors rather than microprocessor chips, and magnetic core memory rather than static or dynamic RAM circuits! Communication between controller nodes and operator stations occurred over thick coaxial cables, implementing master/slave arbitration with a separate device (a “Hiway Traffic Director” or HTD) coordinating all communications between nodes. Like Bob Metcalfe's original version of Ethernet, these coaxial cables were terminated at their end-points by termination resistors, with coaxial “tee” connectors providing branch points for multiple nodes to connect along the network.

A typical TDC2000 operator workstation appears in the next photograph:



Over the years following its 1975 introduction, the Honeywell system grew in sophistication with faster networks (the “Local Control Network” or LCN), more capable controller racks (the “Process Manager” or PM series), and better operator workstations. Many of these improvements were incremental, consisting of add-on components that could work with existing TDC2000 components so that the entire system need not be replaced to accept the new upgrades.

Other control equipment manufacturers responded to the DCS revolution started by Honeywell and Yokogawa by offering their own distributed control systems. The Bailey Network 90 (Net90) DCS, Bailey Infi90 DCS, and the Fisher Provox systems are examples. Foxboro, already an established leader in the control system field with their SPEC 200 analog system, first augmented the SPEC 200 with digital capabilities (the VIDEOSPEC workstation consoles, FOX I/A computer, INTERSPEC and FOXNET data networks), then developed an entirely digital distributed control system, the SPECTRUM.

Some modern distributed control systems offered at the time of this writing (2008) include:

- ABB *800xA*
- Emerson *DeltaV* and *Ovation*
- Foxboro (Invensys) *I/A*
- Honeywell *Experion PKS*
- Yokogawa *CENTUM VP* and *CENTUM CS*


For a visual comparison with the Honeywell TDC2000 DCS, examine the following photograph of an Emerson DeltaV DCS rack, with processor and multiple I/O modules:



A photograph of an Emerson Ovation DCS rack shows a vertically-oriented backplane accepting multiple I/O modules:



Many modern distributed control systems such as the Emerson DeltaV use regular personal computers rather than proprietary hardware as operator workstations. This cost-saving measure leverages existing computer and display technologies without sacrificing control-level reliability (since the control hardware and software is still industrial-grade):



As previously mentioned in the Direct Digital Control (DDC) subsection, programmable logic controllers (PLCs) are becoming more and more popular as PID control platforms due to their ever-expanding speed, functionality, and relatively low cost. It is now possible with modern PLC hardware and networking capabilities to build a truly distributed control system with individual PLCs as the processing nodes, and with redundancy built into each of those nodes so that any single failure does not interrupt critical control functions. Such a system may be purchased at a fraction of the up-front cost of a fully-fledged DCS.

However, what is currently lacking in the PLC world is the same level of hardware and software integration necessary to build a functional distributed control system that comes as ready-to-use as a system pre-built by a DCS manufacturer. In other words, if an enterprise chooses to build their own distributed control system using programmable logic controllers, they must be prepared to do a *lot* of programming work in order to emulate the same level of functionality and power as a pre-engineered DCS³⁸. Any engineer or technician who has experienced the power of a modern DCS – with its self-diagnostic, “smart” instrument management, event auditing, advanced control strategy, pre-engineered redundancy, data collection and analysis, and alarm management capabilities – realizes these features are neither luxuries nor are they trivial to engineer. Woe to anyone who thinks these critical features may be created by incumbent staff at a lesser cost!

³⁸I know of a major industrial manufacturing facility (which shall remain nameless) where a PLC vendor promised the same technical capability as a full DCS at approximately one-tenth the installed cost. Several years and several tens of thousands of man-hours later, the sad realization was this “bargain” did not live up to its promise, and the decision was made to remove the PLCs and go with a complete DCS from another manufacturer. *Caveat emptor!*


29.13.5 Fieldbus control

The DCS revolution started in the mid-1970's was fundamentally a moving of control system "intelligence" from a centralized location to distributed locations. Rather than have a single computer (or a panel full of single-loop controllers) located in a central control room implement PID control for a multitude of process loops, many (smaller) computers located closer to the process areas would execute the PID and other control functions, with network cables shuttling data between those distributed locations and the central control room.

Beginning in the late 1980's, the next logical step in this evolution of control architecture saw the relocation of control "intelligence" to the field instruments themselves. In other words, the new idea was to equip individual transmitters and control valve positioners with the necessary computational power to implement PID control all on their own, using digital networks to carry process data between the field instruments and any location desired. This is the fundamental concept of *fieldbus*.


"Fieldbus" as a technical term has multiple definitions. Many manufacturers use the word "fieldbus" to describe any digital network used to transport data to and from field instruments. In this subsection, I use the word "fieldbus" to describe a design philosophy where field instruments possess all the necessary "intelligence" to control the process, with no need for separate centralized (or even distributed) control hardware. *FOUNDATION Fieldbus* is the first standard to embody this fully-distributed control concept, the technical details of this open standard maintained and promoted by the *Fieldbus Foundation*. The aim of this Foundation is to establish an open, technical standard for *any* manufacturer to follow in the design of their fieldbus instruments. This means a FOUNDATION Fieldbus (FF) transmitter manufactured by Smar will work seamlessly with a FF control valve positioner manufactured by Fisher, communicating effortlessly with a FF-aware host system manufactured by ABB, and so on. This may be thought of in terms of being the digital equivalent of the 3-15 PSI pneumatic signal standard or the 4-20 mA analog electronic signal standard: so long as all instruments "talk" according to the same standard, brands and models may be freely interchanged to build any control system desired.

To illustrate the general fieldbus concept, consider this flow control system:




Here, a fieldbus *coupling device* provides a convenient junction point for cables coming from the transmitter, valve positioner, and host system. FOUNDATION Fieldbus devices both receive DC power and communicate digitally over the same twisted-pair cables. In this case, the host system provides DC power for the transmitter and positioner to function, while communication of process data occurs primarily between the transmitter and positioner (with little necessary involvement of the host system³⁹).

As with distributed control systems, FOUNDATION Fieldbus instruments are programmed using a function block language. In this case, we must have an analog input (for the transmitter's measurement), a PID function block, and an analog output (for the valve positioner) to make a complete flow control system:



³⁹Although it is customary for the host system to be configured as the *Link Active Scheduler* (LAS) device to schedule and coordinate all fieldbus device communications, this is not absolutely necessary. Any suitable field instrument may also serve as the LAS, which means a host system is not even necessary except to provide DC power to the instruments, and serve as a point of interface for human operators, engineers, and technicians.

The analog input (AI) block must reside in the transmitter, and the analog output (AO) block must reside in the valve positioner, since those blocks necessarily relate to the measured and controlled variables, respectively. However, the PID block may reside in *either* field device:



Practical reasons do exist for choosing one location of the PID function block over the other, most notably the difference in communication loading between the two options⁴⁰. However, there

⁴⁰With the PID function block programmed in the flow transmitter, there will be twice as many scheduled communication events per macrocycle than if the function block is programmed into the valve positioner. This is evident by the number of signal lines connecting circled block(s) to circled block(s) in the above illustration.

is no *conceptual* limitation to the location of the PID function block. In a fieldbus control system where the control “intelligence” is distributed all the way to the field instrument devices themselves, there are no limits to system flexibility.

29.14 Practical PID controller features

In order for any PID controller to be practical, it must be able to do more than just implement the PID equation. This section identifies and explains some of the basic features found on most (but not all!) modern PID controllers:

- Manual versus Automatic mode
- Output tracking
- Setpoint tracking
- Alarming
- PV characterization and damping
- Setpoint limits
- Output limits
- PID tuning security

29.14.1 Manual and automatic modes

When a controller continually calculates output values based on PV and SP values over time, it is said to be operating in *automatic* mode. This mode, of course, is what is necessary to regulate any process. There are times, however, when it is desirable to allow a human operator to manually “override” the automatic action of the PID controller. Applicable instances include process start-up and shut-down events, emergencies, and maintenance procedures. A controller that is being “overridden” by a human being is said to be in *manual mode*.

A very common application of manual mode is during maintenance of the sensing element or transmitter. If an instrument technician needs to disconnect a process transmitter for calibration or replacement, the controller receiving that transmitter’s signal cannot be left in automatic mode. If it is, then the controller may⁴¹ take sudden corrective action the moment the transmitter’s signal goes dead. If the controller is first placed in manual mode before the technician disconnects the transmitter, however, the controller will ignore any changes in the PV signal, letting its output signal be adjusted at will by the human operator. If there is another indicator of the same process variable as the one formerly reported by the disconnected transmitter, the human operator may elect to read that other indicator and play the part of a PID controller, manually adjusting the final control element to maintain the alternate indicator at setpoint while the technician completes the transmitter’s maintenance.

An extension of this “mode” concept applies to controllers configured to receive a setpoint from another device (called a *remote* or *cascaded* setpoint). In addition to an automatic and a manual mode selection, a third selection called *cascade* exists to switch the controller’s setpoint from human operator control to remote (or “cascade”) control.

⁴¹The only reason I say “may” instead of “will” is because some modern digital controllers are designed to automatically switch to manual-mode operation in the event of a sensor or transmitter signal loss. Any controller not “smart” enough to shed its operating mode to manual in the event of PV signal loss will react dramatically when that PV signal dies, and this is not a good thing for an operating loop!

29.14.2 Output and setpoint tracking

The provision of manual and automatic operating modes creates a set of potential problems for the PID controller. If, for example, a PID controller is switched from automatic to manual mode by a human operator, and then the output is manually adjusted to some new value, what will the output value do when the controller is switched *back* to automatic mode? In some crude PID controller designs, the result would be an immediate “jump” back to the output value calculated by the PID equation while the controller was in manual. In other words, some controllers never stop evaluating the PID equation – even while in manual mode – and will default to that automatically-calculated output value when the operating mode is switched from manual to automatic.

This can be very frustrating to the human operator, who may wish to use the controller’s manual mode as a way to change the controller’s bias value. Imagine, for example, that a PD controller (no integral action) is operating in automatic mode at some low output value, which happens to be too low to achieve the desired setpoint. The operator switches the controller to manual mode and then raises the output value, allowing the process variable to approach setpoint. When PV nearly equals SP, the operator switches the controller’s mode back to automatic, expecting the PID equation to start working again from this new starting point. In a crude controller, however, the output would jump back to some lower value, right where the PD equation would have placed it for these PV and SP conditions.

A feature designed to overcome this problem – which is so convenient that I consider it an essential feature of any controller with a manual mode – is called *output tracking*. With output tracking, the bias value of the controller shifts every time the controller is placed into manual mode and the output value manually changed. Thus, when the controller is switched from manual mode to automatic mode, the output does *not* immediately jump to some previously-calculated value, but rather “picks up” from the last manually-set value and begins to control from that point as dictated by the PID equation. In other words, output tracking allows a human operator to arbitrarily offset the output of a PID controller by switching to manual mode, adjusting the output value, and then switching back to automatic mode. The output will continue its automatic action from this new starting point instead of the old starting point.

A very important application of output tracking is in the manual correction of integral wind-up (sometimes called *reset windup* or just *windup*). This is what happens to a controller with integral action if for some reason the process variable *cannot* achieve setpoint no matter how far the output signal value is driven by integral action. An example might be on a temperature controller where the source of heat for the process is a steam system. If the steam system shuts down, the temperature controller *cannot* warm the process up to the temperature setpoint value no matter how far open the steam valve is driven by integral action. If the steam system is shut down for too long, the result will be a controller output saturated at maximum value in a futile attempt to warm the process. If and when the steam system starts back up, the controller’s saturated output will now send *too much heating steam* to the process, causing the process temperature to overshoot setpoint until integral action drives the output signal back down to some reasonable level. This situation may be averted, however, if the operator switches the temperature controller to manual mode as soon as the steam system shuts down. Even if this preventive step is not taken, the problem of overshoot may be averted upon steam system start-up if the operator uses output tracking by quickly switching the controller into manual mode, adjusting the output down to a reasonable level, and then switching back into automatic mode so that the controller’s output value is no longer “wound up” at a high

level⁴².

A similar feature to output tracking – also designed for the convenience of a human operator switching a PID controller between automatic and manual modes – is called *setpoint tracking*. The purpose of setpoint tracking is to equalize SP and PV while the controller is in manual mode, so that when the controller gets switched back into automatic mode, it will begin its automatic operation with no error ($PV = SP$).

This feature is most useful during system start-ups, where the controller may have difficulty controlling the process in automatic mode under unusual conditions. Operators often prefer to run certain control loops in manual mode from the time of initial start-up until such time that the process is near normal operating conditions. At that point, when the operator is content with the stability of the process, the controller is assigned the responsibility of maintaining the process at setpoint. With setpoint tracking present in the controller, the controller's SP value will be held equal to the PV value (whatever that value happens to be) for the entire time the controller is in manual mode. Once the operator decides it is proper to switch the controller into automatic mode, the SP value freezes at that last manual-mode PV value, and the controller will continue to control the PV at that SP value. Of course, the operator is free to adjust the SP value to any new value while the controller is in automatic mode, but this is at the operator's discretion.

Without setpoint tracking, the operator would *have to* make a setpoint adjustment either before or after switching the controller from manual mode to automatic mode, in order to ensure the controller was properly set up to maintain the process variable at the desired value. With setpoint tracking, the setpoint value will default to the process variable value when the controller was last in manual mode, which (it is assumed) will be close enough to the desired value to suffice for continued operation.

Unlike output tracking, for which there is virtually no reason not to have the feature present in a PID controller, there may very well be applications where we do not wish to have setpoint tracking. For some processes⁴³, the setpoint value *should* remain fixed at all times, and as such it would be undesirable to have the setpoint value drift around with the process variable value every time the controller was placed into manual mode.

⁴²I once had the misfortune of working on an analog PID controller for a chlorine-based wastewater disinfection system that lacked output tracking. The chlorine sensor on this system would occasionally fail due to sample system plugging by algae in the wastewater. When this happened, the PV signal would fail low (indicating abnormally low levels of chlorine gas dissolved in the wastewater) even though the actual dissolved chlorine gas concentration was adequate. The controller, thinking the PV was well below SP, would ramp the chlorine gas control valve further and further open over time, as integral action attempted to reduce the error between PV and SP. The error never went away, of course, because the chlorine sensor was plugged with algae and simply could not detect the actual chlorine gas concentration in the wastewater. By the time I arrived to address the “low chlorine” alarm, the controller output was already wound up to 100%. After cleaning the sensor, and seeing the PV value jump up to some outrageously high level, the controller would take a long time to “wind down” its output because its integral action was very slow. I could not use manual mode to “unwind” the output signal, because this controller lacked the feature of output tracking. My “work-around” solution to this problem was to re-tune the integral term of the controller to some really fast time constant, watch the output “wind down” in fast-motion until it reached a reasonable value, then adjust the integral time constant back to its previous value for continued automatic operation.

⁴³Boiler steam drum water level control, for example, is a process where the setpoint really should be left at a 50% value at all times, even if there maybe legitimate reasons for occasionally switching the controller into manual mode.

29.14.3 Alarm capabilities

A common feature on many instrument systems is the ability to alert personnel to the onset of abnormal process conditions. The general term for this function is *alarm*. Process alarms may be triggered by process switches directly sensing abnormal conditions (e.g. high-temperature switches, low-level alarms, low-flow alarms, etc.), in which case they are called *hard alarms*. A *soft alarm*, by contrast, is an alarm triggered by some continuous measurement (i.e. a signal from a process transmitter rather than a process switch) exceeding a pre-programmed alarm limit value.

Since PID controllers are designed to input continuous process measurements, it makes sense that a controller could be equipped with programmable alarm limit values as well, to provide “soft” alarm capability without adding additional instruments to the loop⁴⁴. Not only is PV alarming easy to implement in most PID controllers, but *deviation* alarming is easy to implement as well. A “deviation alarm” is a soft alarm triggered by excessive deviation (error) between PV and SP. Such an event indicates control problems, since a properly-operating feedback loop should be able to maintain reasonable agreement between PV and SP at all times.

Alarm capabilities find their highest level of refinement in modern distributed control systems (DCS), where the networked digital controllers of a DCS provide convenient access and advanced management of hard and soft alarms alike. Not only can alarms be accessed from virtually any location in a facility in a DCS, but they are usually time-stamped and archived for later analysis, which is an *extremely* important feature for the analysis of emergency events, and the continual improvement of process safety.

29.14.4 Output and setpoint limiting

In some process applications, it may not be desirable to allow the controller to automatically manipulate the final control element (control valve, variable-speed motor, heater) over its full 0% - 100% range. In such applications, a useful controller feature is an *output limit*. For example, a PID flow controller may be configured to have a minimum output limit of 5%, so that it is not able to close the control valve any further than the 5% open position in order to maintain “minimum flow” through a pump. The valve may still be fully closed (0% stem position) in manual mode, but just not in automatic mode⁴⁵.

Similarly, setpoint values may be internally limited in some PID controllers, such that an operator cannot adjust the setpoint above some limiting value or below some other limiting value. In the event that the process variable *must* be driven outside these limits, the controller may be placed in manual mode and the process “manually” guided to the desired state by an operator.

⁴⁴It is very important to note that soft alarms are not a replacement for hard alarms. There is much wisdom in maintaining both hard and soft alarms for a process, so there will be redundant, non-interactive levels of alarming. Hard and soft alarms should complement each other in any critical process.

⁴⁵Some PID controllers limit manual-mode output values as well, so be sure to check the manufacturer’s documentation for output limiting on your particular PID controller!

29.14.5 Security

There is justifiable reason to prevent certain personnel from having access to certain parameters and configurations on PID controllers. Certainly, operations personnel need access to setpoint adjustments and automatic/manual mode controls, but it may be unwise to grant those same operators unlimited access to PID tuning constants and output limits. Similarly, instrument technicians may require access to a PID controller's tuning parameters, but perhaps should be restricted from editing configuration programs maintained by the engineering staff.

Most digital PID controllers have some form of security access control, allowing for different levels of permission in altering PID controller parameters and configurations. Security may be crude (a hidden switch located on a printed circuit board, which only the maintenance personnel should know about), sophisticated (login names and passwords, like a multi-user computer system), or anything in between, depending on the level of development invested in the feature by the controller's manufacturer.

An interesting solution to the problem of security in the days of analog control systems was the architecture of Foxboro's SPEC 200 analog electronic control system. The controller displays, setpoint adjustments, and auto/manual mode controls were located on the control room panel where anyone could access them. All other adjustments (PID settings, alarm settings, limit settings) could be located in the *nest* area where all the analog circuit control cards resided. Since the "nest" racks could be physically located in a room separate from the control room, personnel access to the nest room served as access security to these system parameters.

At first, the concept of controller parameter security may seem distrustful and perhaps even insulting to those denied access, especially when the denied persons possess the necessary knowledge to understand the functions and consequences of those parameters. It is not uncommon for soft alarm values to be "locked out" from operator access despite the fact that operators understand very well the purpose and functions of these alarms. At some facilities, PID tuning is the exclusive domain of process engineers, with instrument technicians and operators alike barred from altering PID tuning constants even though some operators and many technicians may well understand PID controller tuning.

When considering security access, there is more to regard than just knowledge or ability. At a fundamental level, security is a task of limiting access commensurate with *responsibility*. In other words, security restrictions exist to exclude those not charged with particular responsibilities. Knowledge and ability are necessary conditions of responsibility (i.e. one cannot reasonably be held responsible for something beyond their knowledge or control), but they are not *sufficient* conditions of responsibility (i.e. knowing how to, and being able to perform a task does not confer responsibility for that task getting completed). An operator may very well understand how and why a soft alarm on a controller works, but the responsibility for altering the alarm value may reside with someone else whose job description it is to ensure the alarm values correspond to plant-wide policies.

29.15 Digital PID algorithms

Instrument technicians should not have to concern themselves over the programming details internal to digital PID controllers. Ideally, a digital PID controller should simply perform the task of executing PID control with all the necessary features (setpoint tracking, output limiting, etc.) without the end-user having to know anything about those details. However, in my years of experience I have seen enough examples of poor PID implementation to warrant an explanatory section in this book, both so instrumentation professionals may recognize poor PID implementation when they see it, and also so those with the responsibility of designing PID algorithms may avoid some common mistakes.

29.15.1 Introduction to pseudocode

In order to show digital algorithms, I will use a form of notation called *pseudocode*: a text-based language instructing a digital computing device to implement step-by-step procedures. “Pseudocode” is written to be easily read and understood by human beings, yet similar enough in syntax and structure to real computer programming languages for a human programmer to be able to easily translate to a high-level programming language such as BASIC, C++, or Java. Since pseudocode is not a formal computer language, we may use it to very efficiently describe certain algorithms (procedures) without having to abide by strict “grammatical” rules as we would if writing in a formal language such as BASIC, C++, or Java.

Program loops

Each line of text in the following listing represents a command for the digital computer to follow, one by one, in order from top to bottom. The LOOP and ENDLLOOP markers represent the boundaries of a program *loop*, where the same set of encapsulated commands are executed over and over again in cyclic fashion:

Pseudocode listing ⁴⁶

```

LOOP
  PRINT "Hello World!" // This line prints text to the screen
  OUTPUT audible beep on the speaker // This line beeps the speaker
ENDLOOP

```

In this particular case, the result of this program’s execution is a continuous printing of the words “Hello World!” to the computer’s display with a single “beep” tone accompanying each printed line. The words following a double-slash (//) are called *comments*, and exist only to provide explanatory text for the human reader, not the computer. Admittedly, this example program would be both impractical and annoying to actually run in a computer, but it does serve to illustrate the basic concept of a program “loop” shown in pseudocode.

⁴⁶I have used a typesetting convention to help make my pseudocode easier for human beings to read: all formal commands appear in bold-faced blue type, while all comments appear in italicized red type. All other text appears as normal-faced black type. One should remember that the computer running any program cares not for how the text is typeset: all it cares is that the commands are properly used (i.e. no “grammatical” or “syntactical” errors).

Assigning values

For another example of pseudocode, consider the following program. This code causes a variable (x) in the computer's memory to alternate between two values of 0 and 2 indefinitely:

Pseudocode listing

```
DECLARE x to be an integer variable
SET x = 2 // Initializing the value of x

LOOP
// This SET command alternates the value of x with each pass
SET x = 2 - x
ENDLOOP
```

The first instruction in this listing declares the type of variable x will be. In this case, x will be an *integer* variable, which means it may only represent whole-number quantities and their negative counterparts – no other values (e.g. fractions, decimals) are possible. If we wished to limit the scope of x even further to represent just 0 or 1 (i.e. a single bit), we would have to declare it as a *Boolean* variable. If we required x to be able to represent fractional values as well, we would have to declare it as a *floating-point* variable. Variable declarations are important in computer programming because it instructs the computer how much space in its random-access memory to allocate to each variable, which necessarily limits the range of numbers each variable may represent.

The next instruction initializes x to a value of two. Like the declaration, this instruction need only happen once at the beginning of the program's execution, and never again so long as the program continues to run. The single **SET** statement located between the **LOOP** and **ENDLOOP** markers, however, repeatedly executes as fast as the computer's processor allows, causing x to rapidly alternate between the values of two and zero.

It should be noted that the “equals” sign (=) in computer programming often has a different meaning from that commonly implied in ordinary mathematics. When used in conjunction with the **SET** command, an “equals” sign *assigns* the value of the right-hand quantity to the left-hand variable. For example, the command **SET x = 2 - x** tells the computer to first calculate the quantity $2 - x$ and then set the variable x to this new value. It definitely does *not* mean to imply x is actually equal in value to $2 - x$, which would be a mathematical contradiction. Thus, you should interpret the **SET** command to mean “set equal to . . .”

Testing values (conditional statements)

If we mean to simply test for an equality between two quantities, we may use the same symbol (=) in the context of a different command, such as “IF”:

Pseudocode listing

```
DECLARE x to be an integer variable

LOOP

    // (other code manipulating the value of x goes here)

    IF x = 5 THEN
        PRINT "The value of the number is 5"
        OUTPUT audible beep on the speaker
    ENDIF
ENDLOOP
```

This program repeatedly tests whether or not the variable x is equal to 5, printing a line of text and producing a “beep” on the computer’s speaker if that test evaluates as true. Here, the context of the IF command tells us the equals sign is a test for equality rather than a command to assign a new value to x . If the condition is met ($x = 5$) then all commands contained within the IF/ENDIF set are executed.

Some programming languages draw a more explicit distinction between the operations of equality test versus assignment by using different symbol combinations. In C and C++, for example, a single equals sign (=) represents assignment while a double set of equals signs (==) represents a test for equality. In Structured Text (ST) PLC programming, a single equals sign (=) represents a test for equality, while a colon plus equals sign (:=) represents assignment. The combination of an exclamation point and an equals sign (!=) represents “not equal to,” used as a test condition to check for *inequality* between two quantities.

Branching and functions

A very important feature of any programming language is the ability for the path of execution to change (i.e. the program “flow” to *branch* in another direction) rather than take the exact same path every time. We saw shades of this with the IF statement in our previous example program: the computer would print some text and output a beep sound if the variable x happened to be equal to 5, but would completely skip the PRINT and OUTPUT commands if x happened to be any other value.

An elegant way to modularize a program into separate pieces involves writing portions of the program as separate *functions* which may be “called” as needed by the main program. Let us examine how to apply this concept to the following conditional program:

Pseudocode listing

```
DECLARE x to be an integer variable
LOOP
    // (other code manipulating the value of x goes here)
    IF x = 5 THEN
        PRINT "The value of the number is 5"
        OUTPUT audible beep on the speaker
    ELSEIF x = 7 THEN
        PRINT "The value of the number is 7"
        OUTPUT audible beep on the speaker
    ELSEIF x = 11 THEN
        PRINT "The value of the number is 11"
        OUTPUT audible beep on the speaker
    ENDIF
ENDLOOP
```

This program takes action (printing and outputting beeps) if ever the variable x equals either 5, 7, or 11, but not for any other values of x . The actions taken with each condition are quite similar: print the numerical value of x and output a single beep. In fact, one might argue this code is ugly because we have to keep repeating one of the commands verbatim: the OUTPUT command for each condition where we wish to computer to output a beep sound.

We may streamline this program by placing the PRINT and OUTPUT commands into their own separate “function” written outside the main loop, and then *call* that function whenever we need it. The boundaries of this function’s code are marked by the BEGIN and END labels shown near the bottom of the listing:

Pseudocode listing

```
DECLARE n to be an integer variable
DECLARE x to be an integer variable
DECLARE PrintAndBeep to be a function

LOOP

    // (other code manipulating the value of x goes here)

    IF x = 5 OR x = 7 OR x = 11 THEN
        CALL PrintAndBeep(x)
    ENDIF
ENDLOOP

BEGIN PrintAndBeep (n)
    PRINT "The value of the number is" (n) "!"
    OUTPUT audible beep on the speaker
    RETURN
END PrintAndBeep
```

The main program loop is much shorter than before because the repetitive tasks of printing the value of x and outputting beep sounds has been moved to a separate function. In older computer languages, this was known as a *subroutine*, the concept being that flow through the main program (the “routine”) would branch to a separate sub-program (a “subroutine”) to do some specialized task and then return back to the main program when the sub-program was done with its task.

Note that the program execution flow never reaches the PrintAndBeep function unless x happens to equal 5, 7, or 11. If the value of x never matches any of those specific conditions, the program simply keeps looping between the LOOP and ENDLOOP markers.

Note also how the value of x gets *passed* on to the PrintAndBeep function, then read inside that function under another variable name, n . This was not strictly necessary for the purpose of printing the value of x , since x is the only variable in the main program. However, the use of a separate (“local”) variable within the PrintAndBeep function enables us at some later date to use that function to act on other variables within the main program while avoiding conflict. Take this program for example:

Pseudocode listing

```

DECLARE n to be an integer variable
DECLARE x to be an integer variable
DECLARE y to be an integer variable
DECLARE PrintAndBeep to be a function

LOOP

    // (other code manipulating the value of x and y goes here)

    IF x = 5 OR x = 7 OR x = 11 THEN
        CALL PrintAndBeep(x)
    ENDIF
    IF y = 0 OR y = 2 THEN
        CALL PrintAndBeep(y)
    ENDIF
ENDLOOP

BEGIN PrintAndBeep (n)
    PRINT "The value of the number is" (n) "!"
    OUTPUT audible beep on the speaker
    RETURN
END PrintAndBeep

```

Here, the PrintAndBeep function gets used to print certain values of x , then re-used to print certain values of y . If we had used x within the PrintAndBeep function instead of its own variable (n), the function would only be useful for printing the value of x . Being able to pass values to functions makes those functions more useful.

A final note on branching and functions: most computer languages allow a function to call itself if necessary. This concept is known as *recursion* in computer science.

29.15.2 Position versus velocity algorithms

The canonical “ideal” or “ISA” variety of PID equation takes the following form:

$$m = K_p \left(e + \frac{1}{\tau_i} \int e \, dt + \tau_d \frac{de}{dt} \right)$$

Where,

m = Controller output

e = Error (SP – PV or PV – SP, depending on controller action being direct or reverse)

K_p = Controller gain

τ_i = Integral (reset) time constant

τ_d = Derivative (rate) time constant

The same equation may be written in terms of “gains” rather than “time constants” for the integral and derivative terms. This re-writing exhibits the advantage of consistency from the perspective of PID tuning, where each tuning constant has the same (increasing) effect as its numerical value grows larger:

$$m = K_p \left(e + K_i \int e \, dt + K_d \frac{de}{dt} \right)$$

Where,

m = Controller output

e = Error

K_p = Controller gain

K_i = Integral (reset) gain (repeats per unit time)

K_d = Derivative (rate) gain

However the equation is written, there are two major ways in which it is commonly implemented in a digital computer. One way is the *position* algorithm, where the result of each pass through the program “loop” calculates the actual output value. If the final control element for the loop is a control valve, this value will be the position of that valve’s stem, hence the name *position algorithm*. The other way is the so-called *velocity* algorithm, where the result of each pass through the program “loop” calculates the amount the output value will *change*. Assuming a control valve for the final control element once again, the value calculated by this algorithm is the distance the valve stem will travel *per scan of the program*. In other words, the magnitude of this value describes how *fast* the valve stem will travel, hence the name *velocity algorithm*.

Mathematically, the distinction between the position and velocity algorithms is a matter of differentials: the position equation solves for the output value (m) directly while the velocity equation solves for small increments (differentials) of m , or dm .

A comparison of the position and velocity equations shows both the similarities and the differences:

$$m = K_p \left(e + K_i \int e \, dt + K_d \frac{de}{dt} \right) \quad \text{Position equation}$$

$$dm = K_p \left(de + K_i e \, dt + K_d \frac{d^2e}{dt^2} \right) \quad \text{Velocity equation}$$

Of the two approaches to implementing PID control, the position algorithm makes the most intuitive sense and is the easiest to understand.

We will begin our exploration of both algorithms by examining their application to proportional-only control. This will be a simpler and “gentler” introduction than showing how to implement full PID control. The two respective proportional-only control equations we will consider are shown here:

$$m = K_p e + \text{Bias} \quad \text{Position equation for P-only control}$$

$$dm = K_p de \quad \text{Velocity equation for P-only control}$$

You will notice how a “bias” term is required in the position equation to keep track of the output’s “starting point” each time a new output value is calculated. No such term is required in the velocity equation, because the computer merely calculates *how far the output moves from its last value* rather than the output’s value from some absolute reference.

First, we will examine a simple pseudocode program for implementing the proportional-only equation in its “position” form:

Pseudocode listing for a “position algorithm” proportional-only controller

```

DECLARE PV, SP, and Out to be floating-point variables
DECLARE K_p, Error, and Bias to be floating-point variables
DECLARE Action, and Mode to be boolean variables

LOOP
  SET PV = analog_input_channel_N // Update PV
  SET K_p = operator_input_channel_Gain // From operator interface

  IF Action = 1 THEN
    SET Error = SP - PV // Calculate error assuming reverse action
  ELSE THEN
    SET Error = PV - SP // Calculate error assuming direct action
  ENDIF

  IF Mode = 1 THEN // Automatic mode (if Mode = 1)
    SET Out = K_p * Error + Bias
    SET SP = operator_input_channel_SP // From operator interface
  ELSE THEN // Manual mode (if Mode = 0)
    SET Out = operator_input_channel_Out // From operator interface
    SET SP = PV // Setpoint tracking
    SET Bias = Out // Output tracking
  ENDIF
ENDLOOP

```

The first SET instructions within the loop update the PV to whatever value is being measured by the computer’s analog input channel (channel N in this case), and the K_p variable to whatever value is entered by the human operator through the use of a keypad, touch-screen interface, or networked computer. Next, a set of IF/THEN conditionals determines which way the error should be calculated: $\text{Error} = \text{SP} - \text{PV}$ if the control action is “reverse” ($\text{Action} = 1$) and $\text{Error} = \text{PV} - \text{SP}$ if the control action is “direct” ($\text{Action} = 0$).

The next set of conditional instructions determines what to do in automatic versus manual modes. In automatic mode ($\text{Mode} = 1$), the output value is calculated according to the position equation and the setpoint comes from a human operator’s input. In manual mode ($\text{Mode} = 0$), the output value is no longer calculated by an equation but rather is obtained from the human operator’s input, the setpoint is forced equal to the process variable, and the Bias value is continually made equal to the value of the output. Setting $\text{SP} = \text{PV}$ provides the convenient feature of *setpoint tracking*, ensuring an initial error value of zero when the controller is switched back to automatic mode. Setting the Bias equal to the output provides the essential feature of *output tracking*, where the controller begins automatic operation at an output value precisely equal to the last manual-mode output value.

Next, we will examine a simple pseudocode program for implementing the proportional-only equation in its “velocity” form:

Pseudocode listing for a “velocity algorithm” proportional-only controller

```

DECLARE PV, SP, and Out to be floating-point variables
DECLARE K_p, Error, and last_Error to be floating-point variables
DECLARE Action, and Mode to be boolean variables

LOOP
  SET PV = analog_input_channel_N // Update PV
  SET K_p = operator_input_channel_Gain // From operator interface
  SET last_Error = Error

  IF Action = 1 THEN
    SET Error = SP - PV // Calculate error assuming reverse action
  ELSE THEN
    SET Error = PV - SP // Calculate error assuming direct action
  ENDIF

  IF Mode = 1 THEN // Automatic mode (if Mode = 1)
    SET Out = Out + (K_p * (Error - last_Error))
    SET SP = operator_input_channel_SP // From operator interface
  ELSE THEN // Manual mode (if Mode = 0)
    SET Out = operator_input_channel_Out // From operator interface
    SET SP = PV // Setpoint tracking
  ENDIF
ENDLOOP

```

The code for the velocity algorithm is mostly identical to the code for the position algorithm, with just a few minor changes. The first difference we encounter in reading the code from top to bottom is that we calculate a new variable called “last_Error” immediately prior to calculating a new value for Error. The reason for doing this is to provide a way to calculate the differential *change* in error (*de*) from scan to scan of the program. The variable “last_Error” remembers the value of Error during the previous scan of the program. Thus, the expression “Error – last_Error” is equal to the amount the error has changed from last scan to the present scan.

When the time comes to calculate the output value in automatic mode, we see the SET command calculating the change in output (K_p multiplied by the change in error), then adding this change in output to the existing output value to calculate a new output value. This is how the program translates calculated output increments into an actual output value to drive a final control element. The mathematical expression “K_p * (Error – last_Error)” defines the incremental change in output value, and this increment is then added to the current output value to generate a new output value.

From a human operator’s point of view, the position algorithm and the velocity algorithm are identical with one exception: how each controller reacts to a sudden change in gain (K_p). To understand this difference, let us perform a “thought experiment” where we imagine a condition of

constant error between PV and SP. Suppose the controller is operating in automatic mode, with a setpoint of 60% and a (steady) process variable value of 57%. We should not be surprised that a constant error might exist for a proportional-only controller, since we should be well aware of the phenomenon of *proportional-only offset*.

How will this controller react if the gain is suddenly increased in value while operating in automatic mode? If the controller executes the position algorithm, the result of a sudden gain change will be a sudden change in its output value, since output is a direct function of error and gain. However, if the controller executes the velocity algorithm, the result of a sudden gain change will be no change to the output at all, so long as the error *remains constant*. Only when the error begins to change will there be any noticeable difference in the controller's behavior compared to how it acted before the gain change. This is because the velocity algorithm is a function of gain and *change in error*, not error directly.

Comparing the two responses, the velocity algorithm's response to changes in gain is regarded as "better-mannered" than the position algorithm's response to changes in gain. When tuning a controller, we would rather not have the controller's output suddenly jump in response to simple gain changes⁴⁷, and so the velocity algorithm is generally preferred. If we allow the gain of the algorithm to be set by another process variable⁴⁸, the need for "stable" gain-change behavior becomes even more important.

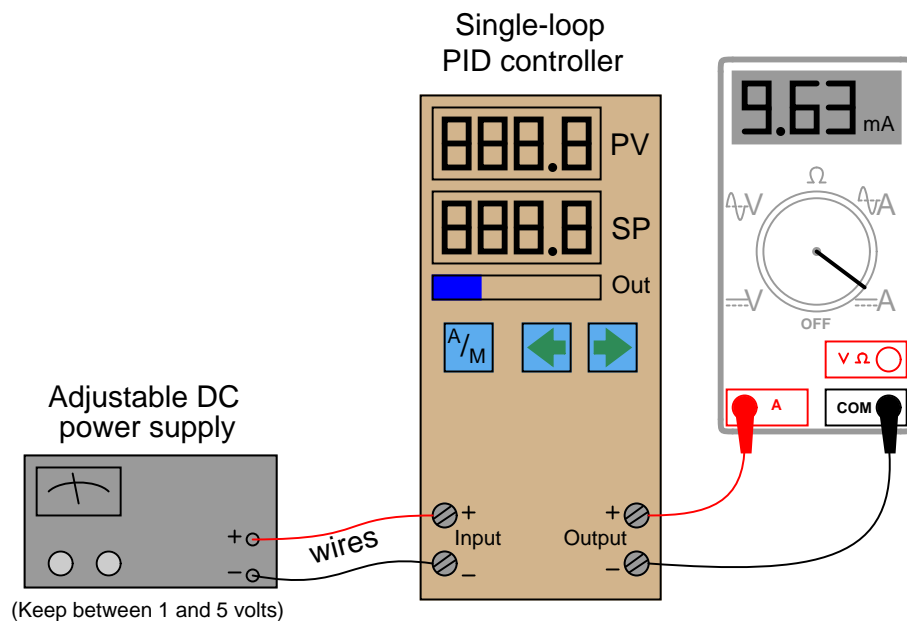
⁴⁷It should be noted that this is precisely what happens when you change the gain in a pneumatic or an analog electronic controller, since all analog PID controllers implement the "position" equation. Although the choice between "position" and "velocity" algorithms in a digital controller is arbitrary, it is *much* easier to build an analog mechanism or circuit implementing the position algorithm than it is to build an analog "velocity" controller.

⁴⁸We call this an *adaptive gain* control system.

29.16 Note to students

PID control can be a frustrating subject for many students, even those with previous knowledge of calculus. At times it can seem like an impossibly abstract concept to master.

Thankfully, there is a relatively simple way to make PID control more “real,” and that is hands-on experience with a real PID controller. I advise you acquire an electronic single-loop PID controller⁴⁹ and set it up with an adjustable DC power supply and milliammeter as such:



Most electronic controllers input a 1 to 5 VDC signal for the process variable (often with a 250 ohm resistor connected across the input terminals to generate a 1-5 VDC drop from a 4-20 mA current signal, which you will not need here). By adjusting the DC power supply between 1 and 5 volts DC, you will simulate a transmitter signal to the controller’s input between 0% and 100%.

The milliammeter reads current output by the controller, 4 mA representing a 0% output signal and 20 mA representing a 100% output signal. With the power supply and milliammeter both connected to the appropriate terminals on the controller, you are all set to simulate input conditions and watch the controller’s output response.

This arrangement does not simulate a process, and so there will be no feedback for you to observe. The purpose of this setup is to simply learn how the controller is supposed to respond to different PV and SP conditions, so that you may gain an intuitive “feel” for the PID algorithm to supplement your theoretical understanding of it. Experimentation with a real process (or even a simulated process) comes later (see section 30.6 beginning on page 2489).

⁴⁹Many instrument manufacturers sell simple, single-loop controllers for reasonable prices, comparable to the price of a college textbook. You need to get one that accepts 1-5 VDC input signals and generates 4-20 mA output signals, and has a “manual” mode of operation in addition to automatic – these features are *very important!* Avoid controllers that can only accept thermocouple inputs, and/or only have time-proportioning (PWM) outputs.

Once you have all components connected, you should check to see that everything works:

- Set power supply to 1, 3, and then 5 volts DC. The controller's PV display should read 0%, 50%, and 100%, respectively. The PV display should follow closely to the power supply voltage signal over time. If the display seems to "lag" behind the power supply adjustment, then it means the controller has damping configured for the input signal. You should keep the damping set to the minimum possible value, so the controller is as responsive as it can be.
- Put the controller in manual mode and set the output to 0%, 50%, and then 100%. The milliammeter should register 4 mA, 12 mA, and 20 mA, respectively.

After checking these basic functions, you may proceed to do the following experiments. For each experiment, I recommend setting the PV input signal to 3 volts DC (50%), and manually setting the output to 50% (12 mA on the milliammeter). When you are ready to test the P,I,D responses of the controller, place the controller into automatic mode and then observe the results.

29.16.1 Proportional-only control action

1. Set the controller PV input to 50% (3 volts) and the output value to 50% in manual mode.
2. Configure the controller for reverse action (this is typically the default setting).
3. Configure the PID settings for proportional action only. This may be done by setting the gain equal to 1 (P.B. = 100%), the integral setting to zero repeats per minute (maximum minutes per repeat), and the derivative setting to zero minutes. Some controllers have the ability to switch to a "proportional-only" algorithm – if your controller has that ability, this is the best way to get set up for this exercise.
4. Switch the controller mode to "automatic."
5. Adjust the PV signal to 75% (4 volts) and observe the output. How far does the output signal move from its starting value of 50%? How does the magnitude of this step relate to the magnitude of the PV step? Does the output signal drift or does it remain the same when you stop changing the PV signal?
6. Adjust the PV signal to 25% (2 volts) and observe the output. How far does the output signal move from its starting value of 50%? How does the magnitude of this step relate to the magnitude of the PV step? Does the output signal drift or does it remain the same when you stop changing the PV signal?
7. Change the controller's gain setting to some different value and repeat the previous two steps. How does the output step magnitude relate to the input step-changes in each case? Do you see the relationship between controller gain and how the output responds to changes in the input?
8. Smoothly vary the input signal back and forth between 0% and 100% (1 and 5 volts). How does the output respond when you do this? Try changing the gain setting again and re-checking.
9. Switch the controller's action from *reverse* to *direct*, then repeat the previous step. How does the output respond now?

29.16.2 Integral-only control action

1. Set the controller PV input to 50% (3 volts) and the output value to 50% in manual mode.
2. Configure the controller for reverse action (this is typically the default setting).
3. Configure the PID settings for integral action only. If the controller has an “I-only” mode, this is the best way to get set up for this exercise. If there is no way to completely turn off proportional action, then I recommend setting the gain value to the minimum non-zero value allowed, and setting the integral constant to an aggressive value (many repeats per minute, or fractions of a minute per repeat). If your controller does have an integral-only option, I recommend setting the integral time constant at 1 minute. Set derivative action at zero minutes.
4. Switch the controller mode to “automatic.”
5. Adjust the PV signal to 75% (4 volts) and observe the output. Which way does the output signal move? Does the output signal drift or does it remain the same when you stop changing the PV signal? How does this action compare with the proportional-only test?
6. Adjust the PV signal to 25% (2 volts) and observe the output. Which way does the output signal move? Does the output signal drift or does it remain the same when you stop changing the PV signal? How does this action compare with the proportional-only test?
7. Change the controller’s integral setting to some different value and repeat the previous two steps. How does the rate of output ramping relate to the input step-changes in each case? Do you see the relationship between the integral time constant and how the output responds to changes in the input?
8. Smoothly vary the input signal back and forth between 0% and 100% (1 and 5 volts). How does the output respond when you do this? Try changing the integral setting again and re-checking.
9. Where must you adjust the input signal to get the output to stop moving? When the output finally does settle, is its value consistent (i.e. does it always settle at the same value, or can it settle at different values)?
10. Switch the controller’s action from *reverse* to *direct*, then repeat the previous two steps. How does the output respond now?

29.16.3 Proportional plus integral control action

1. Set the controller PV input to 50% (3 volts) and the output value to 50% in manual mode.
2. Configure the controller for reverse action (this is typically the default setting).
3. Configure the PID settings with a proportional (gain) value of 1 (P.B. = 100%) and an integral value of 1 repeat per minute (or 1 minute per repeat). Set derivative action at zero minutes.
4. Switch the controller mode to “automatic.”
5. Adjust the PV signal to 75% (4 volts) and observe the output. Which way does the output signal move? Does the output signal drift or does it remain the same when you stop changing the PV signal? How does this action compare with the proportional-only test and with the integral-only test?
6. Adjust the PV signal to 25% (2 volts) and observe the output. Which way does the output signal move? Does the output signal drift or does it remain the same when you stop changing the PV signal? How does this action compare with the proportional-only test and with the integral-only test?
7. Change the controller’s gain setting to some different value and repeat the previous two steps. Can you tell which aspect of the output signal’s response is due to proportional action and which aspect is due to integral action?
8. Change the controller’s integral setting to some different value and repeat those same two steps. Can you tell which aspect of the output signal’s response is due to proportional action and which aspect is due to integral action?
9. Smoothly vary the input signal back and forth between 0% and 100% (1 and 5 volts). How does the output respond when you do this? Try changing the gain and/or integral settings again and re-checking.
10. Switch the controller’s action from *reverse* to *direct*, then repeat the previous two steps. How does the output respond now?

29.16.4 Proportional plus derivative control action

1. Set the controller PV input to 50% (3 volts) and the output value to 50% in manual mode.
2. Configure the controller for reverse action (this is typically the default setting).
3. Configure the PID settings with a proportional (gain) value of 1 (P.B. = 100%) and a derivative value of 1 minute. Set integral action at zero repeats per minute (maximum number of minutes per repeat).
4. Switch the controller mode to “automatic.”
5. Adjust the PV signal to 75% (4 volts) and observe the output. Which way does the output signal move? How does the output signal value compare while you are adjusting the input voltage versus after you reach 4 volts and take your hand off the adjustment knob? How does this action compare with the proportional-only test?
6. Adjust the PV signal to 25% (2 volts) and observe the output. Which way does the output signal move? How does the output signal value compare while you are adjusting the input voltage versus after you reach 4 volts and take your hand off the adjustment knob? How does this action compare with the proportional-only test?
7. Change the controller’s gain setting to some different value and repeat the previous two steps. Can you tell which aspect of the output signal’s response is due to proportional action and which aspect is due to derivative action?
8. Change the controller’s derivative setting to some different value and repeat those same two steps. Can you tell which aspect of the output signal’s response is due to proportional action and which aspect is due to derivative action?
9. Smoothly vary the input signal back and forth between 0% and 100% (1 and 5 volts). How does the output respond when you do this? Try changing the derivative setting again and re-checking.
10. Switch the controller’s action from *reverse* to *direct*, then repeat the previous two steps. How does the output respond now?

29.16.5 Full PID control action

1. Set the controller PV input to 50% (3 volts) and the output value to 50% in manual mode.
2. Configure the controller for reverse action (this is typically the default setting).
3. Configure the PID settings with a proportional (gain) value of 1 (P.B. = 100%), an integral value of 1 repeat per minute (or 1 minute per repeat), and a derivative action of 1 minute.
4. Switch the controller mode to “automatic.”
5. Adjust the PV signal to 75% (4 volts) and observe the output. Which way does the output signal move? Does the output signal drift or does it remain the same when you stop changing the PV signal? How does magnitude of the output signal compare while you are changing the input voltage, versus when the input signal is steady?
6. Adjust the PV signal to 25% (2 volts) and observe the output. Which way does the output signal move? Does the output signal drift or does it remain the same when you stop changing the PV signal? How does magnitude of the output signal compare while you are changing the input voltage, versus when the input signal is steady?
7. Change the controller’s gain setting to some different value and repeat the previous two steps. Can you tell which aspect of the output signal’s response is due to proportional action, which aspect is due to integral action, and which aspect is due to derivative action?
8. Change the controller’s integral setting to some different value and repeat the same two steps. Can you tell which aspect of the output signal’s response is due to proportional action, which aspect is due to integral action, and which aspect is due to derivative action?
9. Change the controller’s derivative setting to some different value and repeat the same two steps. Can you tell which aspect of the output signal’s response is due to proportional action, which aspect is due to integral action, and which aspect is due to derivative action?
10. Smoothly vary the input signal back and forth between 0% and 100% (1 and 5 volts). How does the output respond when you do this? Try changing the gain, integral, and/or derivative settings again and re-checking.
11. Switch the controller’s action from *reverse* to *direct*, then repeat the previous two steps. How does the output respond now?

29.17 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Linear equations:** any function represented by a straight line on a graph may be represented symbolically by the slope-intercept formula $y = mx + b$. Relevant to proportional control algorithms.
- **Zero shift:** any shift in the offset of an instrument is fundamentally additive, being represented by the "intercept" (b) variable of the slope-intercept linear formula $y = mx + b$. Relevant to controller tuning: adjusting the "bias" of a loop controller always adds to or subtracts from its output signal.
- **Span shift:** any shift in the gain of an instrument is fundamentally multiplicative, being represented by the "slope" (m) variable of the slope-intercept linear formula $y = mx + b$. Relevant to controller tuning: adjusting the "gain" of a loop controller always multiplies or divides the response of its output for a given input change.
- **Negative feedback:** when the output of a system is degeneratively fed back to the input of that same system, the result is decreased (overall) gain and greater stability. Relevant to loop controller action: in order for a control system to be stable, the feedback must be negative.
- **Self-balancing pneumatic mechanisms:** all self-balancing pneumatic instruments work on the principle of negative feedback maintaining a nearly constant baffle-nozzle gap. Force-balance mechanisms maintain this constant gap by balancing force against force with negligible motion, like a tug-of-war. Motion-balance mechanisms maintain this constant gap by balancing one motion with another motion, like two dancers moving in unison.
- **Self-balancing opamp circuits:** all self-balancing operational amplifier circuits work on the principle of negative feedback maintaining a nearly zero differential input voltage to the opamp. Making the "simplifying assumption" that the opamp's differential input voltage is exactly zero assists in circuit analysis, as does the assumption that the input terminals draw negligible current.
- **Amplification:** the control of a relatively large signal by a relatively small signal. Relevant to the role of loop controllers exerting influence over a process variable at the command of a measurement signal. In behaving as amplifiers, loop controllers may oscillate if certain criteria are met.
- **Barkhausen criterion:** is overall loop gain is unity (1) or greater, and phase shift is 360° , the loop will sustain oscillations. Relevant to control system stability, explaining why the loop will "cycle" (oscillate) if gain is too high.
- **Time constant:** (τ), defined as the amount of time it takes a system to change 63.2% of the way from where it began to where it will eventually stabilize. The system will be within 1%

of its final value after 5 time constants' worth of time has passed (5τ). Relevant to process control loops, where natural lags contribute to time constants, usually of multiple order.

References

“FOUNDATION Fieldbus”, document L454 EN, Samson AG, Frankfurt, Germany, 2000.

“Identification and Description of Instrumentation, Control, Safety, and Information Systems and Components Implemented in Nuclear Power Plants”, EPRI, Palo Alto, CA: 2001. 1001503.

Kernighan, Brian W. and Ritchie, Dennis M., *The C Programming Language*, Bell Telephone Laboratories, Incorporated, Murray Hill, NJ, 1978.

Lavigne, John R., *Instrumentation Applications for the Pulp and Paper Industry*, The Foxboro Company, Foxboro, MA, 1979.

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Control Volume II*, Third Edition, CRC Press, Boca Raton, FL, 1999.

Mollenkamp, Robert A., *Introduction to Automatic Process Control*, Instrument Society of America, Research Triangle Park, NC, 1984.

“Moore 353 Process Automation Controller User's Manual”, document UM353-1, Revision 11, Siemens Energy and Automation, 2003.

Shinskey, Francis G., *Energy Conservation through Control*, Academic Press, New York, NY, 1978.

Shinskey, Francis G., *Process-Control Systems – Application / Design / Adjustment*, Second Edition, McGraw-Hill Book Company, New York, NY, 1979.

“SPEC 200 Systems”, technical information document TI 200-100, Foxboro, 1980.

“SPEC 200 System Configuration”, technical information document TI 200-105, Foxboro, January 1975.

“SPEC 200 System Wiring”, technical information document TI 200-260, Foxboro, 1972.

Ziegler, J. G., and Nichols, N. B., “Optimum Settings for Automatic Controllers”, *Transactions of the American Society of Mechanical Engineers* (ASME), Volume 64, pages 759-768, Rochester, NY, November 1942.

Chapter 30

Process dynamics and PID controller tuning

To *tune* a feedback control system means to adjust parameters in the controller to achieve robust control over the process. “Robust” in this context is usually defined as stability of the process variable despite changes in load, fast response to changes in setpoint, minimal oscillation following either type of change, and minimal offset (error between setpoint and process variable) over time.

“Robust control” is far easier to define than it is to achieve. With PID (Proportional-Integral-Derivative) control being the most common feedback control algorithm used in industry, it is important for all instrumentation practitioners to understand how to tune these controllers effectively and with a minimum investment of time.

Different types of processes, having different dynamic (time-dependent) behaviors, require different levels of proportional, integral, and derivative control action to achieve stability and robust response. It is therefore imperative for anyone seeking to tune a PID controller to understand the dynamic nature of the process being controlled. For this reason, the chapter begins with an exploration of common process characteristics before introducing techniques useful in choosing practical P, I, and D tuning parameter values.

30.1 Process characteristics

Perhaps the most important rule of controller tuning is to *know the process before attempting to adjust the controller's tuning*. Unless you adequately understand the nature of the process you intend to control, you will have little hope in actually controlling it well. This section of the book is dedicated to an investigation of different process characteristics and how to identify each.

Quantitative PID tuning methods (see section 30.3 beginning on page 2457) attempt to map the characteristics of a process so good PID parameters may be chosen for the controller. The goal of this section is for you to understand various process types by observation and qualitative analysis so you may comprehend why different tuning parameters are necessary for each type, rather than mindlessly following a step-by-step PID tuning procedure.

The three major classifications of process response are *self-regulating*, *integrating*, and *runaway*. Each of these process types is defined by its response to a step-change in the manipulated variable (e.g. control valve position or state of some other final control element). A “self-regulating” process responds to a step-change in the final control element’s status by settling to a new, stable value. An “integrating” process responds by ramping either up or down at a rate proportional to the magnitude of the final control element’s step-change. Finally, a “runaway” process responds by ramping either up or down at a rate that increases over time, headed toward complete instability without some form of corrective action from the controller.

Self-regulating, integrating, and runaway processes have very different control needs. PID tuning parameters that may work well to control a self-regulating process, for example, will *not* work well to control an integrating or runaway process, no matter how similar any of the other characteristics of the processes may be¹. By first identifying the characteristics of a process, we may draw some general conclusions about the P, I, and D setting values necessary to control it well.



Perhaps the best method for testing a process to determine its natural characteristics is to place the controller in *manual mode* and introduce a step-change to the controller output signal. It is critically important that the loop controller be in manual mode whenever process characteristics are being explored. If the controller is left in the automatic mode, the response seen from the process to a setpoint or load change will be partly due to the natural characteristics of the process itself *and* partly due to the corrective action of the controller. The controller’s corrective action thus interferes with our goal of exploring process characteristics. By placing the controller in “manual” mode, we turn off its corrective action, effectively removing its influence by breaking the feedback loop between process and controller, controller and process. In manual mode, the response we see from the process to an output (manipulated variable) or load change is *purely* a function of the natural process dynamics, which is precisely what we wish to discern.

A test of process characteristics with the loop controller in manual mode is often referred to as an *open-loop* test, because the feedback loop has been “opened” and is no longer a complete loop. Open-loop tests are the fundamental diagnostic technique applied in the following subsections.

¹To illustrate, self-regulating processes require significant integral action from a controller in order to avoid large offsets between PV and SP, with minimal proportional action and no derivative action. Integrating processes, in contrast, may be successfully controlled primarily on proportional action, with minimal integral action to eliminate offset. Runaway processes absolutely require derivative action for dynamic stability, but derivative action alone is not enough: some integral action will be necessary to eliminate offset. Even if knowledge of a process’s dominant characteristic does not give enough information for us to quantify P, I, or D values, it will tell us which tuning constant will be most important for achieving stability.

30.1.1 Self-regulating processes

If a liquid flow-control valve is opened in a step-change fashion, flow through the pipe tends to self-stabilize at a new rate very quickly. The following illustration shows a typical liquid flow-control installation, with a process trend showing the flow response following a manual-mode (also known as “open-loop”) step-change in valve position:



The defining characteristic of a self-regulating process is its inherent ability to settle at a new

process variable value without any corrective action on the part of the controller. In other words, a self-regulating process will exhibit a unique process variable value for each possible output (valve) value. The inherently fast response of a liquid flow control process makes its self-regulating nature obvious: the self-stabilization of flow generally takes place within a matter of seconds following the valve's motion. Many other processes besides liquid flow are self-regulating as well, but their slow response times require patience on the part of the observer to tell that the process will indeed self-stabilize following a step-change in valve position.

A corollary to the principle of self-regulation is that a unique output value *will be required* to achieve a new process variable value. For example, to achieve a greater flow rate, the control valve must be opened further and held at that further-open position for as long as the greater flow rate is desired. This presents a fundamental problem for a proportional-only controller. Recall the formula for a proportional-only controller, defining the output value (m) by the error (e) between process variable and setpoint multiplied by the gain (K_p) and added to the bias (b):

$$m = K_p e + b$$

Where,

m = Controller output

e = Error (difference between PV and SP)

K_p = Proportional gain

b = Bias

Suppose we find the controller in a situation where there is no error ($PV = SP$), and the flow rate is holding steady at some value. If we then increase the setpoint value (calling for a greater flow rate), the error will increase, driving the valve further open. As the control valve opens further, flow rate naturally increases to match. This increase in process variable drives the error back toward zero, which in turn causes the controller to decrease its output value back toward where it was before the setpoint change. However, the error can never go all the way back to zero because if it did, the valve would return to its former position, and that would cause the flow rate to self-regulate back to its original value before the setpoint change was made. What happens instead is that the control valve begins to close as flow rate increases, and eventually the process finds some equilibrium point where the flow rate is steady at some value less than the setpoint, creating just enough error to drive the valve open just enough to maintain that new flow rate. Unfortunately, due to the need for an error to exist, this new flow rate will fall shy of our setpoint. We call this error *proportional-only offset*, or *droop*, and it is an inevitable consequence of a proportional-only controller attempting to control a self-regulating process.

For any fixed bias value, there will be only one setpoint value that is perfectly achievable for a proportional-only controller in a self-regulating process. Any other setpoint value will result in some degree of offset in a self-regulating process. If dynamic stability is more important than absolute accuracy (zero offset) in a self-regulating process, a proportional-only controller may suffice. A great many self-regulating processes in industry have been and still are controlled by proportional-only controllers, despite some inevitable degree of offset between PV and SP.

The amount of offset experienced by a proportional-only controller in a self-regulating process may be minimized by increasing the controller's gain. If it were possible to increase the gain of a proportional-only controller to infinity, it would be able to achieve any setpoint desired with zero offset! However, there is a practical limit to the extent we may increase the gain value, and that

limit is *oscillation*. If a controller is configured with too much gain, the process variable will begin to oscillate over time, never stabilizing at any value at all, which of course is highly undesirable for any automatic control system. Even if the gain is not great enough to cause sustained oscillations, excessive values of gain will still cause problems by causing the process variable to oscillate with decreasing amplitude for a period of time following a sudden change in either setpoint or load. Determining the optimum gain value for a proportional-only controller in a self-regulating process is, therefore, a matter of compromise between excessive offset and excessive oscillation.

Recall that the purpose of integral (or “reset”) control action was the elimination of offset. Integral action works by ramping the output of the controller at a rate determined by the magnitude of the offset: the greater the difference between PV and SP for an integral controller, the faster that controller’s output will ramp over time. In fact, the output will stabilize at some value *only* if the error is diminished to zero ($PV = SP$). In this way, integral action works tirelessly to eliminate offset.

It stands to reason then that a self-regulating process *absolutely requires* some amount of integral action in the controller in order to achieve zero offset for all possible setpoint values. The more aggressive (faster) a controller’s integral action, the sooner offset will be eliminated. Just how much integral action a self-regulating process can tolerate depends on the magnitudes of any time lags in the system. The faster a process’s natural response is to a manual step-change in controller output, the better it will respond to aggressive integral controller action once the controller is placed in automatic mode. Aggressive integral control action in a slow process, however, will result in oscillation due to integral wind-up².

It is not uncommon to find self-regulating processes being controlled by *integral-only* controllers. An “integral-only” process controller is an instrument lacking proportional or derivative control modes. Liquid flow control is a nearly ideal candidate process for integral-only control action, due to its self-regulating and fast-responding nature.



Summary:

- Self-regulating processes are characterized by their natural ability to stabilize at a new process variable value following changes in the control element value or load(s).
- Self-regulating processes *absolutely require* integral controller action to eliminate offset between process variable and setpoint, because only integral action is able to create a different controller output value once the error returns to zero.
- Faster integral controller action results in quicker elimination of offset.
- The amount of integral controller action tolerable in a self-regulating process depends on the degree of time lag in the system. Too much integral action will result in oscillation, just like too much proportional control action.

²Recall that wind-up is what happens when integral action “demands” more from a process than the process can deliver. If integral action is too aggressive for a process (i.e. fast integral controller action in a process with slow time lags), the output will ramp too quickly, causing the process variable to overshoot setpoint which then causes integral action to wind the other direction. As with proportional action, too much integral action will cause a self-regulating process to oscillate.

30.1.2 Integrating processes

A good example of an integrating process is liquid level control, where either the flow rate of liquid into or out of a vessel is constant and the other flow rate varies. If a control valve is opened in a step-change fashion, liquid level in the vessel ramps at a rate proportional to the difference in flow rates in and out of the vessel. The following illustration shows a typical liquid level-control installation, with a process trend showing the level response to a step-change in valve position (with the controller in manual mode, for an “open-loop” test):



It is critically important to realize that this ramping action of the process variable over time is a

characteristic of the process itself, not the controller. When liquid flow rates in and out of a vessel are mis-matched, the liquid level within that vessel will change at a rate proportional to the difference in flow rates. The trend shown here reveals a fundamental characteristic of the process, not the controller (this should be obvious once it is realized that the step-change in output is something that would only ever happen with the controller in *manual* mode).

Mathematically, we may express the integrating nature of this process using calculus notation. First, we may express the *rate of change* of volume in the tank over time ($\frac{dV}{dt}$) in terms of the flow rates in and out of the vessel:

$$\frac{dV}{dt} = Q_{in} - Q_{out}$$

For example, if the flow rate of liquid going into the vessel was 450 gallons per minute, and the constant flow rate drawn out of the vessel was 380 gallons per minute, the volume of liquid contained within the vessel would increase over time at a rate equal to 70 gallons per minute: the difference between the in-flow and the out-flow rates.

Another way to express this mathematical relationship between flow rates and liquid volume in the vessel is to use the calculus function of *integration*:

$$\Delta V = \int_0^T (Q_{in} - Q_{out}) dt$$

The amount of liquid volume accumulated in the vessel (ΔV) between time 0 and time T is equal to the sum (\int) of the products (multiplication) of difference in flow rates in and out of the vessel ($Q_{in} - Q_{out}$) during infinitesimal increments of time (dt).

In the given scenario of a liquid level control system where the out-going flow is held constant, this means the level will be stable only at one in-coming flow rate (where $Q_{in} = Q_{out}$). At any other controlled flow rate, the level will either be increasing over time or decreasing over time.

This process characteristic perfectly matches the characteristic of a proportional-only controller, where there is one unique output value when the error is zero ($PV = SP$). We may illustrate this by performing a “thought experiment” on the liquid level-control process shown earlier having a constant draw out the bottom of the vessel. Imagine this process controlled by a proportional-only controller in automatic mode, with the bias value (b) of the controller set to the exact value needed by the control valve to make in-coming flow exactly equal to the constant out-going flow (draw). This means that when the process variable is precisely equal to setpoint ($PV = SP$), the in-flow will match the out-flow and therefore the liquid level will hold constant. If now an operator were to increase the setpoint value (with the controller in automatic mode), it will cause the valve to open further, adding liquid at a faster rate to the vessel. The naturally integrating nature of the process will result in an increasing liquid level. As level increases, the amount of error in the controller decreases, causing the valve to approach its original (bias) position. When the level reaches the new setpoint, proportional-only action will return the controller output to its original (bias) value, thus returning the in-flow control valve to its original position. This return to the bias position makes in-flow once again equal to out-flow, and so the level remains constant at the new setpoint with absolutely no offset (“droop”) even though the controller only exhibits proportional action. Proportional-only offset exists in a self-regulating process because a new valve position is needed to achieve any new setpoint value. Proportional-only offset does not occur on setpoint changes in an integrating process because only one valve position is needed to hold the process variable constant at *any* setpoint value.

The more aggressive the controller's proportional action, the sooner the integrating process will reach new setpoints. Just how much proportional action (gain) an integrating process can tolerate depends on the magnitudes of any time lags in the system as well as the magnitude of noise in the process variable signal. Any process system with time lags will oscillate if the controller has sufficient gain. Noise is a problem because proportional action directly reproduces process variable noise on the output signal: too much gain, and just a little bit of PV noise translates into a control valve whose stem position constantly jumps around.


Purely integrating processes do not require integral control action to eliminate offset as is the case with self-regulating processes, following a setpoint change. The natural integrating action of the process eliminates offset that would otherwise arise from setpoint changes. More than that, the presence of any integral action in the controller will actually force the process variable to overshoot setpoint following a setpoint change in a purely integrating process! Imagine a controller with integral action responding to a step-change in setpoint for the liquid level control process shown earlier. As soon as an error develops, the integral action will begin "winding up" the output value, forcing the valve to open more than proportional action alone would demand. By the time the liquid level reaches the new setpoint, the valve will have reached a position greater than where it originally was before the setpoint change³, which means the liquid level will *not* stop rising when it reaches setpoint, but in fact will overshoot setpoint. Only after the liquid level has spent sufficient time above setpoint will the integral action of the controller "wind" back down to its previous level, allowing the liquid level to finally achieve the new setpoint.

This is not to say that integral control action is completely unnecessary in integrating processes – far from it. If the integrating process is subject to *load* changes, only integral action can return the PV back to the SP value (eliminate offset). Consider, in our level control example, if the out-going flow rate were to change. Now, a new valve position will be required to achieve stable (unchanging) level in the vessel. A proportional-only controller is able to generate a new valve position *only* if an error develops between PV and SP. Without at least some degree of integral action configured in the controller, that error will persist indefinitely. Or consider if the liquid supply pressure upstream of the control valve were to change, resulting in a different rate of incoming flow for the same valve stem position as before. Once again, the controller would have to generate a different output value to compensate for this process change and stabilize liquid level, and the only way a proportional-only controller could do that is to let the process variable drift a bit from setpoint (the definition of an error or offset).

The example of an integrating process used here is just one of many possible processes where we are dealing with either a *mass balance* or an *energy balance* problem. "Mass balance" is the accounting of all mass into and out of a process. Since the Law of Mass Conservation states the impossibility of mass creation or destruction, all mass into and out of a process must be accounted for. If the mass flow rate into a process does not equal the mass flow rate out of a process, the process must be either gaining or losing an internal store of mass. The same may be said for energy: all energy flowing into and out of a process must be accounted for, since the Law of Energy Conservation

³In a proportional-only controller, the output is a function of error (PV – SP) and bias. When PV = SP, bias alone determines the output value (valve position). However, in a controller with integral action, the zero-offset output value is determined by *how long* and *how far* the PV has previously strayed from SP. In other words, there is no fixed bias value anymore. Thus, the output of a controller with integral action will *not* return to its previous value once the new SP is reached. In a purely integrating process, this means the PV will *not* reach stability at the new setpoint, but will continue to rise until all the "winding up" of integral action is un-done.

states the impossibility of energy creation or destruction. If the energy flow rate (input power) into a process does not equal the energy flow rate (output power) out of a process, the process must be either gaining or losing an internal store of energy.





Common examples of integrating processes include the following:

- Liquid level control – *mass balance* – when the flow of liquid either into or out of a vessel is manipulated, and the other flows in or out of the vessel are constant
- Gas pressure control – *mass balance* – when the flow of gas either into or out of a vessel is manipulated, and the other flows in or out of the vessel are constant
- Storage bin level control – *mass balance* – when the conveyor feed rate into the bin is manipulated, and the draw from the bin is constant
- Temperature control – *energy balance* – when the flow of heat into or out of a process is manipulated, and all other heat flows are constant
- Speed control – *energy balance* – when the force (linear) or torque (angular) applied to a mass is manipulated, and all other loads are constant in force or torque

In a self-regulating process, the control element (valve) exerts control over *both* the in-flow and the out-flow of either mass or energy. In the previous subsection, where liquid flow control was the process example, the mass balance consisted of liquid flow into the valve and liquid flow out of the valve. Since the piping was essentially a “series” path for an incompressible fluid, where input flow must equal output flow at any given time, mass in and mass out were *guaranteed* to be in a state of balance, with one valve controlling both. This is why a change in valve position resulted in an almost immediate change and re-stabilization of flow rate: the valve exerts immediate control over both the incoming and the outgoing flow rates, with both in perfect balance. Therefore, nothing “integrates” over time in a liquid flow control process because there can never be an imbalance between in-flow and out-flow.

In an integrating process, the control element (valve) exerts control over *either* the in-flow *or* the out-flow of mass or energy, but never both. Thus, changing valve position in an integrating process causes an imbalance of mass flow and/or energy flow, resulting in the process variable ramping over time as either mass or energy accumulates in (or depletes from) the process.

Our “simple” example of an integrating (level-control) process becomes a bit more complicated if the outgoing flow depends on level, as is the case with a gravity-drained vessel where the outgoing flow is a function of liquid level in the vessel rather than being fixed at a constant rate as it was in the previous example:



If we subject the control valve to a manual step-change increase, the flow rate of liquid into the vessel immediately increases. This causes an imbalance of incoming and outgoing flow, resulting in the liquid level rising over time. As level rises, however, increasing hydrostatic pressure across the

manual valve at the vessel outlet causes the outgoing flow rate to increase. This causes the mass imbalance rate to be less than it was before, resulting in a decreased integration rate (rate of level rise). Thus, the liquid level still rises, but at a slower and slower rate as time goes on. Eventually, the liquid level will become high enough that the pressure across the manual valve forces a flow rate out of the vessel equal to the flow rate into the vessel. At this point, with matched flow rates, the liquid level stabilizes with no corrective action from the controller (remember, the step-change in output was made in manual mode!). Note the final result of letting the outgoing flow be a function of liquid level: *what used to be an integrating process has now become a self-regulating process*, albeit one with a substantial lag time.

Many processes ideally categorized as integrating actually behave in this manner. Although the manipulated variable may control the flow rate into or out of a process, the other flow rates often change with the process variable. Returning to our list of integrating process examples, we see how a PV-variable load in each case can make the process self-regulate:

- Liquid level control – *mass balance* – if the in-flow naturally decreases as liquid level rises and/or the out-flow naturally increases as liquid level rises, the vessel’s liquid level will tend to self-regulate instead of integrate
- Gas pressure control – *mass balance* – if in-flow naturally decreases as pressure rises and/or the out-flow naturally increases as pressure rises, the vessel’s pressure will tend to self-regulate instead of integrate
- Storage bin level control – *mass balance* – if the draw from the bin increases with bin level (greater weight pushing material out at a faster rate), the bin’s level will tend to self-regulate instead of integrate
- Temperature control – *energy balance* – if the process naturally loses heat at a faster rate as temperature increases and/or the process naturally takes in less heat as temperature rises, the temperature will tend to self-regulate instead of integrate
- Speed control – *energy balance* – if drag forces on the object increase with speed (as they usually do for any fast-moving object), the speed will tend to self-regulate instead of integrate

We may generalize all these examples of integrating processes turned self-regulating by noting the one aspect common to all of them: some natural form of *negative feedback* exists internally to bring the system back into equilibrium. In the mass-balance examples, the physics of the process ensure a new balance point will eventually be reached because the in-flow(s) and/or out-flow(s) naturally change in ways that oppose any change in the process variable. In the energy-balance examples, the laws of physics again conspire to ensure a new equilibrium because the energy gains and/or losses naturally change in ways that oppose any change in the process variable. The presence of a control system is, of course, the ultimate example of negative feedback working to stabilize the process variable. However, the control system may not be the *only* form of negative feedback at work in a process. All self-regulating processes are that way because they intrinsically possess some degree of negative feedback acting as a sort of natural, proportional-only control system.


This one detail completely alters the fundamental characteristic of a process from integrating to self-regulating, and therefore changes the necessary controller parameters. Self-regulation guarantees at least some integral controller action is *necessary* to attain new setpoint values. A purely integrating process, by contrast, requires no integral controller action at all to achieve new setpoints, and in fact is *guaranteed* to suffer overshoot following setpoint changes if the controller is programmed with any integral action at all! Both types of processes, however, need some amount of integral action in the controller in order to recover from *load* changes.

Summary:


- Integrating processes are characterized by a ramping of the process variable in response to a step-change in the control element value or load(s).
- This integration occurs as a result of either *mass flow imbalance* or *energy flow imbalance* in and out of the process.
- Integrating processes are ideally controllable with proportional controller action alone.
- Integral controller action guarantees setpoint overshoot in a purely integrating process.
- Some integral controller action will be required in integrating processes to compensate for load changes.
- The amount of proportional controller action tolerable in an integrating process depends on the degree of time lag and process noise in the system. Too much proportional action will result in oscillation (time lags) and/or erratic control element motion (noise).
- An integrating process will become self-regulating if sufficient negative feedback is naturally introduced. This usually takes the form of loads varying with the process variable.

30.1.3 Runaway processes

A classic “textbook” example of a runaway process is an inverted pendulum: a vertical stick balanced on its end by moving the bottom side-to-side. Inverted pendula are typically constructed in a laboratory environment by fixing a stick to a cart by a pivot, then equipping the cart with wheels and a reversible motor to give it lateral control ability. A sensor (usually a potentiometer) detects the stick’s angle from vertical, reporting that angle to the controller as the process variable. The cart’s motor is the final control element:



The defining characteristic of a runaway process is its tendency to accelerate away from a condition of stability with no corrective action applied. Viewed on a process trend, a runaway process tends to respond as follows to an open-loop step-change:



A synonym for “runaway” is *negative self-regulation* or *negative lag*, because the process variable curve over time for a runaway process resembles the mathematical inverse of a self-regulating curve

with a lag time: it races away from the horizontal, while a self-regulating process variable draws closer and closer to the horizontal over time.

The “SegwayTM” personal transport device is a practical example of an inverted pendulum, with wheel motion controlled by a computer attempting to maintain the body of the vehicle in a vertical position. As the human rider leans forward, it causes the controller to spin the wheels with just the right amount of acceleration to maintain balance. There are many examples of runaway processes in motion-control applications, especially automated controls for vertical-flight vehicles such as helicopters and vectored-thrust aircraft such as the Harrier military fighter jet.

Some chemical reaction processes are runaway as well, especially *exothermic* (heat-releasing) reactions. Most chemical reactions increase in rate as temperature rises, and so exothermic reactions tend to accelerate with time (either becoming hotter or becoming colder) unless checked by some external influence. This poses a significant challenge to process control, as many exothermic reactions used to manufacture products must be temperature-controlled to ensure efficient production of the desired product. Off-temperature chemical reactions may not “favor” production of the desired products, producing unwanted byproducts and/or failing to efficiently consume the reactants. Furthermore, safety concerns usually surround exothermic chemical reactions, as no one wants their process to melt down or explode.

What makes a runaway process behave as it does is internal *positive feedback*. In the case of the inverted pendulum, gravity works to pull an off-center pendulum even farther off center, accelerating it until it falls down completely. In the case of exothermic chemical reactions, the direct relationship between temperature and reaction rate forms a positive feedback loop: the hotter the reaction, the faster it proceeds, releasing even more heat, making it even hotter. It should be noted that *endothermic* chemical reactions (absorbing heat rather than releasing heat) tend to be self-regulating for the exact same reason exothermic reactions tend to be runaway: reaction rate usually has a positive correlation with reaction temperature.

It is easy to demonstrate for yourself how challenging a runaway process can be to control. Simply try to balance a long stick vertically in the palm of your hand. You will find that the only way to maintain stability is to react swiftly to any changes in the stick’s angle – essentially applying a healthy dose of *derivative* control action to counteract any motion from vertical.

Fortunately, runaway processes are less common in the process industries. I say “fortunately” because these processes are notoriously difficult to control and usually pose more danger than inherently self-regulating processes. Many runaway processes are also nonlinear, making their behavior less intuitive to human operators.


Just as integrating processes may be forced to self-regulate by the addition of (natural) negative feedback, intrinsically runaway processes may also be forced to self-regulate given the presence of sufficient natural negative feedback. An interesting example of this is a pressurized water nuclear fission reactor.

Nuclear fission is a process by which the nuclei of specific types of atoms (most notably uranium-235 and plutonium-239) undergo spontaneous disintegration upon the absorption of an extra neutron, with the release of significant thermal energy and additional neutrons. A quantity of fissile material such as ^{235}U or ^{239}Pu is subjected to a source of neutron particle radiation, which initiates the fission process, releasing massive quantities of heat which may then be used to boil water into steam and drive steam turbine engines to generate electricity. The “chain reaction” of neutrons splitting fissile atoms, which then eject more neutrons to split more fissile atoms, is inherently exponential in nature. The more atoms split, the more neutrons are released, which then proceed to split even more atoms. The rate at which neutron activity within a fission reactor grows or decays over time is determined by the *multiplication factor*⁴, and this factor is easily controlled by the insertion of neutron-absorbing *control rods* into the reactor core.

Thus, a fission chain-reaction naturally behaves as an inverted pendulum. If the multiplication factor is greater than 1, the reaction grows exponentially. If the multiplication factor is less than 1, the reaction dies exponentially. In the case of a nuclear weapon, the desired multiplication factor is as large as physically possible to ensure explosive reaction growth. In the case of an operating nuclear power plant, the desired multiplication factor is exactly 1 to ensure stable power generation.

⁴When a nucleus of uranium or plutonium undergoes fission (“splits”), it releases more neutrons capable of splitting additional uranium or plutonium nuclei. The ratio of new nuclei “split” versus old nuclei “split” is the multiplication factor. If this factor has a value of one (1), the chain reaction will sustain at a constant power level, with each new generation of atoms “split” equal to the number of atoms “split” in the previous generation. If this multiplication factor exceeds unity, the rate of fission will increase over time. If the factor is less than one, the rate of fission will decrease over time. Like an inverted pendulum, the chain reaction has a tendency to “fall” toward infinite activity or toward no activity, depending on the value of its multiplication factor.

A simplified diagram of a pressurized-water reactor (PWR) is shown here:



Water under high pressure (too high of pressure to boil) circulates through the reactor vessel, carrying heat away from the nuclear core, then transferring the heat energy to a heat exchanger (“steam generator”) where a second water loop is allowed to boil into steam and drive turbine engines to spin electrical generators. Control rods inserted into the core by linear actuators adjust the multiplication factor of the reactor.

If the multiplication factor of a fission reactor were solely controlled by the positions of these control rods, it would be a classic “runaway” process, with the reactor’s power level tending to increase toward infinity or decrease toward zero if the rods were at any position other than one yielding a multiplication factor of precisely unity (1). This would make nuclear reactors extremely difficult (if not impossible) to safely control. Fortunately, there are ways to engineer negative feedback directly into the design of the reactor core so that neutron activity *naturally* self-stabilizes without active control rod action. In water-cooled reactors, the water itself achieves this goal. Pressurized water plays a dual role in a fission reactor: it not only transfers heat out of the reactor core and into a boiler to produce steam, but it also offsets the multiplication factor inversely proportional to temperature. As the reactor core heats up, the water’s density changes, affecting the probability⁵ of neutrons being captured by fissile nuclei. This is called a *negative temperature*

⁵The mechanism by which this occurs varies with the reactor design, and is too detailed to warrant a full explanation here. In pressurized light-water reactors – the dominant design in the United States of America – this action occurs

coefficient for the reactor, and it forces the otherwise runaway process of nuclear fission to become self-regulating.

With this self-regulating characteristic in effect, control rod position essentially determines the reactor's steady-state temperature. The further the control rods are withdrawn from the core, the hotter the core will run. The cooling water's natural negative temperature coefficient prevents the fission reaction from "running away" either to destruction or to shutdown.

Some nuclear fission reactor designs are capable of "runaway" behavior, though. The ill-fated reactor at Chernobyl (Ukraine, Russia) was of a design where its power output could "run away" under certain operating conditions, and that is exactly what happened on April 26, 1986. The Chernobyl reactor used solid graphite blocks as the main neutron-moderating substance, and as such its cooling water did not provide enough natural negative feedback to overcome the intrinsically runaway characteristic of nuclear fission. This was especially true at low power levels where the reactor was being tested on the day of the accident. A combination of poor management decisions, unusual operating conditions, and unstable design characteristics led to the reactor's destruction with massive amounts of radiation released into the surrounding environment. It stands at the time of this writing as the world's worst nuclear accident⁶.

Summary:


- Runaway processes are characterized by an exponential ramping of the process variable in response to a step-change in the control element value or load(s).
- This "runaway" occurs as a result of some form of *positive feedback* happening inside the process.
- Runaway processes cannot be controlled with proportional or integral controller action alone, and always requires derivative action for stability.
- Some integral controller action will be required in runaway processes to compensate for load changes.
- A runaway process will become self-regulating if sufficient negative feedback is naturally introduced, as is the case with water-moderated fission reactors.

due to the water's ability to *moderate* (slow down) the velocity of neutrons. Slow neutrons have a greater probability of being "captured" by fissile nuclei than fast neutrons, and so the water's moderating ability will have a direct effect on the reactor core's multiplication factor. As a light-water reactor core increases temperature, the water becomes less dense and therefore less effective at moderating (slowing down) fast neutrons emitted by "splitting" nuclei. These fast(er) neutrons then "miss" the nuclei of atoms they would have otherwise split, effectively reducing the reactor's multiplication factor without any need for regulatory control rod motion. The reactor's power level therefore self-stabilizes as it warms, rather than "running away" to dangerously high levels, and may thus be classified as a *self-regulating* process.

⁶Discounting, of course, the intentional discharge of nuclear weapons, whose sole design purpose is to self-destruct in a "runaway" chain reaction.

30.1.4 Steady-state process gain


When we speak of a controller's *gain*, we refer to the aggressiveness of its proportional control action: the ratio of output change to input change. However, we may go a step further and characterize each component within the feedback loop as having its own gain (a ratio of output change to input change):



The gains intrinsic to the measuring device (transmitter), final control device (e.g. control valve), and the process itself are all important in helping to determine the necessary controller gain to achieve robust control. The greater the combined gain of transmitter, process, and valve, the less gain is needed from the controller. The less combined gain of transmitter, process, and valve, the more gain will be needed from the controller. This should make some intuitive sense: the more “responsive” a process appears to be, the less aggressive the controller needs to be in order to achieve stable control (and vice-versa).

These combined gains may be empirically determined by means of a simple test performed with the controller in manual mode, also known as an “open-loop” test. By placing the controller in manual mode (and thus disabling its automatic correction of process changes) and adjusting the output signal by some fixed amount, the resulting change in process variable may be measured and compared. If the process is self-regulating, a definite ratio of PV change to controller output change may be determined.

For instance, examine this process trend graph showing a manual “step-change” and process variable response:



Here, the output step-change is 10% of scale, while the resulting process variable step-change is about 7.5%. Thus, the “gain” of the process⁷ (together with transmitter and final control element) is approximately 0.75, or 75% ($\text{Gain} = \frac{7.5\%}{10\%}$). Incidentally, it is irrelevant that the PV steps *down* in response to the controller output stepping *up*. All this means is the process is reverse-responding, which necessitates *direct* action on the part of the controller in order to achieve negative feedback. When we calculate gains, we usually ignore directions (mathematical signs) and speak in terms of absolute values.


We commonly refer to this gain as the *steady-state gain* of the process, because the determination of gain is made after the PV settles to its self-regulating value.

Since from the controller’s perspective the individual gains of transmitter, final control element, and physical process meld into one over-all gain value, the process may be made to appear more or less responsive (more or less steady-state gain) just by altering the gain of the transmitter and/or the gain of the final control element.

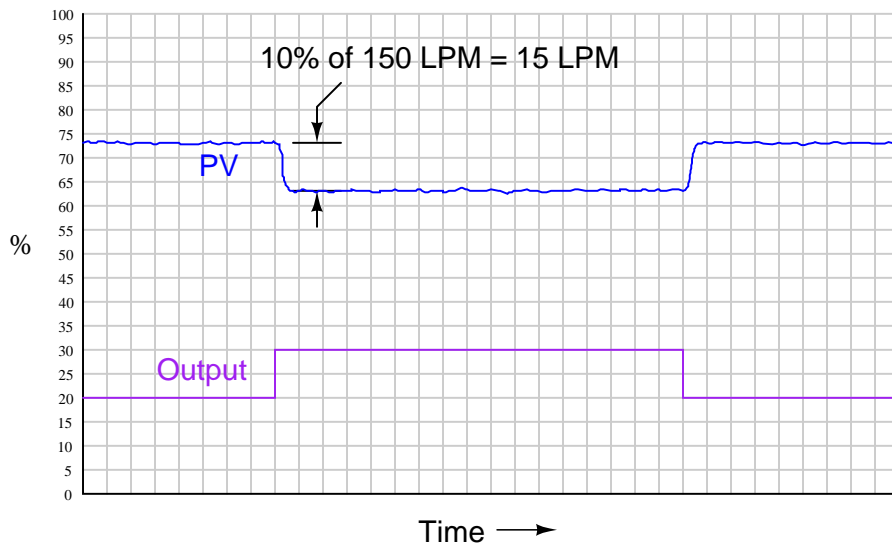
Consider, for example, if we were to reduce the span of the transmitter in this process. Suppose this was a flow control process, with the flow transmitter having a calibrated range of 0 to 200 liters per minute (LPM). If a technician were to re-range the transmitter to a new range of 0 to 150 LPM, what effect would this have on the apparent process gain?

⁷The general definition of gain is the ratio of output change over input change ($\frac{\Delta \text{Out}}{\Delta \text{In}}$). Here, you may have noticed we calculate process gain by dividing the process variable change (7.5%) by the controller output change (10%). If this seems “inverted” to you because we placed the *output* change value in the denominator of the fraction instead of the numerator, you need to keep in mind the perspective of our gain measurement. We are not calculating the gain of the controller, but rather the gain of the *process*. Since the output of the controller is the “input” to the process, it is entirely appropriate to refer to the 10% manual step-change as the change of *input* when calculating process gain.

To definitively answer this question, we must re-visit the process trend graph for the old calibrated range:



We see here that the 7.5% PV step-change equates to a change of 15 LPM given the flow transmitter's span of 200 LPM. However, if a technician re-ranges the flow transmitter to have just three-quarters that amount of span (150 LPM), the exact same amount of output step-change will *appear* to have a more dramatic effect on flow, even though the physical response of the process has the same as it was before:



From the controller's perspective – which only “knows” percent⁸ of signal range – the process gain appears to have increased from 0.75 to 1, with nothing more than a re-ranging of the transmitter. Since the process is now “more responsive” to controller output signals than it was before, there may be a tendency for the loop to oscillate in automatic mode even if it did not oscillate previously with the old transmitter range. A simple fix for this problem is to decrease the controller's gain by the same factor that the process gain increased: we need to make the controller's gain $\frac{3}{4}$ what it was before, since the process gain is now $\frac{4}{3}$ what it was before.


The exact same effect occurs if the final control element is re-sized or re-ranged. A control valve that is replaced with one having a different C_v value, or a variable-frequency motor drive that is given a different speed range for the same 4-20 mA control signal, are two examples of final control element changes which will result in different overall gains. In either case, a given change in controller output signal percentage results in a different amount of influence on the process thanks to the final control element being more or less influential than it was before. Re-tuning of the controller may be necessary in these cases to preserve robust control.

If and when re-tuning is needed to compensate for a change in loop instrumentation, all control modes should be proportionately adjusted. This is automatically done if the controller uses the *Ideal* or *ISA* PID equation, or if the controller uses the *Series* or *Interacting* PID equation⁹. All that needs to be done to an Ideal-equation controller in order to compensate for a change in process gain is to change that controller's proportional (P) constant setting. Since this constant directly affects all terms of the equation, the other control modes (I and D) will be adjusted along with the proportional term. If the controller happens to be executing the *Parallel* PID equation, you will have to manually alter all three constants (P, I, and D) in order to compensate for a change in process gain.

⁸While this is true of analog-signal transmitters, it is not necessarily true of digital-signal transmitters such as Fieldbus or wireless (digital radio). The reason for this distinction is that in a digital-signal transmitter, the reported process variable value is scaled in engineering units rather than percent. Applied to this case, if the flow transmitter gets re-ranged from 0-200 LPM to 0-150 LPM, the controller sees no change in process gain because a change of 10 LPM is still reported as a change in 10 LPM regardless of the transmitter's range.

⁹For more information on different PID equations, refer to Section 29.10 beginning on page 2314.

A very important aspect of process gain is how *consistent* it is over the entire measurement range. It is entirely possible (and in fact very likely) that a process may be more responsive (have higher gain) in some areas of control than in others. Take for instance this hypothetical trend showing process response to a series of manual-mode step-changes:



Note how the PV changes about 5% for the first 5% step-change in output, corresponding to a process gain of 1. Then, the PV changes about 7.5% for the next 5% output step-change, for a process gain of 1.5. The final increasing 5% step-change yields a PV change of about 12.5%, a process gain of 2.5. Clearly, the process being controlled here is not equally responsive throughout the measurement range. This is a concern to us in tuning the PID controller because any set of tuning constants that work well to control the process around a certain setpoint may not work as well if the setpoint is changed to a different value, simply because the process may be more or less responsive at that different process variable value.


Inconsistent process gain is a problem inherent to many different process types, which means it is something you will need to be aware of when investigating a process prior to tuning the controller. The best way to reveal inconsistent process gain is to perform a series of step-changes to the controller output while in manual mode, “exploring” the process response throughout the safe range of operation.

Compensating for inconsistent process gain is much more difficult than merely detecting its presence. If the gain of the process continuously grows from one end of the range to the other (e.g. low gain at low output values and high gain at high output values, or vice-versa), a control valve with a different characteristic may be applied to counter-act the process gain.

If the process gain follows some pattern more closely related to PV value rather than controller output value, the best solution is a type of controller known as an *adaptive gain controller*. In an adaptive gain controller, the proportional setting is made to vary in a particular way as the process changes, rather than be a fixed constant set by a human technician or engineer.


30.1.5 Lag time

If a square-wave signal is applied to an RC passive integrator circuit, the output signal will appear to have a “sawtooth” shape, the crisp rising and falling edges of the square wave replaced by damped curves:





In a word, the output signal of this circuit *lags* behind the input signal, unable to keep pace with the steep rising and falling edges of the square wave.

Most mechanical and chemical processes exhibit a similar tendency: an “inertial” opposition to rapid changes. Even instruments themselves naturally¹⁰ damp sudden stimuli. We could have just as easily subjected a pressure transmitter to a series of pressure pulses resembling square waves, and witnessed the output signal exhibit the same damped response:



¹⁰It is also possible to *configure* many instruments to deliberately damp their response to input conditions. This is called *damping*, and it is covered in more detail in section 18.4 beginning on page 1265.

The gravity-drained level-control process highlighted in an earlier subsection exhibits a very similar response to a sudden change in control valve position:



For any particular flow rate into the vessel, there will be a final (self-regulating) point where the liquid level “wants” to settle¹¹. However, the liquid level does not *immediately* achieve that new

¹¹ Assuming a constant discharge valve position. If someone alters the hand valve’s position, the relationship between incoming flow rate and final liquid level changes.

level if the control valve jumps to some new position, owing to the “capacity” of the vessel and the dynamics of gravity flow.

Any physical behavior exhibiting the same “settling” behavior over time may be said to illustrate a *first-order lag*. A classic “textbook” example of a first-order lag is the temperature of a cup of hot liquid, gradually equalizing with room temperature. The liquid’s temperature drops rapidly at first, but then slows its approach to ambient temperature as time progresses. This natural tendency is described by *Newton’s Law of Cooling*, mathematically represented in the form of a *differential equation* (an equation containing a variable along with one or more of its derivatives). In this case, the equation is a *first-order* differential equation, because it contains the variable for temperature (T) and the first derivative of temperature ($\frac{dT}{dt}$) with respect to time:

$$\frac{dT}{dt} = -k(T - T_{ambient})$$

Where,

T = Temperature of liquid in cup

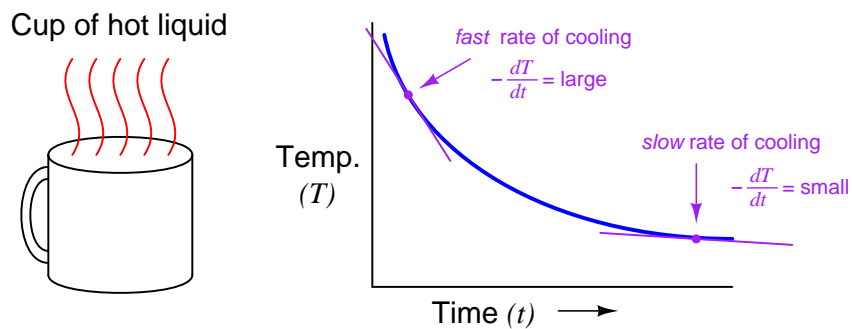
$T_{ambient}$ = Temperature of the surrounding environment

k = Constant representing the thermal conductivity of the cup

t = Time

All this equation tells us is that the rate of cooling ($\frac{dT}{dt}$) is directly proportional ($-k$) to the difference in temperature between the liquid and the surrounding air ($T - T_{ambient}$). The hotter the temperature, the faster the object cools (the faster rate of temperature fall):

Newton’s Cooling Law



The proportionality constant in this equation (k) represents how readily thermal energy escapes the hot cup. A cup with more thermal insulation, for example, would exhibit a smaller k value (i.e. the rate of temperature loss $\frac{dT}{dt}$ will be less for any given temperature difference between the cup and ambient $T - T_{ambient}$).

A general solution to this equation is as follows:

$$T = (T_{initial} - T_{final}) \left(e^{-\frac{t}{\tau}} \right) + T_{final}$$

Where,

T = Temperature of liquid in cup at time t

$T_{initial}$ = Starting temperature of liquid ($t = 0$)

T_{final} = Ultimate temperature of liquid (ambient)


e = Euler's constant

τ = "Time constant" of the system

This mathematical analysis introduces a descriptive quantity of the system: something called a *time constant*. The "time constant" of a first-order system is the amount of time necessary for the system to come to within 36.8% (e^{-1}) of its final value (i.e. the time required for the system to go 63.2% of the way from the starting point to its ultimate settling point: $1 - e^{-1}$). After two time-constants' worth of time, the system will have come to within 13.5% (e^{-2}) of its final value (i.e. gone 86.5% of the way: $1 - e^{-2}$); after three time-constants' worth of time, to within 5% (e^{-3}) of the final value, (i.e. gone 95% of the way: $1 - e^{-3}$). After five time-constants' worth of time, the system will be within 1% (e^{-5} , rounded to the nearest whole percent) of its final value, which is often close enough to consider it "settled" for most practical purposes.

Time	Percent of final value	Percent change remaining
0	0.000%	100.000%
τ	63.212%	36.788%
2τ	86.466%	13.534%
3τ	95.021%	4.979%
4τ	98.168%	1.832%
5τ	99.326%	0.674%
6τ	99.752%	0.248%
7τ	99.909%	0.091%
8τ	99.966%	0.034%
9τ	99.988%	0.012%
10τ	99.995%	0.005%


The concept of a “time constant” may be shown in graphical form for both falling and rising variables:



One "time constant" (τ) is the amount of time required for the variable to change 63.2% of the way from its starting point to its ultimate (terminal) value

Students of electronics will immediately recognize this concept, since it is widely used in the analysis and application of capacitive and inductive circuits. However, you should recognize the fact that the concept of a “time constant” for capacitive and inductive electrical circuits is only one case of a more general phenomenon. Literally *any* physical system described by the same first-order differential equation may be said to have a “time constant.” Thus, it is perfectly valid for us to speak of a hot cup of coffee as having a time constant (τ), and to say that the coffee’s temperature will be within 1% of room temperature after five of those time constants have elapsed.

In the world of process control, it is more customary to refer to this as a *lag time* than as a *time constant*, but these are really interchangeable terms. The term “lag time” makes sense if we consider a first-order system *driven* to achieve a constant rate of change. For instance, if we subjected our RC circuit to a ramping input voltage rather than a “stepped” input voltage – such that the output ramped as well instead of passively settling at some final value – we would find that the amount of time separating equal input and output voltage values was equal to this time constant (in an RC circuit, $\tau = RC$):




Lag time is thus defined as the difference in time between when the process variable ramps to a certain value and when it *would have* ramped to that same value were it not for the existence of first-order lag in the system. The system’s output variable *lags behind* the ramping input variable by a fixed amount of time, regardless of the ramping rate. If the process in question is an RC circuit, the lag time will still be the product of ($\tau = RC$), just as the “time product” defined for a stepped input voltage. Thus, we see that “time constant” and “lag time” are really the exact same concept, merely manifesting in different forms as the result of two different input conditions (*stepped* versus *ramped*).

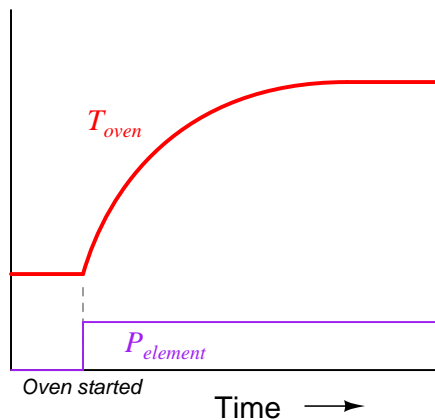
When an engineer or a technician describes a process being “fast” or “slow,” they generally refer to the magnitude of this lag time. This makes lag time very important to our selection of PID controller tuning values. Integral and derivative control actions in particular are sensitive to the amount of lag time in a process, since both those actions are time-based. “Slow” processes (i.e. process types having large lag times) cannot tolerate aggressive integral action, where the controller “impatiently” winds the output up or down at a rate that is too rapid for the process to respond to. Derivative action, however, is generally useful on processes having large lag times.

30.1.6 Multiple lags (orders)

Simple, self-regulating processes tend to be first-order: that is, they have only one mechanism of lag. More complicated processes often consist of multiple sub-processes, each one with its own lag time. Take for example a convection oven, heating a potato. Being instrumentation specialists in addition to cooks, we decide to monitor both the oven temperature and the potato temperature using thermocouples and remote temperature indicators:



The oven itself is a first-order process. If we graph its temperature over time as the heater power is suddenly stepped up to some fixed value¹², we will see a classic first-order response:




The potato forms another first-order process, absorbing heat from the air within the oven (heat transfer by convection), gradually warming up until its temperature (eventually) reaches that of

¹²We will assume here the heating element reaches its final temperature immediately upon the application of power, with no lag time of its own.

the oven¹³. From the perspective of the heating element to the oven air temperature, we have a first-order process. From the perspective of the heating element to the potato, however, we have a *second-order* process.

Intuition might lead you to believe that a second-order process is just like a first-order process – except slower – but that intuition would be wrong. Cascading two first-order lags creates a fundamentally different time dynamic. In other words, two first-order lags do not simply result in a *longer* first-order lag, but rather a *second-order* lag with its own unique characteristics.


If we superimpose a graph of the potato temperature with a graph of the oven temperature (once again assuming constant power output from the heating element, with no thermostatic control), we will see that the *shape* of this second-order lag is different. The curve now has an “S” shape, rather than a consistent downward concavity:



This, in fact, is one of the tell-tale signature of multiple lags in a process: an “S”-shaped curve rather than the characteristically abrupt initial rise of a first-order curve.

¹³Given the presence of water in the potato which turns to steam at 212 °F, things are just a bit more complicated than this, but let’s ignore the effects of water in the potato for now!

Another tell-tale signature of multiple lags is that the lagging variable does not immediately reverse its direction of change following a reversal in the final control element signal. We can see this effect by cutting power to the heating element before either the oven air or potato temperatures have reached their final values:




Note how the air temperature trend *immediately* reverses direction following the cessation of power to the heating element, but how the potato temperature trend continues to rise for a short amount of time¹⁴ before reversing direction and cooling. Here, the contrast between first-order and second-order lag responses is rather dramatic – the second-order response is clearly not just a longer version of the first-order response, but rather something quite distinct unto itself.


This is why multiple-order lag processes have a greater tendency to *overshoot* their setpoints while under automatic control: the process variable exhibits a sort of “inertia” whereby it fails to switch directions simultaneously with the controller output.

¹⁴The amount of time the potato’s temperature will continue to rise following the down-step in heating element power is equal to the time it takes for the oven’s air temperature to equal the potato’s temperature. The reason the potato’s temperature keeps rising after the heating element turns off is because the air inside the oven is (for a short time) still hotter than the potato, and therefore the potato continues to absorb thermal energy from the air for a time following power-off.

If we were able to ramp the heater power at a constant rate and graph the heater element, air, and potato temperatures, we would clearly see the separate lag times of the oven and the potato as offsets in time at any given temperature:



As another example, let us consider the control of level in three cascaded, gravity-drained vessels:



From the perspective of the level transmitter on the last vessel, the control valve is driving a *third-order* process, with three distinct lags cascaded in series. This would be a challenging process to control, and not just because of the possibility of the intermediate vessels overflowing (since their levels are not being measured)!


When we consider the dynamic response of a process, we are usually concerned primarily with the physical process itself. However, the instruments attached to that process also influence lag

orders and lag times. As discussed in the previous subsection, almost every physical function exhibits some form of lag. Even the instruments we use to measure process variables have their own (usually very short) lag times. Control valves may have substantial lag times, measured in the tens of seconds for some large valves. Thus, a “slow” control valve exerting control over a first-order process effectively creates a second-order loop response. Thermowells used with temperature sensors such as thermocouples and RTDs can also introduce lag times into a loop (especially if the sensing element is not fully contacting the bottom of the well!).

This means it is nearly impossible to have a control loop with a purely first-order response. Many real loops come close to being first-order, but only because the lag time of the physical process swamps (dominates) the relatively tiny lag times of the instruments. For inherently fast processes such as liquid flow and liquid pressure control, however, the process response is so fast that even short time lags in valve positioners, transmitters, and other loop instruments significantly alter the loop’s dynamic character.

Multiple-order lags are relevant to the issue of PID loop tuning because they encourage oscillation. The more lags there are in a system, the more delayed and “detached” the process variable becomes from the controller’s output signal.

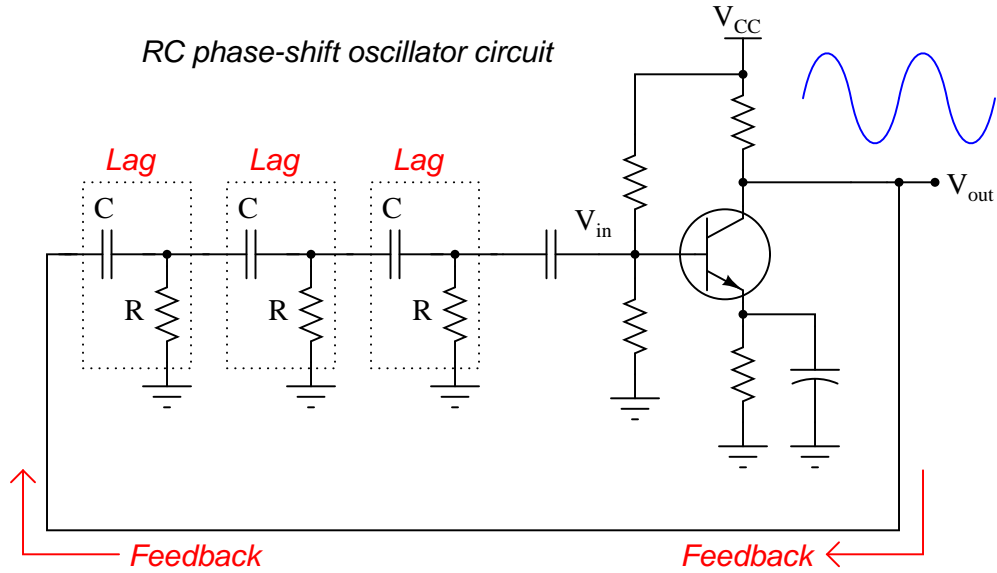
A system with lag time exhibits *phase shift* when driven by a sinusoidal stimulus: the outgoing waveform lags behind the input waveform by a certain number of degrees at one frequency. The exact amount of phase shift depends on frequency – the higher the frequency, the more phase shift (to a maximum of -90° for a first-order lag):



The phase shifts of multiple, cascaded lag functions (or processes, or physical effects) add up. This means each lag in a system contributes an additional negative *phase shift* to the loop. This can be detrimental to negative feedback, which by definition is a 180° phase shift. If sufficient lags exist in a system, the total loop phase shift may approach 360° , in which case the feedback becomes *positive* (regenerative): a necessary¹⁵ condition for oscillation.

¹⁵The so-called *Barkhausen criterion* for oscillation in a feedback system is that the total loop gain is at least unity

It is worthy to note that multiple-order lags are constructively applied in electronics when the express goal is to create oscillations. If a series of RC “lag” networks are used to feed the output of an inverting amplifier circuit back to its input with sufficient signal strength intact¹⁶, and those networks introduce another 180 degrees of phase shift, the total loop phase shift will be 360° (i.e. positive feedback) and the circuit will self-oscillate. This is called an *RC phase-shift oscillator* circuit:



The amplifier works just like a proportional-only process controller, with action set for negative feedback. The resistor-capacitor networks act like the lags inherent to the process being controlled. Given enough controller (amplifier) gain, the cascaded lags in the process (RC networks) create the perfect conditions for self-oscillation. The amplifier creates the first 180° of phase shift (being inverting in nature), while the RC networks collectively create the other 180° of phase shift to give a total phase shift of 360° (positive, or *regenerative* feedback).


In theory, the most phase shift a single RC network can create is -90° , but even that is not practical¹⁷. This is why more than two RC phase-shifting networks are required for successful operation of an RC phase-shift oscillator circuit.

(1) and the total loop phase shift is 360° .

¹⁶The conditions necessary for self-sustaining oscillations to occur is a total phase shift of 360° and a total loop gain of 1. Merely having positive feedback or having a total gain of 1 or more will not guarantee self-sustaining oscillations; both conditions must simultaneously exist. As a measure of how close any feedback system is to this critical confluence of conditions, we may quantify a system's *phase margin* (how many degrees of phase shift the system is away from 360° while at a loop gain of 1) and/or a system's *gain margin* (how many decibels of gain the system is away from 0 dB while at a phase shift of 360°). The less phase or gain margin a feedback system has, the closer it is to a condition of instability.

¹⁷At maximum phase shift, the gain of any first-order RC network is zero. Both phase shift and attenuation in an RC lag network are frequency-dependent: as frequency increases, phase shift grows larger (from 0° to a maximum of -90°) and the output signal grows weaker. At its theoretical maximum phase shift of exactly -90° , the output signal would be reduced to nothing!

As an illustration of this point, the following circuit is incapable¹⁸ of self-oscillation. Its lone RC phase-shifting network cannot create the -180° phase shift necessary for the overall loop to have positive feedback and oscillate:




The RC phase-shift oscillator circuit design thus holds a very important lesson for us in PID loop tuning. It clearly illustrates how multiple orders of lag are a more significant obstacle to robust control than a single lag time of *any* magnitude. A purely first-order process will tolerate enormous amounts of controller gain without ever breaking into oscillations, because it lacks the phase shift necessary to self-oscillate. This means – barring any other condition limiting our use of high gain, such as process noise – we may use very aggressive proportional-only action (e.g. gain values of 20 or more) to achieve robust control on a first-order process¹⁹. Multiple-order processes are less forgiving of high controller gains, because they *are* capable of generating enough phase shift to self-oscillate.

¹⁸In its pure, theoretical form at least. In practice, even a single-lag circuit may oscillate given enough gain due to the unavoidable presence of parasitic capacitances and inductances in the wiring and components causing multiple orders of lag (and even some dead time). By the same token, even a “pure” first-order process will oscillate given enough controller gain due to unavoidable lags and dead times in the field instrumentation (especially the control valve). The point I am trying to make here is that there is more to the question of stability (or instability) than loop gain.

¹⁹Truth be told, the same principle holds for purely integrating processes as well. A purely integrating process *always* exhibits a phase shift of -90° at any frequency, because that is the nature of integration in calculus. A purely first-order lag process will exhibit a phase shift anywhere from 0° to -90° depending on frequency, but never more lagging than -90° , which is not enough to turn negative feedback into positive feedback. In either case, so long as we don’t have process noise to deal with, we can increase the controller’s gain all the way to *eleven*. If that last sentence (a joke) does not make sense to you, be sure to watch the 1984 movie *This is Spinal Tap* as soon as possible. Seriously, I have used controller gains as high as 50 on low-noise, first-order processes such as furnace temperature control. With such high gain in the controller, response to setpoint and load changes is quite swift, and integral action is almost unnecessary because the offset is naturally so small.


30.1.7 Dead time

Lag time refers to a damped response from a process, from a change in manipulated variable (e.g. control valve position) to a measured change in process variable: the initial effect of a change in controller output is immediately seen, but the final effect takes time to develop. *Dead time*, by contrast, refers to a period of time during which a change in manipulated variable produces *no effect whatsoever* in the process variable: the process appears “dead” for some amount of time before showing a response. The following graph contrasts first-order and multiple-order lag times against pure dead time, as revealed in response to a manual step-change in the controller’s output (an “open-loop” test of the process characteristics):



Although the first-order response does take some time to settle at a stable value, there is no time delay between when the output steps up and the first-order response *begins* to rise. The same may be said for the multiple-order response, albeit with a slower rate of initial rise. The dead-time response, however, is actually delayed some time after the output makes its step-change. There is a period of time where the dead-time response does *absolutely nothing* following the output step-change.

Dead time is also referred to as *transport delay*, because the mechanism of dead time is often a time delay caused by the transportation of material at finite speed across some distance. The following cookie-baking process has dead time by virtue of the time delay inherent to the cookies' journey from the oven to the temperature sensor:




Dead time is a far worse problem for feedback control systems than lag time. The reason why is best understood from the perspective of phase shift: the delay (measured in degrees of angular displacement) between input and output for a system driven by a sinusoidal stimulus. Excessive phase shift in a feedback system makes possible self-sustaining oscillations, turning what is supposed to be negative feedback into positive feedback. Systems with lag produce phase shift that is frequency-dependent (the greater the frequency, the more the output “lags” behind the input), but this phase shift has a natural limit. For a first-order lag function, the phase shift has an absolute maximum value of -90° ; second-order lag functions have a theoretical maximum phase shift of -180° ; and so on. Dead time functions also produce phase shift that increases with frequency, but there is no ultimate limit to the amount of phase shift. This means a single dead-time element in a feedback control loop is capable of producing *any* amount of phase shift given the right frequency²⁰. What is more, the gain of a dead time function usually does not diminish with frequency, unlike the gain of a lag function.

Recall that a feedback system will self-oscillate if two conditions are met: a total phase shift of 360° (or -360° : the same thing) and a total loop gain of at least one. Any feedback system meeting these criteria²¹ will oscillate, be it an electronic amplifier circuit or a process control loop. In the interest of achieving robust process control, we need to prevent these conditions from ever occurring simultaneously.

²⁰A sophisticated way of saying this is that a dead-time function has no *phase margin*, only *gain margin*. All that is needed in a feedback system with dead time is sufficient gain to make the system oscillate.

²¹Sometimes referred to as the *Barkhausen criterion*.

A visual comparison between the phase shifts and gains exhibited by lag versus dead time functions may be seen here, the respective phase functions and gains exhibited by the electrical entities of a simple RC network (lag time) and an LC “delay line” network (dead time):




As frequency increases, the lag time function’s phase shift asymptotically approaches -90° while its attenuation asymptotically approaches zero. Ultimately, when the phase shift reaches its maximum of -90° , the output signal amplitude is reduced to nothing. By contrast, the dead time

function's phase shift grows linearly with frequency (to -180° and beyond!) while its attenuation remains unchanged. Clearly, dead time better fulfills the dual criteria of sufficient phase shift and sufficient loop gain needed for feedback oscillation than lag time, which is why dead time invites oscillation in a control loop more than lag time.

Pure dead-time processes are rare. Usually, an industrial process will exhibit at least some degree of lag time in addition to dead time. As strange as it may sound, this is a fortunate for the purpose of feedback control. The presence of lag(s) in a process guarantees a degradation of loop gain with frequency increase, which may help avoid oscillation. The greater the ratio between dead time and lag time in a loop, the more unstable it tends to be.

The appearance of dead time may be created in a process by the cascaded effect of multiple lags. As mentioned in an earlier subsection, multiple lags create a process response to step-changes that is "S"-shaped, responding gradually at first instead of immediately following the step-change. Given enough lags acting in series, the beginning of this "S" curve may be so flat that it appears "dead:"



While dead time may be impossible to eliminate in some processes, it should be minimized wherever possible due to its detrimental impact on feedback control. Once an open-loop (manual-mode step-change) test on a process confirms the existence of dead time, the source of dead time should be identified and eliminated if at all possible.

One technique applied to the control of dead-time-dominant processes is a special variation of the PID algorithm called *sample-and-hold*. In this variation of PID, the controller effectively alternates between "automatic" and "manual" modes according to a pre-programmed cycle. For a short period of time, it switches to "automatic" mode in order to "sample" the error ($PV - SP$) and calculate a new output value, but then switches right back into "manual" mode ("hold") so as to give time for the effects of those corrections to propagate through the process dead time before taking another sample and calculating another output value. This sample-and-hold cycle of course slows the controller's response to changes such as setpoint adjustments and load variations, but it does allow for more aggressive PID tuning constants than would otherwise work in a continuously sampling controller

because it effectively blinds²² the controller from “seeing” the time delays inherent to the process.

All digital instruments exhibit dead time due to the nature of their operation: processing signals over discrete time periods. Usually, the amount of dead time seen in a modern digital instrument is too short to be of any consequence, but there are some special cases meriting attention. Perhaps the most serious case is the example of *wireless* transmitters, using radio waves to communicate process information back to a host system. In order to maximize battery life, a wireless transmitter must transmit its data sparingly. Update times (i.e. dead time) measured in minutes are not uncommon for battery-powered wireless process transmitters.

²²An interesting analogy is that of a narcoleptic human operator manually controlling a process with a lot of dead time. If we imagine this person helplessly falling asleep at periodic intervals, then waking up to re-check the process variable and make another valve adjustment before falling asleep again, we see that the dead time of the process disappears from the perspective of the operator. The operator never realizes the process even has dead time, because they don't remain awake long enough to notice. So long as the poor operator's narcolepsy occurs at just the right intervals (i.e. not too short so as to notice dead time, and not too long so as to miss important changes in setpoint or load), good control of the process is possible.


30.1.8 Hysteresis

A detrimental effect to feedback control is a characteristic known as *hysteresis*: a lack of responsiveness to a change in direction. Although hysteresis typically resides with instruments rather than the physical process they connect to, it is most easily detected by a simple open-loop (“step-change”) test with the controller in manual mode just like all the important process characteristics (self-regulating versus integrating, steady-state gain, lag time, dead time, etc.).

The most common source of hysteresis is found in pneumatically-actuated control valves possessing excess stem friction. The “jerky” motion of such a valve to smooth increases or decreases in signal is sometimes referred to as *stiction*. Similarly, a pneumatically-actuated control valve with excess friction will be unresponsive to small reversals in signal direction. To illustrate, this means the control valve’s stem position will not be the same at a *rising* signal value of 50% (typically 12 mA, or 9 PSI) as it will be at a *falling* signal value of 50%.

Control valve stiction may be quite severe in valves with poor maintenance histories, and/or lacking positioners to correct for deviations between controller signal value and actual stem position. I have personally encountered control valves with hysteresis values in excess of 10%²³, and have heard of even more severe cases.


Detecting hysteresis in a control loop is as simple as performing “up-and-down” tests of the controller output signal in manual mode. The following trend shows how hysteresis might appear in a self-regulating process such as liquid flow control:



Note how the PV responds to large up-and-down output step-changes, but stops responding as soon as the magnitude of these open-loop step-changes falls below a threshold equal to the control valve’s hysteresis.

²³A 10% hysteresis value means that the signal must be changed by 10% following a reversal of direction before any movement is seen from the valve stem.

Applied to an integrating process such as liquid level control, the same type of test reveals the control valve's hysteresis by the largest step-change that does not alter the PV's slope:




It is not as simple to perform this test on a process with slow lag or dead times, of course, or on a process possessing a “runaway” (rather than self-regulating or integrating) characteristic, in which case a better test for valve hysteresis would be to monitor valve stem position rather than the PV when executing the step-changes.

Hysteresis is a problem in feedback control because it essentially acts like a variable dead time. Recall that “dead time” was defined as a period of time during which a change in manipulated variable produces no effect in the process variable: the process appears “dead” for some amount of time before showing a response. If a change in controller output (manipulated variable) is insufficient to overcome the hysteresis inherent to a control valve or other component in a loop, the process variable will not respond to that output signal change at all. Only when the manipulated variable signal continues to change sufficiently to overcome hysteresis will there be a response from the process variable, and the time required for that to take place depends on how soon the controller’s output happens to reach that critical value. If the controller’s output moves quickly, the “dead time” caused by hysteresis will be short. If the controller’s output changes slowly over time, the “dead time” caused by hysteresis will be longer.

Another problem caused by hysteresis in a feedback loop occurs in combination with integral action, whether it be programmed into the controller or is inherent to the process (i.e. an *integrating* process). It is highly unlikely that a “sticky” control valve will happen to “stick” at exactly the right stem position required to equalize PV and SP. Therefore, the probability at any time of an error developing between PV and SP, or of an offset developing between the valve position and the equilibrium position required by an integrating process, is very great. This leads to a condition of guaranteed instability. For a self-regulating process with integral action in the controller, the almost guaranteed existence of $PV - SP$ error means the controller output will ceaselessly ramp up and down as the valve first slips and sticks to give a positive error, then slips and sticks to give a negative


error. For an integrating process with proportional action in the controller, the process variable will ceaselessly ramp up and down as the valve first sticks too far open, then too far closed to equalize process in-flow and out-flow which is necessary to stabilize the process variable. In either case, this particular form of instability is called a *slip-stick cycle*.

The following process trend shows a slip-stick cycle in a self-regulating process, controlled by an integral-only controller:



Note how the output ceaselessly ramps in a futile attempt to drive the process variable to setpoint. Once sufficient pressure change accumulates in the valve actuator to overcome static stem friction, the valve “slips to and sticks at” a new stem position where the PV is unequal to setpoint, and the controller’s integral action begins to ramp the output in the other direction.

The next trend shows a slip-stick cycle in an integrating process, controlled by a proportional-only controller:



Note how the process variable's slope changes every time the valve "slips to and sticks at" a new stem position unequal to the balance point for this integrating process. The process's natural integrating action then causes the PV to ramp, causing the controller's proportional action to similarly ramp the output signal until the valve has enough accumulated force on its stem to jump to a new position.

It is very important to note that the problems created by a "sticky" control valve *cannot* be completely overcome by controller tuning²⁴. For instance, if one were to de-tune the integral-only controller (i.e. longer time constant, or fewer repeats per minute) in the self-regulating process, it would *still* exhibit the same slip-stick behavior, only over a longer period (lower frequency) than before. If one were to de-tune the proportional-only controller (i.e. greater proportional band, or less gain) in the integrating process, the exact same thing would happen: a decrease in cycle frequency, but no elimination of the cycling. Furthermore, de-tuning the controller in either process would also result in less responsive (poorer) corrective action to setpoint and load changes. The only solution²⁵ to either one of these problems is to reduce or eliminate the friction inside the control valve.

²⁴Some integral controllers are equipped with a useful feature called *integral deadband* or *reset deadband*. This is a special PID function inhibiting integration whenever the process variable comes close enough to setpoint, the "deadband" value specifying how close the PV must come to SP before integration stops. If this deadband value is set equal to or wider than the error caused by the valve's stiction, the controller will stop its integral-driven cycling. The trade-off, of course, is that the controller will no longer work to eliminate all error, but rather will be content with an error equal to or less than the specified deadband.

²⁵An alternate solution is to install a positioner on the control valve, which acts as a secondary (cascaded) controller seeking to equalize stem position with the loop controller's output signal at all times. However, this just "shifts" the problem from one controller to another. I have seen examples of control valves with severe packing friction which will *self-oscillate* their own positioners (i.e. the positioner will "hunt" back and forth for the correct valve stem position given a constant signal from the loop controller)! If valve stem friction can be minimized, it should be minimized.

30.2 Before you tune . . .

Much has been written about the benefits of robust PID control. Increased productivity, decreased equipment strain, and increased process safety are some of the advantages touted of proper PID tuning. What is often overlooked, though, are the negative consequences of poor PID controller tuning. If robust PID control can increase productivity, then poor PID control can decrease productivity. If a well-tuned system helps equipment run longer and safer, then a poorly tuned system may increase failure frequency and safety incidents. The instrumentation professional should be mindful of this dichotomy when proceeding to tune a PID control system. One should never think there is “nothing to lose” by trying different PID settings. Tuning a PID controller is as serious a matter as reconfiguring any field instrument.

PID tuning parameters are easy to access, which makes them a tempting place to begin for technicians looking to improve the performance of a feedback loop. Another temptation driving technicians to focus on controller tuning as a first step is the prestige associated with being able to tame an unruly feedback loop with a few adjustments to the controller’s PID tuning constants. For those who do not understand PID control (and this constitutes the vast majority of the human population, even in the industrial world), there is something “magic” about being able to achieve robust control behavior simply by making small adjustments to numbers in a computer (or to knobs in an analog controller). The reality, though, is that many poorly-behaving control systems are that way not due (at least purely) to a deficit of proper PID tuning values, but rather to problems external to the controller which no amount of “tuning” will solve. Adjusting PID tuning constants *as a first step* is almost always a bad idea.

This section aims to describe and explain some of the recommended considerations prior to making adjustments to the tuning of a loop controller. These considerations include:

- Identifying operational needs (i.e. “How do the operators want the system to respond?”)
- Identifying process and system hazards before manipulating the loop
- Identifying whether it is a tuning problem, a field instrument problem, and/or a design problem

30.2.1 Identifying operational needs

As defined elsewhere in this book, “robust” control is a stability of the process variable despite changes in load, fast response to changes in setpoint, minimal oscillation following either type of change, and minimal offset (error between setpoint and process variable) over time. However, these criteria are not equally valued in all processes, and neither are they equally attainable with simple PID control in all processes. It may be critical, for example, in a boiler water level control process to have fast response to changes in load, but minimal offset over time is not as important. It may be completely permissible to have a persistent 5% error between PV and SP in such a system, so long as the water level does not deviate much over 20% for any length of time due to load changes. In another process, such as liquid level control inside one stage (“effect”) of a multi-stage (“multi-effect”) evaporator system, a priority may be placed upon relatively steady flow control through the valve rather than steady level in the process. A level controller tuned for aggressive response to setpoint changes will cause large fluctuations in liquid flow rate to all successive stages (“effects”) of the evaporator process in the event of a sudden load or setpoint change, which would be more detrimental to product quality than some deviation from setpoint in that one effect.

Thus, we must first determine what the operational needs of a control system are before we aim to adjust the performance of that control system. The operations personnel (operators, unit managers, process engineers) are your best resources here. Ultimately, they are your “internal customers.” Your task is to give the customers the system performance they need to do their jobs best.

Keep in mind the following process control objectives, knowing that it will likely be impossible to achieve *all* of them with any particular PID tuning. Try to rank the relative importance of these objectives, then concentrate on achieving those most important, at the expense of those least important:

- Minimum change in PV (dynamic stability) with changes in load
- Fast response to setpoint changes (minimum dynamic error)
- Minimum overshoot/undershoot/oscillation following sudden load or setpoint changes
- Minimum error (PV – SP) over time
- Minimum valve velocity (i.e. minimal effect to upstream or downstream processes)


The control actions best suited for rapid response to load and/or setpoint changes are proportional and derivative. Integral action takes effect only *after* error has had time to develop, and as such cannot act as immediately as either proportional or derivative.

If the priority is to minimize overshoot, undershoot, and/or oscillations, the controller’s response will likely need to be more sluggish than is typical. New setpoint values will take longer to achieve, and load changes will not be responded to with quite the same vigor. Derivative action may be helpful in some applications to “tame” the oscillatory tendencies of proportional and integral actions.

Minimum error over time can really only be addressed by integral action. No other controller action pays specific “attention” to the magnitude and duration of error. This is not to say that the process will work well on integral-only control, but rather that integral action will be absolutely necessary (i.e. a P-only or PD controller will not suffice).

Minimum valve velocity is a priority in processes where the manipulated variable has an effect on some *other* process in the system. For example, liquid level control in a multi-stage (multi-“effect”) evaporator system where the discharge flow from one evaporator becomes the incoming flow for another evaporator, is a system where sharp changes in the manipulated variable of one control loop can upset downstream processes:

A multi-effect evaporator system



In other words, an aggressively-tuned level controller on an upstream evaporator (e.g. Effect 1) may achieve its goal of holding liquid level very steady in that evaporator by varying its out-going flow, but it will do so at the expense of causing level variations in all downstream evaporators. Cases such as this call for controller tuning (at least in the upstream effects) responding slowly to errors. Proportional action will very likely be limited to low gain values (high proportional band values), and derivative action (if any is used at all) should be set to respond only to the process variable, not to error ($PV - SP$). This leaves the main work of stabilizing the loop to integral action, even though we know that integral action tends to overshoot following setpoint changes in an integrating process such as liquid level control. Understand that tuning a PID loop with the goal of minimizing valve motion *will* result in longer deviations from setpoint than if the controller were tuned to respond faster to process or setpoint changes.

30.2.2 Identifying process and system hazards

When students practice PID control in an Instrumentation program, they usually do so using computer simulation software and/or “toy” processes constructed in a lab environment. A potential disadvantage to this learning environment is a failure to recognize real problems that may develop when tuning an actual production process. Rarely will you find a completely isolated feedback loop in industry: generally there are interactive effects between control loops in a process, which means one cannot proceed to tune a loop with impunity.

A very important question to ask the operations personnel before tuning a loop is, “How far and how fast am I allowed to let the process variable increase and decrease?” Processes and process equipment may become dangerously unstable, for example, if certain temperatures become too high (or too low, as is the case in process liquids that solidify when cold). It is not uncommon for certain control loops in a process to be equipped with alarms, either hard or soft, that automatically *shut down* equipment if exceeded. Clearly, these “shutdown” limits must be avoided during the tuning of the process loop.

One should also examine the control strategy before proceeding to tune. Is this a cascaded loop? If so, the slave controller needs to be tuned before the master. Does this loop incorporate feedforward action to act on load changes? If so, the effectiveness of that feedforward loop (gain, dynamic compensation) should be checked and adjusted before the feedback loop is tuned. Are there limits in this loop? Is this a selector or override control strategy? If so, you need to be able to clearly tell which loop components are selected, and which signals are being limited, at any given time.

Another consideration is whether or not the process is in a “normal” condition before you attempt to improve its performance. Ask the operations personnel if this is a typical day, or if there is some abnormal condition in effect (equipment shutdown, re-routing of flows, significantly different production rates, etc.) that might skew the response of the process loop to be tuned. Once again we see a need for input from the operations personnel, because they know the day-to-day behavior of the system better than anyone else.

30.2.3 Identifying the problem(s)

One of the questions I advise instrument technicians to ask of operators when diagnosing any process problem is simply, “How long has this problem existed?” The age of a problem can be a very important indicator of possible causes. If you were told that a problem suddenly developed after the last night shift, you would be inclined to suspect an equipment failure, or something else that could happen *suddenly* (e.g. a hand valve someone opened or shut when they shouldn’t have). Alternatively, if you were told a problem has been in existence since the day the process was constructed, you would be more inclined to suspect an issue with system design or improper installation. This same diagnostic technique – obtaining a “history” of the “patient” – applies to loop tuning as well. A control loop that suddenly stopped working as it should might be suffering from an instrument failure (or an unauthorized change of controller parameters), whereas a chronically misbehaving loop would more likely be suffering from poor design, bad instrument installation, or a controller that was never tuned properly.

In either case, poor control is just as likely to be caused by field instrument problems as it is by incorrect PID tuning parameters. No PID settings can fully compensate for faulty field instrumentation, but it is possible for some instrument problems to be “masked” by controller tuning. Your first step in actually manipulating the control loop should be a check of instrument health. Thankfully, this is relatively easy to do by performing a series of “step-change” tests with the controller in manual mode. By placing the controller in manual and making small changes in output signal (remember to check with operations to see how far you are allowed to move the output, and how far you can let the PV drift!), you can determine much about the process and the loop instrumentation, including:

- Whether the PV signal is “noisy” (first turn off all damping in the controller and transmitter)
- How much “stiction” is in the control valve
- Whether the process is integrating, runaway, or self-regulating
- Process gain (and whether this gain is stable or if it changes as PV changes)
- Process lag time and lag degree (first-order versus multiple-order)
- Process dead time

Such an open-loop test might reveal potential problems without pinpointing the exact nature or location of those problems. For example, a large lag time may be intrinsic to the process, or it may be the result of a poorly-installed sensor (e.g. a thermocouple not pushed to touch the bottom of its thermowell) or even a control valve in need of a volume booster or positioner. Dead time measured in an open-loop test may also be intrinsic to the process (transport delay), intrinsic to the sensor (e.g. a gas chromatograph where each analysis cycle takes several *minutes* of time), or it could be the result of stiction in the valve. The only way to definitively identify the problem is to test the instruments themselves, ideally in the field location.

An indispensable tool for identifying loop problems is a *trend recorder*, showing all the relevant variables in a control loop graphed over time. In order to obtain the best “view” of the process, you need to make sure the graphing trend display has sufficient resolution and responsiveness. If the

trend fails to show fine details such as noise in the process, it is possible that the graphing device will be insufficient for your needs.

If this is the case, you may still perform response tests of the loop, but you will have to use some other instrument(s) to graph the controller and process actions. A modern tool useful for this purpose is a portable computer with a data acquisition device connected, giving the computer the ability to read instrument signal voltages. Many data graphing programs exist for taking acquired data and plotting it over the time domain. Data acquisition modules with sample rates in the thousands of samples per second are available for very modest prices.

30.2.4 Final precautions

Be prepared to document your work! This means capturing and recording “screen shot” images of process trend graphs, both for the initial open-loop tests and the closed-loop PID trials. It also means documenting the original PID settings of the controller, and all PID setting values attempted during the tuning process (linked to their respective trend graphs, so it will be easy to tell which sets of PID constants produced which process responses). If there are any instrument configuration settings (e.g. damping time values in process transmitters) changed during the tuning exercise, both the original values and all your changes need to be documented as well.

As a final word, I would like to cast a critical vote against auto-tuning controllers. With all due respect for the engineers who work hard to make controllers “intelligent” enough to adjust their own PID settings, there is no controller in the world able to account for all the factors described in this “Before you tune . . .” section. Feel free to use the automatic tuning feature of a controller, but only *after* you have ensured all instrument and process problems are corrected, and *after* you have confirmed the tuning goal of the controller matches the behavioral goal of the control loop as defined by the operators (e.g. fast response versus minimum overshoot, etc.). Some people in the automation business are over-confident with regard to the capabilities of auto-tuning controllers. We would all do well to recognize this feature as a *tool*, and just like any other tool it is only as useful as the person handling it is knowledgeable regarding how and why it works. Wielding any tool in ignorance is a recipe for disaster.

30.3 Quantitative PID tuning procedures

A *quantitative* PID tuning procedure is a step-by-step approach leading directly to a set of numerical values to be used in a PID controller. These procedures may be split into two categories: *open loop* and *closed loop*. An “open loop” tuning procedure is implemented with the controller in manual mode: introducing a step-change to the controller output and then mathematically analyzing the results of the process variable response to calculate appropriate PID settings for the controller to use when placed into automatic mode. A “closed loop” tuning procedure is implemented with the controller in automatic mode: adjusting tuning parameters to achieve an easily-defined result, then using those PID parameter values and information from a graph of the process variable over time to calculate new PID parameters.


Quantitative PID tuning got its start with a paper published in the November 1942 *Transactions of the American Society of Mechanical Engineers* written by two engineers named Ziegler and Nichols. “Optimum Settings For Automatic Controllers” is a seminal paper, and deserves to be read by every serious student of process control theory. That Ziegler’s and Nichols’ recommendations for PID controller settings may still be found in modern references more than 60 years after publication is a testament to its impact in the field of industrial control. Although dated in its terminology and references to pneumatic controller technology (some controllers mentioned as not even having adjustable proportional response, and others as having only discrete degrees of reset adjustment rather than continuously variable settings!), the PID algorithm described by its authors and the effects of P, I, and D adjustments on process control behavior are as valid today as they were then.

This section is devoted to a discussion of quantitative PID tuning procedures in general, and the “Ziegler-Nichols” methods in specific. It is the opinion of this author that the Ziegler-Nichols tuning methods are useful primarily as historical references, and indeed suffer from serious practical impediments. The most serious reservation I have with the Ziegler-Nichols methods (and in fact any algorithmic procedure for PID tuning) is that these methods tend to absolve the practitioner of responsibility for understanding the process they intend to tune. Any time you provide people with step-by-step instructions to perform complex tasks, there will be a great many readers of those instructions tempted to mindlessly follow the instructions, even to their doom. PID tuning is one of these “complex tasks,” and there is significant likelihood for a person to do more harm than good if all they do is implement a step-by-step approach rather than understand what they are doing, why they are doing it, and what it means if the results do not meet with satisfaction. Please bear this in mind as you study any PID tuning procedure, Ziegler-Nichols or otherwise.

30.3.1 Ziegler-Nichols closed-loop (“Ultimate Gain”)

Closed-loop refers to the operation of a control system with the controlling device in “automatic” mode, where the flow of the information from sensing element to transmitter to controller to control element to process and back to sensor represents a continuous (“closed”) feedback loop. If the total amount of signal amplification provided by the instruments is too much, the feedback loop will self-oscillate at the system’s natural (resonant) frequency. While oscillation is almost always considered undesirable in a control system, it may be used as an exploratory test of process dynamics if the controller acts purely on proportional action (no integral or derivative action): providing data useful for calculating effective PID controller settings. Thus, a “closed-loop” PID tuning procedure entails disabling any integral or derivative actions in the controller, then raising the gain value of the controller just far enough that self-sustaining oscillations ensue. The minimum amount of controller gain necessary to sustain sinusoidal oscillations is called the *ultimate sensitivity* (S_u) or *ultimate gain* (K_u) of the process, while the time (period) between successive oscillation peaks is called the *ultimate period* (P_u) of the process. We may then use the measured values of K_u and P_u to calculate reasonable controller tuning parameter values (K_p , τ_i , and/or τ_d).

When performing such a test on a process loop, it is important to ensure the oscillation peaks do not reach the limits of the instrumentation, either measurement or final control element. In other words, in order for the oscillation to accurately reveal the process characteristics of ultimate sensitivity and ultimate period, the oscillations must be naturally limited and not artificially limited by either the transmitter or the control valve saturating. Oscillations characterized by either the transmitter or the final control element reaching their range limits should be avoided in order to obtain the best closed-loop oscillatory test results. An illustration is shown here as a model of what to avoid:



Here the controller gain is set too high, the result being saturation at the positive peaks of the output waveform. The controller gain should be decreased until symmetrical, sinusoidal waves

result.

If the controller in question is proportional-only (i.e. capable of providing no integral or derivative control actions), Ziegler and Nichols' recommendation is to set the controller gain²⁶ to one-half the value of the ultimate sensitivity determined in the closed-loop test, which I will call ultimate *gain* (K_u) from now on:

$$K_p = 0.5K_u$$


Where,

K_p = Controller gain value that you should enter into the controller for good performance

K_u = "Ultimate" gain determined by increasing controller gain until self-sustaining oscillations are achieved

Generally, a controller gain of one-half the experimentally determined "ultimate" gain results in reasonably quick response to setpoint and process load changes. Oscillations of the process variable following such setpoint and load changes typically damp with each successive wave peak being approximately one-quarter the amplitude of the one preceding. This is known as *quarter-wave damping*. While certainly not ideal, it is a compromise between fast response and stability.

The following process trend shows what "quarter-wave damping" looks like with the controller in automatic mode, with the process variable (PV) exhibiting decaying oscillations following a step-change in setpoint (SP):



²⁶Note that this is truly the *gain* of the controller, not the *proportional band*. If you were to enter a proportional band value one-half the proportional band value necessary to sustain oscillations, the controller would (obviously) oscillate completely out of control!

Ziegler and Nichols were careful to qualify quarter-wave damping as less than optimal for some applications. In their own words (page 761):

“The statement that a sensitivity setting of one half the ultimate with attendant 25 per cent amplitude ratio gives optimum control must be modified in some cases. For example, the actual level maintained by a liquid-level controller might not be nearly as important as the effect of sudden valve movements on further portions of the process. In this case the sensitivity should be lowered to reduce the amplitude ratio even though the offset is increased by so doing. On the other hand, a pressure-control application giving oscillations with very short period could be set to give an 80 or 90 per cent amplitude ratio. Due to the short period, a disturbance would die out in reasonable time, even though there were quite a few oscillations. The offset would be reduced somewhat though it should be kept in mind that it can never be reduced to less than one half of the amount given at our previously defined optimum sensitivity of one half the ultimate.”

Some would argue (myself included) that quarter-wave damping exercises the control valve needlessly, causing undue stem packing wear and consuming large quantities of compressed air over time. Given the fact that all modern process controllers have integral (reset) capability, unlike the simple pneumatic controllers of Ziegler and Nichols’ day, there is really no need to tolerate prolonged offset (failure of the process variable to exactly equalize with setpoint over time) as a necessary cost of avoiding valve oscillation.

If the controller in question has integral (reset) action in addition to proportional, Ziegler and Nichols’ recommendation is to set the controller gain to slightly less than one-half the value of the ultimate sensitivity, and to set the integral time constant²⁷ to a value slightly less than the ultimate period:

$$K_p = 0.45K_u$$

$$\tau_i = \frac{P_u}{1.2}$$

Where,

K_p = Controller gain value that you should enter into the controller for good performance

K_u = “Ultimate” gain determined by increasing controller gain until self-sustaining oscillations are achieved

τ_i = Controller integral setting that you should enter into the controller for good performance (minutes per repeat)

P_u = “Ultimate” period of self-sustaining oscillations determined when the controller gain was set to K_u (minutes)

²⁷Either minutes per repeat or seconds per repeat. If the controller’s integral rate is expressed in units of repeats per minute (or second), the formula would be $K_i = \frac{1.2}{P_u}$.

If the controller in question has all three control actions present (full PID), Ziegler and Nichols' recommendation is to set the controller tuning constants as follows:

$$K_p = 0.6K_u$$

$$\tau_i = \frac{P_u}{2}$$

$$\tau_d = \frac{P_u}{8}$$

Where,

K_p = Controller gain value that you should enter into the controller for good performance

K_u = "Ultimate" gain determined by increasing controller gain until self-sustaining oscillations are achieved

τ_i = Controller integral setting that you should enter into the controller for good performance (minutes per repeat)

τ_d = Controller derivative setting that you should enter into the controller for good performance (minutes)

P_u = "Ultimate" period of self-sustaining oscillations determined when the controller gain was set to K_u (minutes)

An important caveat with any tuning procedure based on ultimate gain is the potential to cause trouble in a process while experimentally determining the ultimate gain. Recall that "ultimate" gain is the amount of controller gain (proportional action) resulting in self-sustaining oscillations of constant amplitude. In order to precisely determine this gain setting, one must spend some time provoking the process with sudden setpoint changes (to induce oscillation) and experimenting with greater and greater gain settings until constant oscillation amplitude is achieved. Any more gain than the "ultimate" value, of course, leads to ever-growing oscillations which may be brought under control only by decreasing controller gain or switching to manual mode (thereby stopping all feedback in the system). The problem with this is, one never knows for certain when ultimate gain is achieved until this critical value has been exceeded, as evidenced by ever-growing oscillations. In other words, *the system must be brought to the brink of total instability in order to determine its ultimate gain value*. Not only is this time-consuming to achieve – especially in systems where the natural period of oscillation is long, as is the case with many temperature and composition control applications – but potentially hazardous to equipment and certainly detrimental to process quality²⁸. In fact, one might argue that any process tolerant of such abuse probably doesn't need to be well-tuned at all!

Despite its practical limitations, the rules given by Ziegler and Nichols do shed light on the relationship between realistic P, I, and D tuning parameters and the operational characteristics of the process. Controller gain should be some fraction of the gain necessary for the process to self-oscillate. Integral time constant should be proportional to the process time constant; i.e. the


²⁸Imagine informing the lead operations manager or a unit supervisor in a chemical processing facility you wish to over-tune the temperature controller in the main reaction furnace or the pressure controller in one of the larger distillation columns until it nearly oscillates out of control, and that doing so may necessitate hours of unstable operation before you find the perfect gain setting. Consider yourself fortunate if your declaration of intent does not result in security personnel escorting you out of the control room.

“slower” the process is to respond, the “slower” (less aggressive) the controller’s integral response should be. Derivative time constant should likewise be proportional to the process time constant, although this has the opposite meaning from the perspective of aggressiveness: a “slow” process deserves a long derivative time constant; i.e. *more aggressive* derivative action.

30.3.2 Ziegler-Nichols open-loop

In contrast to the first tuning technique presented by Ziegler and Nichols in their landmark 1942 paper where the process was made to oscillate using proportional-only automatic control and the parameters of that oscillation served to define PID tuning parameters, their second tuning technique did not even rely on the presence of a controller. Instead, this second technique consisted of making a manual “step-change” of the control element (valve) and analyzing the resulting effect on the process variable, much the same way as described in the Process Characterization section of this chapter (section 30.1 beginning on page 2408).

After making the step-change in output signal with the controller in manual mode, the process variable trend is closely analyzed for two salient features: the *dead time* and the *reaction rate*. Dead time (L)²⁹ is the amount of time delay between the output step-change and the first indication of process variable change. Reaction rate is the maximum rate at which the process variable changes following the output step-change (the maximum time-derivative of the process variable):



Dead time and reaction rate are responses common to self-regulating and integrating processes alike. Whether or not the process variable ends up stabilizing at some new value, its rate of rise will

²⁹Unfortunately, Ziegler and Nichols chose to refer to dead time by the word *lag* in their paper. In modern technical parlance, “lag” refers to a first-order inverse-exponential function, which is fundamentally different from dead time.

reach some maximum value following the output step-change, and this will be the reaction rate of the process³⁰. The unit of measurement for reaction rate is *percent per minute*:

$$R = \frac{\Delta PV}{\Delta t} = \frac{[\text{Percent rise}]}{[\text{Minutes run}]}$$

While dead time in a process tends to be constant regardless of the output step-change magnitude, reaction rate tends to vary directly with the magnitude of the output step-change. For example, an output step-change of 10% will generally cause the PV to rise at a rate twice as steep compared to the effects of a 5% output step-change. In order to ensure our predictive calculations capture only what is inherent to the process and not our own arbitrary open-loop tuning actions, we must include the output step-change magnitude (Δm) in those calculations as well³¹.

If the controller in question is proportional-only (i.e. capable of providing no integral or derivative control actions), Ziegler and Nichols' recommendation is to set the controller gain as follows:

$$K_p = \frac{\Delta m}{RL}$$

Where,

- K_p = Controller gain value that you should enter into the controller for good performance
- Δm = Output step-change magnitude made while testing in open-loop (manual) mode (percent)
- R = Process reaction rate = $\frac{\Delta PV}{\Delta t}$ (percent per minute)
- L = Process dead time (minutes)

If the controller in question has integral (reset) action in addition to proportional, Ziegler and Nichols' recommendation is to set the controller gain to 90% of the proportional-only value, and to set the integral time constant to a value just over three times the measured dead time value:

$$K_p = 0.9 \frac{\Delta m}{RL}$$

$$\tau_i = 3.33L$$

Where,

- K_p = Controller gain value that you should enter into the controller for good performance
- Δm = Output step-change magnitude made while testing in open-loop (manual) mode (percent)
- R = Process reaction rate = $\frac{\Delta PV}{\Delta t}$ (percent per minute)
- L = Process dead time (minutes)
- τ_i = Controller integral setting that you should enter into the controller for good performance (minutes per repeat)

³⁰Right away, we see a weakness in the Ziegler-Nichols open-loop method: it makes absolutely no distinction between self-regulating and integrating process types. We know this is problematic from the analysis of each process type in sections 30.1.1 and 30.1.2.

³¹Ziegler and Nichols' approach was to define a normalized reaction rate called the *unit reaction rate*, equal in value to $\frac{R}{\Delta m}$. I opt to explicitly include Δm in all the tuning parameter equations in order to avoid the possibility of confusing reaction rate with unit reaction rate.

If the controller has full PID capability, Ziegler and Nichols' recommendation is to set the controller gain to 120% of the proportional-only value, to set the integral time constant to twice the measured dead time value, and to set the derivative time constant to one-half the measured dead time value:

$$K_p = 1.2 \frac{\Delta m}{RL}$$

$$\tau_i = 2L$$

$$\tau_d = 0.5L$$

Where,

K_p = Controller gain value that you should enter into the controller for good performance

Δm = Output step-change magnitude made while testing in open-loop (manual) mode (percent)

R = Process reaction rate = $\frac{\Delta PV}{\Delta t}$ (percent per minute)

L = Process dead time (minutes)

τ_i = Controller integral setting that you should enter into the controller for good performance (minutes per repeat)

τ_d = Controller derivative setting that you should enter into the controller for good performance (minutes)

As you can see, the Ziegler-Nichols open-loop tuning method relies heavily on dead time (L) as a descriptive parameter for the process. This may be problematic in processes having insubstantial dead time, as the small L values obtained during the open-loop test will predict large controller gain (K_p) and aggressive integral (τ_i) time constant values, often too large to be practical. The open-loop method, however, is less disruptive to an operating process than the closed-loop method (which necessitated over-tuning the controller to the brink of total instability).

Another limitation, common to both the closed-loop and open-loop tuning methods, is that other factors in the process such as noise and hysteresis are completely overlooked. Noise is troublesome for large controller gain values (because the controller's proportional action reproduces that noise on the output) and is especially troublesome for derivative action which amplifies any noise it sees. Hysteresis causes integral action to continually "hunt" up and down, leading to cycling of the process variable. The lesson here is that no algorithmic PID tuning method can replace informed judgment on the part of the person tuning the loop. The methods proposed by Ziegler and Nichols (and others!) are merely starting points, and should never be taken as a definitive answer for controller tuning.

30.4 Heuristic PID tuning procedures

In contrast to quantitative tuning procedures where definite numerical values for P, I, and D controller settings are obtained through data collection and analysis, a *heuristic* tuning procedure is one where general rules are followed to obtain approximate or qualitative results. The majority of PID loops in the world have been tuned with such methods, for better or for worse. My goal in this section is to optimize the effectiveness of such tuning methods.

When I was first educated on the subject of PID tuning, I learned this rather questionable tuning procedure:

1. Configure the controller for proportional action only (integral and derivative control actions set to minimum effect), setting the gain near or at 1.
2. Increase controller gain until self-sustaining oscillations are achieved, “bumping” the setpoint value up or down as necessary to provoke oscillations.
3. When the ultimate gain is determined, reduce the aggressiveness of proportional action by a factor of two.
4. Repeat steps 2 and 3, this time adjusting integral action instead of proportional.
5. Repeat steps 2 and 3, this time adjusting derivative action instead of proportional.

The first three steps of this procedure are identical to the steps recommended by Ziegler and Nichols for closed-loop tuning. The last two steps are someone else’s contribution. The results of this method are generally poor, *and I strongly recommend against using it!*

While this particular procedure is crude and ineffective, it does illustrate a useful principle in trial-and-error PID tuning: we can tune a PID-controlled process by incrementally adjusting the aggressiveness of a controller’s P, I, and/or D actions until we see oscillations (suggesting the action has become too aggressive), then reducing the aggressiveness of the action until stable control is achieved. The “trick” to doing this effectively and efficiently is knowing which action(s) to focus on, which action(s) to avoid, and how to tell which of the actions is too aggressive when things do begin to oscillate. The following portions of this subsection describe the utility of each control action, the limitations of each, and how to recognize an overly-aggressive condition.

Much improvement may be made to any “trial-and-terror” PID tuning procedure if one is aware of the process characteristics and recognizes the applicability of P, I, and D actions to those process characteristics. Random experimentation with P, I, and D parameter values is tedious at best and dangerous at worst! As always, the key is to *understand* the role of each action, their applicability to different process types, and their limitations. The competent loop-tuner should be able to visually analyze trends of PV, SP, and Output, and be able to discern the degrees of P, I, and D action in effect at any given time in that trend.

Here is an improved PID tuning technique employing heuristics (general rules) regarding P, I, and D actions. It is assumed that you have taken all necessary safety precautions (e.g. you know the hazards of the process and the limits you are allowed to change it) and other steps recommended in the “Before You Tune . . .” section (30.2) beginning on page 2451:

1. Perform open-loop (manual-mode) tests of the process to determine its natural characteristics (e.g. self-regulating versus integrating versus runaway, steady-state gain, noisy versus calm, dead time, time constant, lag order) and to ensure no field instrument or process problems exist (e.g. control valve with excessive friction, inconsistent process gain, large dead time). *Correct all problems before proceeding*³².
2. Identify any controller actions that may be problematic (e.g. derivative action on a noisy process), noting to use them sparingly or not at all.
3. Identify whether the process will depend mostly on proportional action or integral action for stability. This will be the controller’s “dominant” action when tuned. You may find the chart on page 2469 useful for this.
4. Start with all terms of the controller set for minimal response (minimal P, minimal I, no D).
5. Set the dominant action to some safe value³³ (e.g. gain less than 1, integral time much longer than time constant of process) and check the loop’s response to setpoint and/or load changes in automatic mode.
6. Increase aggressiveness of this action until a point is reached where any more causes excessive overshoot or oscillation.
7. Increase aggressiveness of the other action(s) as needed to achieve the best compromise between stability and quick response.
8. If the loop ever shows signs of being too aggressive (e.g. oscillations), use the technique of phase-shift comparison between PV and Output trends (see page 2470) to identify which controller action to attenuate.
9. Repeat the last three steps as often as needed.

Note: the loop should *never* “porpoise” (see page 2474) when responding to a setpoint or load change! Some oscillation around setpoint may be tolerable (especially when optimizing the tuning for fast response to setpoint and load changes), but “porpoising” is always something to avoid.

³²This is very important: no degree of controller “tuning” will fix a poor control valve, noisy transmitter, or ill-designed process. If your open-loop tests reveal significant process problems, you must remedy them before attempting to tune the controller.

³³It is important to know which PID equation your controller implements in order to adjust just one action (P, I, or D) of the controller without affecting the others. Most PID controllers, for example, implement either the “Ideal” or “Series” equations, where the gain value (K_p) multiplies every action in the controller including integral and derivative. If you happen to be tuning such a controller for integral-dominant control, you cannot set the gain to zero (in order to minimize proportional action) because this will nullify integral action too! Instead, you must set K_p to some value small enough that the proportional action is minimal while allowing integral action to function.

30.4.1 Features of P, I, and D actions

Purpose of each action

- **Proportional action** is the “universal” control action, capable of providing at least marginal control quality for any process.
- **Integral action** is useful for eliminating offset caused by load variations and process self-regulation.
- **Derivative action** is useful for canceling lags, but useless by itself.

Limitations of each action

- **Proportional action** will cause oscillations if sufficiently aggressive, in the presence of lags and/or dead time. The more lags (higher-order), the worse the problem. It also directly reproduces process noise onto the output signal.
- **Integral action** will cause oscillation if sufficiently aggressive, in the presence of lags and/or dead time. Any amount of integral action will guarantee overshoot following setpoint changes in purely integrating processes.
- **Derivative action** dramatically amplifies process noise, and will cause oscillations in fast-acting processes.

Special applicability of each action






- **Proportional action** works exceptionally well when aggressively applied to processes lacking the phase shift necessary to oscillate: self-regulating processes dominated by first-order lag, and purely integrating processes.
- **Integral action** works exceptionally well when aggressively applied to fast-acting, self-regulating processes. Has the unique ability to ignore process noise.
- **Derivative action** works exceptionally well to speed up the response of processes dominated by large lag times, and to help stabilize runaway processes. Small amounts of derivative action will sometimes allow more aggressive P and/or I actions to be used than otherwise would be possible without unacceptable overshoot.

Gain and phase shift of each action

- **Proportional action** acts on the *present*, adding no phase shift to a sinusoidal signal. Its gain is constant for any signal frequency.
- **Integral action** acts on the *past*, adding a -90° phase shift to a sinusoidal signal. Its gain decreases with increasing frequency.
- **Derivative action** acts on the *future*, adding a $+90^\circ$ phase shift to a sinusoidal signal. Its gain increases with increasing frequency.

30.4.2 Tuning recommendations based on process dynamics

Knowing which control actions to focus on first is a matter of characterizing the process (identifying whether it is self-regulating, integrating, runaway, noisy, has lag or dead time, or any combination of these traits based on an open-loop response test³⁴) and then selecting the best actions to fit those characteristics. The following table shows some general recommendations for fitting PID tuning to different process characteristics

	Pure self-regulating	<i>May be controlled with aggressive integral action, and perhaps with a bit of proportional action. Use absolutely no derivative action!</i>
	Self-reg w/ pure 1 st order lag	<i>Responds well to aggressive proportional action, with integral action needed only for recovery from load changes.</i>
	Self-reg w/ multiple lags	<i>Proportional action needed for quick response to setpoint changes, integral action needed for recovery from load changes, and derivative needed to prevent overshoot. Proportional and integral actions are limited by tendency to oscillate.</i>
	Integrating w/ lag(s)	<i>Proportional action should be aggressive as possible without generating oscillations. Integral action needed only for recovery from load changes.</i>
	Pure integrating	<i>Responds well to aggressive proportional action, with integral action needed only for recovery from load changes.</i>

General rules:

- Use no derivative action if the process signal is “noisy”
- Use proportional action sparingly if the process signal is “noisy”
- The slower the time lag(s), the less integral action to use (a good approximation is to set the integration time τ_i equal to the measured lag time of the process)
- The higher-order the time lag(s), the less proportional action (gain) to use
- Self-regulating processes *need* integral action
- Integrating processes *need* proportional action


³⁴Recall that an open-loop response test consists of placing the loop controller in manual mode, introducing a step-change to the controller output (manipulated variable), and analyzing the time-domain response of the process variable as it reacts to that perturbation.

- Dead time requires a reduction of all PID constants below what would normally work

Once you have determined the basic character of the process, and understand from that characterization what the needs of the process will be regarding P, I, and/or D control actions, you may “experiment” with different tuning values of P, I, and D until you find a combination yielding robust control.

30.4.3 Recognizing an over-tuned controller by phase shift


When performing heuristic tuning of a PID controller, it is important to be able to identify a condition where one or more of the “actions” (P, I, or D) is configured too aggressively for the process. The characteristic indication of over-tuning is the presence of sinusoidal oscillations. At best, this means damped oscillations following a sudden setpoint or load change. At worst, this means oscillations that *never* decay:



At this point, the question is: *which action of the controller is causing this oscillation?* We know all three actions (P, I, and D) are fully capable of causing process oscillation if set too aggressively, so in the lack of clarifying information it could be any of them (or some combination!).

One clue is the trend of the controller’s *output* compared with the trend of the process variable (PV). Modern digital control systems all have the ability to trend manipulated variables as well as process variables, allowing personnel to monitor exactly how a loop controller is managing a process. If we compare the two trends, we may examine the *phase shift* between PV and output of a loop controller to discern what it’s dominant action is.

For example, the following trend graphs show the PV and output signals for a loop controller with proportional-dominant response. Both direct- and reverse-acting versions are shown:



Since we know proportional action is *immediate*, there should be no phase shift³⁵ between the PV and output waveforms. This makes sense from a mathematical perspective: if we substitute a sine function for the error variable in a proportional-only controller equation, we see that different gain values (K_p) simply result in an output signal of the same phase, just amplified or attenuated:

$$m = K_p e + b$$


$$m = K_p(\sin t) + b$$

If ever you see a process oscillating like this (sinusoidal waveforms, with the PV and output signals in-phase), you know that the controller's response to the process is dominated by proportional action, and that the gain needs to be reduced (i.e. increase proportional band) to achieve stability.

³⁵For reverse-acting controllers, I am ignoring the obvious 180° phase shift necessary for negative feedback control when I say “no phase shift” between PV and output waveforms. I am also ignoring dead time resulting from the scanning of the PID algorithm in the digital controller. For some controllers, this scan time may be significant enough to see on a trend!

Integral and derivative actions, however, introduce phase shift between the PV waveform and the output waveform. The direction of phase shift will reveal which time-based action (either I or D) dominates the controller's response and is therefore most likely the cause of oscillation.

For example, the following trend graph shows the PV and output signals for a loop controller with integral-dominant response. Both direct- and reverse-acting versions are shown:



Here, the output waveform is phase-shifted 90° behind (lagging) the PV waveform compared to what it would be by proportional action alone. This -90° phase shift is most difficult to see for reverse-acting controllers, since the natural 180° phase shift caused by reverse action makes the additional -90° shift look like a $+90^\circ$ shift (i.e. a -270° shift). One method I use for discerning the direction of phase shift for reverse acting controllers is to imagine what the output waveform would look like for a proportional-only controller, and compare to that. Another method is to perform a “thought experiment” looking at the PV waveform, noting the velocity of the controller's output at points of maximum error (when integral action would be expected to move the output signal fastest). In the previous example of an integral-dominant reverse-acting controller, we see that the output signal is indeed descending at its most rapid pace when the PV is highest above setpoint (greatest positive error), which is exactly what we would expect reverse-acting integration to do.


Mathematically, this makes sense as well. Integration of a sine function results in a negative cosine waveform:

$$-\cos t = \int \sin t \, dt$$

Another way of stating this is to say integration always adds a -90° phase shift:

$$\sin(t - 90^\circ) = \int \sin t \, dt$$

If derivative action is set too aggressive for the needs of the process, the resulting output waveform will be phase-shifted $+90^\circ$ from that of a purely proportional response:



Once again, the direction of phase shift is easiest to discern in the direct-acting case, since we expect the PV and output sine waves to be perfectly in-phase for proportional-dominant response. A $+90^\circ$ phase shift is very clear to see in the direct-acting example because the peaks of the output waveform clearly precede the corresponding peaks of the PV waveform. The reverse-acting control example is more difficult due to the added 180° of phase shift intrinsic to reverse action.

The derivative of a sine function is always a cosine function (or alternatively stated, a sine function with a $+90^\circ$ phase shift):


$$\cos t = \frac{d}{dt} \sin t$$

$$\sin(t + 90^\circ) = \frac{d}{dt} \sin t$$

You may encounter cases where *multiple* PID terms are set too aggressively for the needs of the process, in which case the phase shift will be split somewhere between 0° and $\pm 90^\circ$ owing to the combined actions. In such cases one must “round up” or “round down” the phase shift to the nearest value of -90° , 0° , or $+90^\circ$ in order to determine which of the three actions (I, P, or D, respectively) is most responsible for the oscillations.

30.4.4 Recognizing a “porpoising” controller

An interesting case of over-tuning is when the process variable “porpoises³⁶” on its way to setpoint following a step-change in setpoint. The following trend shows such a response:



“Porpoising” is universally poor behavior for a loop, because it combines the negative consequences of over-tuning (instability and excessive valve travel) with the negative consequence of under-tuning (delay achieving setpoint). There is no practical purpose served by a loop “porpoising,” and so this behavior should be avoided if at all possible.

Thankfully, identifying the cause of “porpoising” is rather easy to do. Only two control actions are capable of causing this response: proportional and derivative. Integral action simply *cannot* cause porpoising. In order for the process variable to “porpoise,” the controller’s output signal must reverse direction before the process variable ever reaches setpoint. Integral action, however, will always drive the output in a consistent direction when the process variable is on one side of setpoint. Only proportional and derivative actions are capable of producing a directional change in the output signal prior to reaching setpoint.

Solely examining the process variable waveform will not reveal whether it is proportional action, derivative action, or both responsible for the “porpoising” behavior. A trial reduction in the derivative³⁷ tuning parameter is one way to identify the culprit, as is phase-shift analysis between the PV and output waveforms during the “porpoising” period.

³⁶The term “porpoise” comes from the movements of a porpoise swimming rapidly toward the water’s surface as it chases along the bow of a moving ship. In order to generate speed, the animal undulates its body up and down to powerfully drive forward with its horizontal tail, tracing a sinusoidal path on its way up to breaching the surface of the water.

³⁷You could try reducing the controller’s gain as a first step, but if the controller implements the Ideal or Series algorithm, reduction in gain will *also* reduce derivative action, which may mask an over-tuned derivative problem.

30.5 Tuning techniques compared

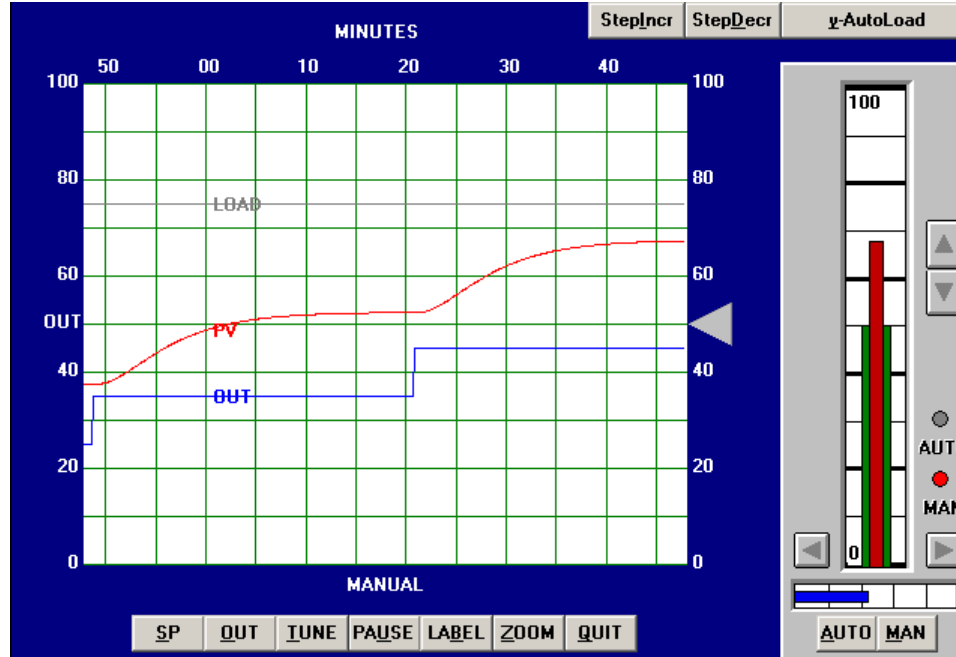
In this section I will show screenshots from a process loop simulation program illustrating the effectiveness of Ziegler-Nichols open-loop (“Reaction Rate”) and closed-loop (“Ultimate”) PID tuning methods, and then contrast them against the results of my own heuristic tuning. As you will see in some of these cases, the results obtained by either Ziegler-Nichols method tends toward instability (excessive oscillation of the process variable following a setpoint change). This is not necessarily an indictment of Ziegler’s and Nichols’ recommendations as much as it is a demonstration of the power of understanding. Ziegler and Nichols presented a simple step-by-step procedure for obtaining *approximate* PID tuning constant values based on closed-loop and open-loop process responses, which could be applied by anyone regardless of their level of understanding PID control theory. If I were tasked with drafting a procedure to instruct anyone to quantitatively determine PID constant values without an understanding of process dynamics or process control theory, I doubt my effort would be an improvement. Ultimately, robust PID control is attainable only at the hands of someone who understands how PID works, what each mode does (and why), and is able to distinguish between intrinsic process characteristics and instrument limitations. The purpose of this section is to clearly demonstrate the limitations of ignorantly-followed procedures, and contrast this “mindless” approach against the results of simple experimentation directed by qualitative understanding.

Each of the examples illustrated in this section were simulations run on a computer program called *PC-ControLab* developed by Wade Associates, Inc. Although these are simulated processes, in general I have found similar results using both Ziegler-Nichols and heuristic tuning methods on real processes. The control criteria I used for heuristic tuning were fast response to setpoint changes, with minimal overshoot or oscillation.

30.5.1 Tuning a “generic” process

Ziegler-Nichols open-loop tuning procedure

The first process tuned in simulation was a “generic” process, unspecific in its nature or application. Performing an open-loop test (two 10% output step-changes made in manual mode, both increasing) on this process resulted in the following behavior:




From the trend, we can see that this process is self-regulating, with multiple lags and some dead time. The reaction rate (R) is 20% over 15 minutes, or 1.333 percent per minute. Dead time (L) appears to be approximately 2 minutes. Following the Ziegler-Nichols recommendations for PID tuning based on these process characteristics (also including the 10% step-change magnitude Δm):

$$K_p = 1.2 \frac{\Delta m}{RL} = 1.2 \frac{10\%}{\frac{20\%}{15 \text{ min}} 2 \text{ min}} = 4.5$$

$$\tau_i = 2L = (2)(2 \text{ min}) = 4 \text{ min}$$

$$\tau_d = 0.5L = (0.5)(2 \text{ min}) = 1 \text{ min}$$

Applying the PID values of 4.5 (gain), 4 minutes per repeat (integral), and 1 minute (derivative) gave the following result in automatic mode (with a 10% setpoint change):




The result is reasonably good behavior with the PID values predicted by the Ziegler-Nichols open-loop equations, and would be acceptable for applications where some setpoint overshoot were tolerable.

We may tell from analyzing the phase shift between the PV and OUT waveforms that the dominant control action here is proportional: each negative peak of the PV lines up fairly close with each positive peak of the OUT, for this reverse-acting controller. If we were interested in minimizing overshoot and oscillation, the logical choice would be to reduce the gain value somewhat.

Ziegler-Nichols closed-loop tuning procedure

Next, the closed-loop, or “Ultimate” tuning method of Ziegler and Nichols was applied to this process. Eliminating both integral and derivative control actions from the controller, and experimenting with different gain (proportional) values until self-sustaining oscillations of consistent amplitude³⁸ were obtained, gave a gain value of 11:



From the trend, we can see that the ultimate period (P_u) is approximately 7 minutes in length. Following the Ziegler-Nichols recommendations for PID tuning based on these process characteristics:

$$K_p = 0.6K_u = (0.6)(11) = 6.6$$

$$\tau_i = \frac{P_u}{2} = \frac{7 \text{ min}}{2} = 3.5 \text{ min}$$


$$\tau_d = \frac{P_u}{8} = \frac{7 \text{ min}}{8} = 0.875 \text{ min}$$

It should be immediately apparent that these tuning parameters will yield poor control. While the integral and derivative values are close to those predicted by the open-loop (Reaction Rate) method, the gain value calculated here is even larger than what was calculated before. Since we

³⁸The astute observer will note the presence of some limiting (saturation) in the output waveform, as it attempts to go below zero percent. Normally, this is unacceptable while determining the ultimate gain of a process, but here it was impossible to make the process oscillate at consistent amplitude without saturating on the output signal. The gain of this process falls off quite a bit at the ultimate frequency, such that a high controller gain is necessary to sustain oscillations, causing the output waveform to have a large amplitude.

know proportional action was excessive in the last tuning attempt, and this one recommends an even higher gain value, we can expect our next trial to oscillate even worse.

Applying the PID values of 6.6 (gain), 3.5 minutes per repeat (integral), and 0.875 minute (derivative) gave the following result in automatic mode:




This time the loop stability is a bit worse than with the PID values given by the Ziegler-Nichols open-loop tuning equations, owing mostly to the increased controller gain value of 6.6 (versus 4.5). Proportional action is still the dominant mode of control here, as revealed by the minimal phase shift between PV and OUT waveforms (ignoring the 180 degrees of shift inherent to the controller's reverse action).

In all fairness to the Ziegler-Nichols technique, the excessive controller gain value probably resulted more from the saturated output waveform than anything else. This led to more controller gain being necessary to sustain oscillations, leading to an inflated K_p value.

Heuristic tuning procedure

From the initial open-loop (manual output step-change) test, we could see this process contains multiple lags in addition to about 2 minutes of dead time. Both of these factors tend to limit the amount of gain we can use in the controller before the process oscillates. Both Ziegler-Nichols tuning attempts confirmed this fact, which led me to try much lower gain values in my initial heuristic tests. Given the self-regulating nature of the process, I knew the controller needed integral action, but once again the aggressiveness of this action would be necessarily limited by the lag and dead times. Derivative action, however, would prove to be useful in its ability to help “cancel” lags, so I suspected my tuning would consist of relatively tame proportional and integral values, with a relatively aggressive derivative value.

After some experimenting, the values I arrived at were 1.5 (gain), 10 minutes (integral), and 5 minutes (derivative). These tuning values represent a proportional action only one-third as aggressive as the least-aggressive Ziegler-Nichols recommendation, an integral action less than half as aggressive as the Ziegler-Nichols recommendations, and a derivative action *five times* more aggressive than the most aggressive Ziegler-Nichols recommendation. The results of these tuning values in automatic mode are shown here:




With this PID tuning, the process responded with much less overshoot of setpoint than with the results of either Ziegler-Nichols technique.

30.5.2 Tuning a liquid level process

Ziegler-Nichols open-loop tuning procedure

The next simulated process I attempted to tune was a liquid level-control process. Performing an open-loop test (one 10% increasing output step-change, followed by a 10% decreasing output step-change, both made in manual mode) on this process resulted in the following behavior:




From the trend, the process appears to be purely integrating, as though the control valve were throttling the flow of liquid into a vessel with a constant out-flow. The reaction rate (R) on the first step-change is 50% over 10 minutes, or 5 percent per minute. Dead time (L) appears virtually nonexistent, estimated to be 0.1 minutes simply for the sake of having a dead-time value to use in the Ziegler-Nichols equations. Following the Ziegler-Nichols recommendations for PID tuning based on these process characteristics (also including the 10% step-change magnitude Δm):

$$K_p = 1.2 \frac{\Delta m}{RL} = 1.2 \frac{10\%}{\frac{50\%}{10 \text{ min}} 0.1 \text{ min}} = 24$$

$$\tau_i = 2L = (2)(0.1 \text{ min}) = 0.2 \text{ min}$$

$$\tau_d = 0.5L = (0.5)(0.1 \text{ min}) = 0.05 \text{ min}$$

Applying the PID values of 24 (gain), 0.2 minutes per repeat (integral), and 0.05 minutes (derivative) gave the following result in automatic mode:



The process variable certainly responds rapidly to the five increasing setpoint changes and also to the one large decreasing setpoint change, but the valve action is hopelessly chaotic. Not only would this “jittery” valve motion prematurely wear out the stem packing, but it would also result in vast over-consumption of compressed air to continually stroke the valve from one extreme to the other. Furthermore, we see evidence of “overshoot” at every setpoint change, most likely from excessive integral action.

We can see from the valve’s wild behavior even during periods when the process variable is holding at setpoint that the problem is not a loop oscillation, but rather the effects of process noise on the controller. The extremely high gain value of 24 is amplifying PV noise by that factor, and reproducing it on the output signal.


Ziegler-Nichols closed-loop tuning procedure

Next, I attempted to perform a closed-loop “Ultimate” gain test on this process, but I was not successful. Even the controller’s maximum possible gain value would not generate oscillations, due to the extremely crisp response of the process (minimal lag and dead times) and its integrating nature (constant phase shift of -90°).

Heuristic tuning procedure

From the initial open-loop (manual output step-change) test, we could see this process was purely integrating. This told me it could be controlled primarily by proportional action, with very little integral action required, and no derivative action whatsoever. The presence of some process noise is the only factor limiting the aggressiveness of proportional action. With this in mind, I experimented with increasingly aggressive gain values until I reached a point where I felt the output signal noise was at a maximum acceptable limit for the control valve. Then, I experimented with integral action to ensure reasonable elimination of offset.

After some experimenting, the values I arrived at were 3 (gain), 10 minutes (integral), and 0 minutes (derivative). These tuning values represent a proportional action only one-eighth as aggressive as the Ziegler-Nichols recommendation, and an integral action *fifty times* less aggressive than the Ziegler-Nichols recommendation. The results of these tuning values in automatic mode are shown here:



You can see on this trend five 10% increasing setpoint value changes, with crisp response every time, followed by a single 50% decreasing setpoint step-change. In all cases, the process response clearly meets the criteria of rapid attainment of new setpoint values and no overshoot or oscillation.


If it was decided that the noise in the output signal was too detrimental for the valve, we would have the option of further reducing the gain value and (possibly) compensating for slow offset recovery with more aggressive integral action. We could also attempt the insertion of a damping constant into either the level transmitter or the controller itself, so long as this added lag did not cause oscillation problems in the loop³⁹. The best solution would be to find a way to isolate the level transmitter from noise, so that the process variable signal was much “quieter.” Whether or not this is possible depends on the process and on the particular transmitter used.

³⁹We would have to be *very* careful with the addition of damping, since the oscillations could create may not appear on the trend. Remember that the insertion of damping (low-pass filtering) in the PV signal is essentially an act of “lying” to the controller: telling the controller something that differs from the real, measured signal. If our PV trend shows us this damped signal and not the “raw” signal from the transmitter, it is possible for the process to oscillate and the PV trend to be deceptively stable!

30.5.3 Tuning a temperature process

Ziegler-Nichols open-loop tuning procedure

This next simulated process is a temperature control process. Performing an open-loop test (two 10% increasing output step-changes, both made in manual mode) on this process resulted in the following behavior:



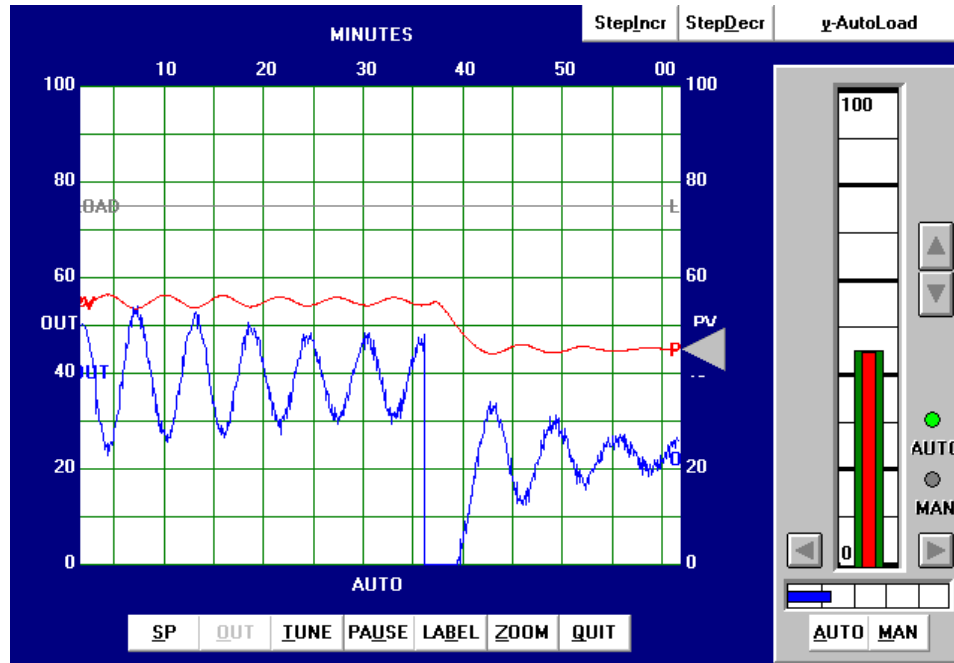
From the trend, the process appears to be self-regulating with a slow time constant (lag) and a substantial dead time. The reaction rate (R) on the first step-change is 30% over 30 minutes, or 1 percent per minute. Dead time (L) looks to be approximately 1.25 minutes. Following the Ziegler-Nichols recommendations for PID tuning based on these process characteristics (also including the 10% step-change magnitude Δm):

$$K_p = 1.2 \frac{\Delta m}{RL} = 1.2 \frac{10\%}{\frac{30\%}{30 \text{ min}} 1.25 \text{ min}} = 9.6$$

$$\tau_i = 2L = (2)(1.25 \text{ min}) = 2.5 \text{ min}$$

$$\tau_d = 0.5L = (0.5)(1.25 \text{ min}) = 0.625 \text{ min}$$


Applying the PID values of 9.6 (gain), 2.5 minutes per repeat (integral), and 0.625 minutes (derivative) gave the following result in automatic mode:



As you can see, the results are quite poor. The PV is still oscillating with a peak-to-peak amplitude of almost 20% from the last process upset at the time of the 10% downward SP change. Additionally, the output trend is rather noisy, indicating excessive amplification of process noise by the controller.

Ziegler-Nichols closed-loop tuning procedure

Next, the closed-loop, or “Ultimate” tuning method of Ziegler and Nichols was applied to this process. Eliminating both integral and derivative control actions from the controller, and experimenting with different gain (proportional) values until self-sustaining oscillations of consistent amplitude were obtained, gave a gain value of 15:




From the trend, we can see that the ultimate period (P_u) is approximately 5.2 minutes in length. Following the Ziegler-Nichols recommendations for PID tuning based on these process characteristics:

$$K_p = 0.6K_u = (0.6)(15) = 9$$

$$\tau_i = \frac{P_u}{2} = \frac{5.2 \text{ min}}{2} = 2.6 \text{ min}$$

$$\tau_d = \frac{P_u}{8} = \frac{5.2 \text{ min}}{8} = 0.65 \text{ min}$$

These PID tuning values are quite similar to those predicted by the open loop (“Reaction Rate”) method, and so we would expect to see very similar results:




As expected, we still see excessive oscillation following a 10% setpoint change, as well as excessive “noise” in the output trend.

Heuristic tuning procedure

From the initial open-loop (manual output step-change) test, we could see this process was self-regulating with a slow lag and substantial dead time. The self-regulating nature of the process demands at least some integral control action to eliminate offset, but too much will cause oscillation given the long lag and dead times. The existence of over 1 minute of process dead time also prohibits the use of aggressive proportional action. Derivative action, which is generally useful in overcoming lag times, will cause problems here by amplifying process noise. In summary, then, we would expect to use mild proportional, integral, *and* derivative tuning values in order to achieve good control with this process. Anything too aggressive will cause problems for this process.

After some experimenting, the values I arrived at were 3 (gain), 5 minutes (integral), and 0.5 minutes (derivative). These tuning values represent a proportional action only one-third as aggressive as the Ziegler-Nichols recommendation, and an integral action about half as aggressive as the Ziegler-Nichols recommendation. The results of these tuning values in automatic mode are shown here:



As you can see, the system's response has almost no overshoot (with either a 10% setpoint change or a 15% setpoint change) and very little "noise" on the output trend. Response to setpoint changes is relatively crisp considering the naturally slow nature of the process: each new setpoint is achieved within about 7.5 minutes of the step-change.


30.6 Note to students

Learning how to tune PID controllers is a skill born of much practice. Regardless of how thoroughly you may study the subject of PID control on paper, you really do not understand it until you have spent a fair amount of time actually tuning real controllers.

In order to gain this experience, though, you need access to working processes and the freedom to disturb those processes over and over again. If your school's lab has several "toy" processes built to facilitate this type of learning experience, that is great. However, your learning will grow even more if you have a way to practice PID tuning at your own convenience.

30.6.1 Electrically simulating a process

Thankfully, there is a relatively simple way to build your own “process” for PID tuning practice. First, you need to obtain an electronic single-loop PID controller⁴⁰ and connect it to a resistor-capacitor network such as this:




The $250\ \Omega$ resistor converts the controller’s 4-20 mA signal into a 1-5 VDC signal, which then drives the passive integrator (lag) RC networks. The two stages of RC “lag” simulate a self-regulating process with a second-order lag and a steady-state gain of 1. The potentiometers establish the lag times for each stage, providing a convenient way to alter the process characteristics for more tuning practice. Feel free to extend the circuit with additional RC lag networks for even more delay (and an even harder-to-tune process!).

Since this simulated “process” is direct-acting (i.e. increasing manipulated variable signal results in an increasing process variable signal), the controller must be configured for *reverse* action (i.e. increasing process variable signal results in a decreasing manipulated variable signal) in order to achieve negative feedback. You are welcome to configure the controller for direct action just to see what the effects will be, but I assure you control will be impossible: the PV will saturate beyond 100% or below 0% no matter how the PID values are set.


⁴⁰Many instrument manufacturers sell simple, single-loop controllers for reasonable prices, comparable to the price of a college textbook. You need to get one that accepts 1-5 VDC input signals and generates 4-20 mA output signals, and has a “manual” mode of operation in addition to automatic – these features are *very important!* Avoid controllers that can only accept thermocouple inputs, and/or only have time-proportioning (PWM) outputs. Additionally, I strongly recommend you take the time to experimentally learn the actions of proportional, integral, and derivative as outlined in section 29.16 beginning on page 2399 before you embark on any PID tuning exercises.

30.6.2 Building a “Desktop Process” unit

A more sophisticated approach to gaining hands-on experience tuning PID controllers is to actually build a working “process” that the controller can regulate. A relatively simple way to do this for students is to build what I like to call *Desktop Processes*, where a loop controller is used to control the speed of a motor/generator set made from small DC “hobby” electric motors. An illustration of a “Desktop Process” is shown here:





You must build your own variable-speed drive (VSD) circuit to convert the controller's 4-20 mA output signal into a DC voltage powerful enough to drive the motor. This same circuit should also contain components for "scaling" and filtering the tachogenerator's DC voltage signal so it may be read by the controller's input. Fortunately, the following circuit is a proven and simple design for doing just that:



Diodes included in this design protect against reverse-polarity power supply connections and inductive "kickback" resulting from de-energizing inductive loads.

Photographs showing a complete “Desktop Process” unit in operation, including close-ups of the motor/generator set and the variable speed drive circuit board, appear here:



As you can see from the photograph, the motor and generator are held in a short length of split

PVC pipe. This is a simple way to clamp and align both machines so their shafts turn on the same centerline. The coupling between the two shafts is nothing more than a piece of rubber tube (or wire insulation, or heat-shrink tubing, or even electrical tape!).

An optional accessory to add to a Desktop Process is a data acquisition unit capable of measuring the DC voltage motor speed and controller output signals, plotting them on a computer display for further analysis. This becomes very useful when fine-tuning PID response, allowing students to visually recognize oscillation, overshoot, windup, and other phenomena of closed-loop control.

The controller model shown in these photographs happens to be a Siemens 353, but any loop controller capable of receiving a 1-5 volt DC input signal and generating 4-20 mA DC output signal will work just fine. In fact, I've connected this very same VSD and motor/generator set to different controllers⁴¹ to compare operation.


Interesting experiments to perform with a Desktop Process – other than PID tuning practice – include the following:

- Introducing process loads by touching the spinning motor shaft (slowing it down using your finger) and compare the responses between the controller's "manual" and "automatic" modes. This proves to be a very effective way for students to comprehend the difference between these two modes of operation. I have yet to encounter a student who does not immediately grasp the concept after doing this experiment for themselves, feeling the motor's shaft speed respond to their finger load in both modes, also watching the controller's output response.
- Switching the controller mode from reverse action to direct action to see how a process "runs away" when the loop feedback is positive rather than negative.
- Switching the VSD action from direct to reverse, then reconfiguring the controller's action to complement it, maintaining negative feedback in the system.
- Try switching between auto and manual modes in the controller, comparing the response with and without the feature of *setpoint tracking*. Again, this is a concept many students struggle to grasp in theory, but immediately comprehend when they see it in action.

⁴¹Among these different controllers were a Distech ESP-410 building (HVAC) controller and a small PLC programmed with a custom PID control algorithm. In fact, a Desktop Process is ideal for courses where students create their own control algorithms in PLC or data acquisition hardware. The significance of controller scan rate becomes very easy to comprehend when controlling a process like this with such a short time constant. The contrast between a DDC controller with a 500 millisecond scan rate and a PLC with a 50 millisecond scan rate, for example, is marked.

30.6.3 Simulating a process by computer

A fascinating solution for realistic PID tuning in the classroom was offered to me by Blair MacNeil of Cape Breton University (located in the town of Sydney, on the island of Nova Scotia, Canada) in 2010 by way of email correspondence. Professor MacNeil uses Moore 353 loop controllers connected to analog computer I/O (“data acquisition”) modules, with a personal computer running VisSim Realtime software to simulate the dynamics of a real process. With the power of a personal computer simulating the process, virtually any process dynamic (as well as any instrument fault) may be generated for the benefit of the loop controller to control:



This approach provides realistic process dynamics for the loop controller to manage, yet requires little in the way of capital expense or physical space to implement. Different process models, instrument faults, and control strategies may be easily implemented in the personal computer’s software, making it far more flexible as a teaching tool than any analog electronic simulation network or real process connected to the controller. It is also completely safe to operate, with absolutely no danger of harming anything in the event of a process “excursion” or other upset.

30.7 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Conservation of mass:** mass is an intrinsic property of matter, and as such cannot be created or destroyed. Relevant to the *mass balance* of a process, meaning that all mass flowing into a process must equal all mass flowing out of a process (unless mass is being accumulated or released from the process). This is relevant to the determination of a process' characteristic as either self-regulating or integrating: whether the mass balance of the process naturally equalizes or not as the process variable changes value.
- **Conservation of energy:** energy cannot be created or destroyed, only converted between different forms. Relevant to the *energy balance* of a process, meaning that all energy flowing into a process must equal all energy flowing out of a process (unless energy is being accumulated or released from the process). This is relevant to the determination of a process' characteristic as either self-regulating or integrating: whether the energy balance of the process naturally equalizes or not as the process variable changes value.
- **Negative feedback:** when the output of a system is degeneratively fed back to the input of that same system, the result is decreased (overall) gain and greater stability. Relevant to loop controller action: in order for a control system to be stable, the feedback must be negative.
- **Amplification:** the control of a relatively large signal by a relatively small signal. Relevant to the role of loop controllers exerting influence over a process variable at the command of a measurement signal. In behaving as amplifiers, loop controllers may oscillate if certain criteria are met.
- **Barkhausen criterion:** is overall loop gain is unity (1) or greater, and phase shift is 360° , the loop will sustain oscillations. Relevant to control system stability, explaining why the loop will "cycle" (oscillate) if gain is too high.
- **Time constant:** (τ), defined as the amount of time it takes a system to change 63.2% of the way from where it began to where it will eventually stabilize. Also known as the *lag time* of a system. The system will be within 1% of its final value after 5 time constants' worth of time has passed (5τ). Relevant to process control loops, where natural lags contribute to time constants, usually of multiple order.
- **Phase shift:** the angular difference between two sinusoidal waves of the same frequency. Relevant to the analysis of controller trend graphs: zero phase shift between PV and Output is the hallmark of proportional action; lagging phase shift is the hallmark of integral action; leading phase shift is the hallmark of derivative action.
- **Differentiation (calculus):** where one variable is proportional to the rate-of-change of two others. Differentiation always results in a positive phase shift if the input is a wave. Relevant

to the output of a controller, for determining by (leading) phase shift whether derivative action is dominant.

- **Integration (calculus):** where one variable is proportional to the accumulation of the product of two others. Integration always results in a negative phase shift if the input is a wave. Relevant to the output of a controller, for determining by (lagging) phase shift whether integral action is dominant.

References

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Control Volume II*, Third Edition, CRC Press, Boca Raton, FL, 1999.

McMillan, Greg, “Is Wireless Process Control Ready for Prime Time?”, *Control*, May 2009, pp. 54-57.

Mollenkamp, Robert A., *Introduction to Automatic Process Control*, Instrument Society of America, Research Triangle Park, NC, 1984.

Palm, William J., *Control Systems Engineering*, John Wiley & Sons, Inc., New York, NY, 1986.

Shinskey, Francis G., *Energy Conservation through Control*, Academic Press, New York, NY, 1978.

Shinskey, Francis G., *Process-Control Systems – Application / Design / Adjustment*, Second Edition, McGraw-Hill Book Company, New York, NY, 1979.


St. Clair, David W., *Controller Tuning and Control Loop Performance, a primer*, Straight-Line Control Company, Newark, DE, 1989.

Ziegler, J. G., and Nichols, N. B., “Optimum Settings for Automatic Controllers”, *Transactions of the American Society of Mechanical Engineers (ASME)*, Volume 64, pages 759-768, Rochester, NY, November 1942.

Chapter 31

Basic process control strategies

In a simple control system, a process variable (PV) is measured and compared with a setpoint value (SP). A manipulated variable (MV, or output) signal is generated by the controller and sent to a final control element, which then influences the process variable to achieve stable control. The algorithm by which the controller develops its output signal is typically PID (Proportional-Integral-Derivative), but other algorithms may be used as well:




This form of simple control may be improved upon and expanded for a greater range of process applications by interconnecting multiple controllers and/or redirecting measurement and control signals in more complex arrangements. An exploration of some of the more common control system configurations is the subject of this chapter.

31.1 Supervisory control

In a manually-controlled process, a human operator directly actuates some form of final control element (usually a valve) to influence a process variable. Simple automatic (“regulatory”) control relieves human operators of the need to continually adjust final control elements by hand, replacing this task with the occasional adjustment of setpoint values. The controller then manipulates the final control element to hold the process variable at the setpoint value determined by the operator.

The next step in complexity after simple automatic control is to automate the adjustment of the setpoint for a process controller. A common implementation of this concept is the automatic cycling of setpoint values according to a timed schedule. An example of this is a temperature controller for a heat-treatment furnace used to temper metal samples:



Here, a computer “supervises” the furnace’s temperature by communicating setpoint values to the temperature indicating controller (TIC) over a digital network interface such as Ethernet. From the temperature controller’s perspective, this is a *remote* setpoint signal, as opposed to a *local* setpoint value which would be set by a human operator at the controller faceplate. Since the heat-treatment of metals requires particular temperature ranges and rates of change over time, this control system relieves the human operator of having to manually adjust setpoint values again and again during heat-treatment cycles. Instead, the computer schedules different setpoint values at different times (even setpoint values that change steadily at a certain rate over a period of time) according to the needs of the particular metal type and treatment type. Such a control scheme is quite common for heat-treating processes, and it is referred to as *ramp and soak*¹.

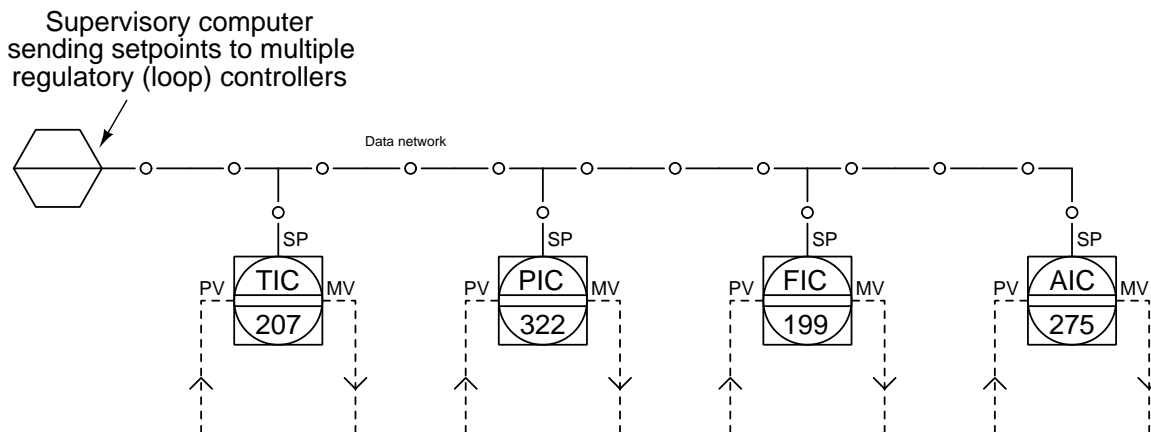
¹In honor of the system’s ability to slowly “ramp” temperature up or down at a specified rate, then “soak” the metal at a constant temperature for set periods of time. Many single-loop process controllers have the ability to perform ramp-and-soak setpoint scheduling without the need of an external “supervisory” computer.

Process controllers configured for supervisory setpoint control typically have three operating modes:

- **Manual mode:** Controller takes no automatic action. Output value set by human operator.
- **Automatic mode with local SP:** Controller automatically adjusts its output to try to keep $PV = SP$. Setpoint value set “locally” by human operator.
- **Automatic mode with remote SP:** Controller automatically adjusts its output to try to keep $PV = SP$. Setpoint value set “remotely” by supervising computer.

Supervisory setpoint control is also used in the chemical processing industries to optimize production efficiencies by having a powerful computer provide setpoint adjustments to regulatory controls based on mathematical models of the process and optimization constraints. In simple terms, this means having a computer make setpoint adjustments to the normal PID loop controllers instead (or in addition to) human operators making setpoint changes. This forms a two-layer process control system: the “base” or “regulatory” layer of control (PID loop controllers) and the “high” or “supervisory” level of control (the powerful computer with the mathematical process models).

Such “optimizing” control systems are usually built over a digital network for reasons of convenience. A single network cable not only is able to communicate the frequent setpoint changes from the supervisory computer to the multitude of process loop controllers, but it may also carry process variable information from those controllers back to the supervisory computer so it has data for its optimization algorithms to operate on:




The complexity of these optimization algorithms is limited only by the computational power of the supervisory computer and the creativity of the programmers and engineers who implement it. A modern trend in process optimization for industries able to produce varying proportions of different products from the same raw material feed is to have computer algorithms select and optimize production not only for maximum cost efficiency, but also for maximum market sales and minimum storage of volatile product².

²I once attended a meeting of industry representatives where one person talked at length about a highly automated

31.2 Cascade control

A simple control system drawn in block diagram form looks like this:




Information from the measuring device (e.g. transmitter) goes to the controller, then to the final control device (e.g. control valve), influencing the process which is sensed again by the measuring device. The controller's task is to inject the proper amount of negative feedback such that the process variable stabilizes over time. This flow of information is collectively referred to as a feedback control *loop*.

To *cascade* controllers means to connect the output signal of one controller to the setpoint of another controller, with each controller sensing a different aspect of the same process. The first controller (called the *primary*, or *master*) essentially “gives orders” to the second controller (called the *secondary* or *slave*) via a *remote setpoint* signal.

lumber mill where logs were cut into lumber not only according to minimum waste, but also according to the real-time market value of different board types and stored inventory. The joke was, if the market value of wooden toothpicks suddenly spiked up, the control system would shred every log into toothpicks in an attempt to maximize profit!

Thus, a cascade control system consists of two feedback control loops, one nested inside the other:




A very common example of cascade control is a *valve positioner*, which receives a command signal from a regular process controller, and in turn works to ensure the valve stem position precisely matches that command signal. The control valve’s stem position is the process variable (PV) for the positioner, just as the command signal is the positioner’s setpoint (SP). Valve positioners therefore act as “slave” controllers to “master” process controllers controlling pressure, temperature, flow, or some other process variable.

The purpose of cascade control is to achieve greater stability of the primary process variable by regulating a secondary process variable in accordance with the needs of the first. An essential requirement of cascaded control is that the secondary process variable be faster-responding (i.e. shorter lag and dead times) than the primary process variable.

An analogy for understanding cascade control is that of *delegation* in a work environment. If a supervisor delegates some task to a subordinate, and that subordinate performs the task without further need of guidance or assistance from the supervisor, the supervisor’s job is made easier. The subordinate takes care of all the little details that would otherwise burden the supervisor if the supervisor had no one to delegate to. This analogy also makes it clear why the secondary process variable must be faster-responding than the primary process variable: the supervisor-subordinate management structure fails to work if the supervisor does not maintain focus on long-term goals (i.e. longer-term than the completion time of the tasks given to subordinates). If a supervisor focuses on achieving goals that are shorter-term than the time required for subordinates to complete


their assignments, the supervisor will inevitably call for “course changes” that are too quick for the subordinates to execute. This will lead to the subordinates “lagging” behind the supervisor’s orders, to the detriment of everyone’s satisfaction.

An example of cascade control applied to a real industrial process is shown here, for a *dryer* system where heated air is used to evaporate water from a granular solid. The primary process variable is the outlet air exiting the dryer, which should be maintained at a high enough temperature to ensure water will not remain in the upper layers of the solid material. This outlet temperature is fairly slow to react, as the solid material mass creates a large lag time:



There are several parameters influencing the temperature of the outlet air other than the moisture content of the drying material. These include air flow, ambient air temperature, and variations in steam temperature. Each one of these variables is a *load* on the process variable we are trying to control (outlet air temperature). If any of these parameters were to suddenly change, the effect would be slow to register at the outlet temperature even though there would be immediate impact at the bottom of the dryer where the heated air enters. Correspondingly, the control system would be slow to correct for any of these changing loads.

We may better compensate for these loads by installing a second temperature transmitter at the inlet duct of the dryer, with its own controller to adjust steam flow at the command of the primary controller:




Now, if any of the loads related to incoming air flow or temperature vary, the secondary controller (TC-1b) will *immediately* sense the change in dryer inlet temperature and compensate by adjusting steam flow through the heat exchanger. Thus, the “slave” control loop (1b) helps stabilize the “master” control loop (1a) by reacting to load changes long before any effect might manifest at the dryer outlet.

A helpful way to think of this is to consider the slave controller as *shielding* the master controller from the loads previously mentioned (incoming air flow, ambient temperature, and steam temperature). Of course, these variables still act as loads to the slave controller, as it must continuously adjust the steam valve to compensate for changes in air flow, ambient air temperature, and steam temperature. However, so long as the slave controller does a good job of stabilizing the air temperature entering the dryer, the master controller will never “see” the effects of those load changes. Responsibility for incoming air temperature has been delegated to the slave controller, and as a result the master controller is conveniently isolated from the loads impacting that loop.

To re-emphasize an important point, one of the non-negotiable requirements for cascade control is that the secondary (slave) loop must be *faster-responding* than the primary (master) loop. Cascade control cannot function if this speed relationship is reversed. Temperature controller TC-1b is able to be a slave to controller TC-1a because the natural response time of the temperature at the dryer’s bottom is much shorter than at the dryer’s top with respect to any changes in steam valve position.

A common implementation of cascade control is where a flow controller receives a setpoint from some other process controller (pressure, temperature, level, analytical, etc.), fluid flow being one of the fastest-responding process types in existence. A feedwater control system for a steam boiler – shown here in pneumatic form – is a good example:




The “secondary” or “slave” flow controller works to maintain feedwater flow to the boiler at whatever flow rate is desired by the level controller. If feedwater pressure happens to increase or decrease, any resulting changes in flow will be quickly countered by the flow controller without the level controller having to react to a consequent upset in steam drum water level. Thus, cascade control works to guard against steam drum level instability resulting from changes in the feedwater flow caused by factors outside the boiler. As stated previously, the slave (flow) controller effectively *shields* the master (level) controller from loads in the feedwater supply system, so that master controller doesn’t have to deal with those loads.

This level/flow cascade control system also embodies the principle of the secondary (slave) loop being faster-responding than the primary (master) loop. Water flow is an inherently fast process, the flow rate responding immediately to changes in valve position. By contrast, water level is a much slower-responding type of process. If you perform a “thought experiment” where the feedwater valve is suddenly opened fully, it is easy to see that the feedwater flow rate will immediately reach its full

(100%) value while the steam drum's water level will merely begin to rise, taking time to reach its full (100%) value.

It is worth noting that the inclusion of a flow control “slave” loop to this boiler water level control system also helps to overcome a potential problem of the control valve: nonlinear behavior. In the control valves chapter, we explore the phenomenon of *installed valve characteristics* (Section 27.13.1 beginning on page 2190), specifically noting how changes in pressure drop across a control valve influences its throttling behavior. The result of these pressure changes is a non-linearization of valve response, such that the valve tends to be more responsive near its closed position and less responsive near its open position. One of the benefits of cascaded flow control is that this problem becomes confined to the secondary (flow control) loop, and is effectively removed from the primary control loop. To phrase it simply, distorted valve response becomes “the flow controller’s problem” rather than something the level controller must manage. The result is a level control system with more predictable response.

A classic example of cascade control strategy is found in *motion control* applications, where an electric motor is used as the final control element to precisely position a piece of machinery. In this capacity, the motor is usually called a *servo*. Robotic systems make extensive use of servo motors and cascaded control loops to modulate power to those motors. The following illustration shows a triple-cascade control system³ for a motor-actuated elevator, precisely controlling the position of the elevator through cascaded velocity and motor current control:



Hypothetically, the position of the elevator could be controlled with a single PID controller sensing platform position and directly sending power to the motor. However, much more precise control of the platform is achievable by sensing position, velocity, and motor current, and controlling each one of those variables with its own loop. In motion control systems, each successive variable is the time-derivative of its precursor. Velocity, for instance, is the time-derivative of position ($v = \frac{dx}{dt}$). Motor current, which is usually proportional to motor torque, which in turn is proportional to the angular acceleration (α) of the winch and consequently the linear acceleration of the platform (a), is the time-derivative of velocity ($a = \frac{dv}{dt}$). If it were not for cascading, a single PID controller would have to control position by manipulating the *acceleration* of the platform (i.e. motor current). This would make the process characteristic “runaway” in nature, as any fixed amount of current will cause the platform to accelerate⁴.

³Interestingly, servo motor control is one application where *analog* loop controllers have historically been favored over digital loop controllers, simply for their superior speed. An opamp-based P, PI, or PID controller is lightning-fast because it has no need to digitize any analog process variables (analog-to-digital conversion) nor does it require time for a clock to sequence step-by-step through a written program as a microprocessor does. Servomechanism processes are inherently fast-responding, and so the controller(s) used to control servos must be faster yet.

⁴At one specific current level, the motor will develop just enough torque to hold the platform’s weight, at which point the acceleration will be zero. Any amount of current above this value will cause an upward acceleration, while any amount of current below this value will cause a downward acceleration.

Here with servomechanisms we see how cascading not only has the effect of “shielding” certain load variables from the master controller’s view, but it also simplifies the dynamic characteristics of the process from that same point of view. Instead of the position controller having to regulate an inherently “runaway” process, it now sees the process as having an “integrating” characteristic, since any constant output signal from the position controller results in the platform holding to a constant velocity (i.e. platform position will change at a constant rate over time, rather than at an accelerating rate).

A necessary step in implementing cascade control is to ensure the secondary (“slave”) controller is well-tuned *before* any attempt is made to tune the primary (“master”) controller. Just a moment’s thought is all that is needed to understand why this precedence in tuning must be: it is a simple matter of dependence. The slave controller does not depend on good tuning in the master controller in order to control the slave loop. If the master controller were placed in manual (effectively turning off its automatic response), the slave controller would simply control to a constant setpoint. However, the master controller most definitely depends on the slave controller being well-tuned in order to fulfill the master’s “expectations.” If the slave controller were placed in manual mode, the master controller would not be able to exert any control over its process variable whatsoever. Clearly then, the slave controller’s response is essential to the master controller being able to control its process variable, therefore the slave controller should be tuned first when initially commissioning or optimizing a cascade control system.

Just like supervisory control systems where a process controller receives a “remote” setpoint signal from some other system, the secondary (“slave”) controller in a cascade system typically has three different operating modes:

- **Manual mode:** Controller takes no automatic action. Output value set by human operator.
- **Automatic mode:** Controller automatically adjusts its output to try to keep $PV = SP$. Setpoint value set “locally” by human operator.
- **Cascade mode:** Controller automatically adjusts its output to try to keep $PV = SP$. Setpoint value set “remotely” by primary (master) controller.


This means it is possible to defeat a cascade control system by placing the secondary controller in the wrong mode (automatic) just as it is possible to defeat any control system by placing the controller in manual mode. If a controller is “slaved” to another controller, it must be left in *cascade* mode in order for the control strategy to function as designed.

31.3 Ratio control


Most people reading this book have likely had the experience of adjusting water temperature using two hand valves as they took a shower: one valve controlling the flow of hot water and the other valve controlling the flow of cold water. In order to adjust water temperature, the *proportion* of one valve opening to the other must be changed. Increasing or decreasing total water flow rate without upsetting the outlet temperature is a matter of adjusting both valves in the same direction, maintaining that same proportion of hot to cold water flow.

Although you may not have given it much thought while taking your shower, you were engaged in a control strategy known as *ratio control*, where the ratio of one flow rate to another is controlled for some desired outcome. Many industrial processes also require the precise mixing of two or more ingredients to produce a desired product. Not only do these ingredients need to be mixed in proper proportion, but it is usually desirable to have precise control over the total flow rate as well.

A simple example of ratio control is in the production of paint, where a base liquid must be mixed with one or more pigments to achieve a desired consistency and color. A manually controlled paint mixing process, similar to the hot and cold water valve “process” in some home showers, is shown here. Two flowmeters, a ratio calculating relay, and a display provide the human operator with a live measurement of pigment-to-base ratio:




One alteration we could make to this mixing system is to link the two manual control valve handles together in such a way that the ratio of base to pigment was *mechanically* established. All the human operator needs to do now is move the one link to increase or decrease mixed paint production:




Adjusting the pigment-to-base ratio is now a matter of adjusting the linkage ratio, a task most likely performed by a mechanic or someone else skilled in the alignment of mechanical linkages. The convenience of total flow adjustment gained by the link comes at the price of inconvenient ratio adjustment.

Mechanical link ratio-control systems are commonly used to manage simple burners, proportioning the flow rates of fuel and air for clean, efficient combustion. A photograph of such a system appears here, showing how the fuel gas valve and air damper motions are coordinated by a single rotary actuator:



As you can see in this photo, the fuel gas valve is actuated by means of a cam, allowing precise “tuning” of the valve characteristics for consistent fuel/air ratio across a wide range of firing rates. Making ratio adjustments in such a linkage system is obviously a task for a skilled mechanic or technician.


A more automated approach to the general problem of ratio control involves the installation of a flow control loop on one of the lines and a flow-sensing transmitter on the other line. The signal coming from the uncontrolled flow transmitter becomes the setpoint for the flow control loop:



Here, the flow transmitter on the uncontrolled line measures the flow rate of base, sending a flow rate signal to the pigment flow controller which acts to match flow rates. If the calibrations of each flow transmitter are precisely equal to one another, the ratio of pigment to base will be 1:1 (equal). The flow of base liquid into the mixing system is called a *wild flow* or *wild variable*, since this flow rate is not controlled by the ratio control system. The only purpose served by the ratio control system is to match the pigment flow rate to the wild (base) flow rate, so the same ratio of pigment to base will always be maintained regardless of total flow rate. Thus, the flow rate of pigment will be held *captive* to match the “wild” base flow rate, which is why the controlled variable in a ratio system is sometimes called the *captive variable* (in this case, a *captive flow*).


As with the mechanically-linked manual ratio mixing system, this ratio control system provides convenient total flow control at the expense of convenient ratio adjustment. In order to alter the ratio of pigment to base, someone must re-range one or more flow transmitters. To achieve a 2:1 ratio of base to pigment, for example, the base flow transmitter’s range would have to be double that of the pigment flow transmitter. This way, an equal percentage of flow registered by both flow transmitters (as the ratio controller strives to maintain equal percentage values of flow between pigment and base) would actually result in twice the amount of base flow than pigment flow.

We may incorporate convenient ratio adjustment into this system by adding another component (or function block) to the control scheme: a device called a *signal multiplying relay* (or alternatively, a *ratio station*). This device (or computer function) takes the flow signal from the base (wild) flow transmitter and multiplies it by some constant value (k) before sending the signal to the pigment (captive) flow controller as a setpoint:



With identical flow range calibrations in both flow transmitters, this multiplying constant k directly determines the pigment-to-base ratio (i.e. the ratio will be 1:1 when $k = 1$; the ratio will be 2:1 when $k = 2$, etc.). If the k value is easily adjusted by a human operator, mixing ratio becomes a very simple parameter to change at will, just as the total production rate is easy to adjust by moving the base flow control valve.

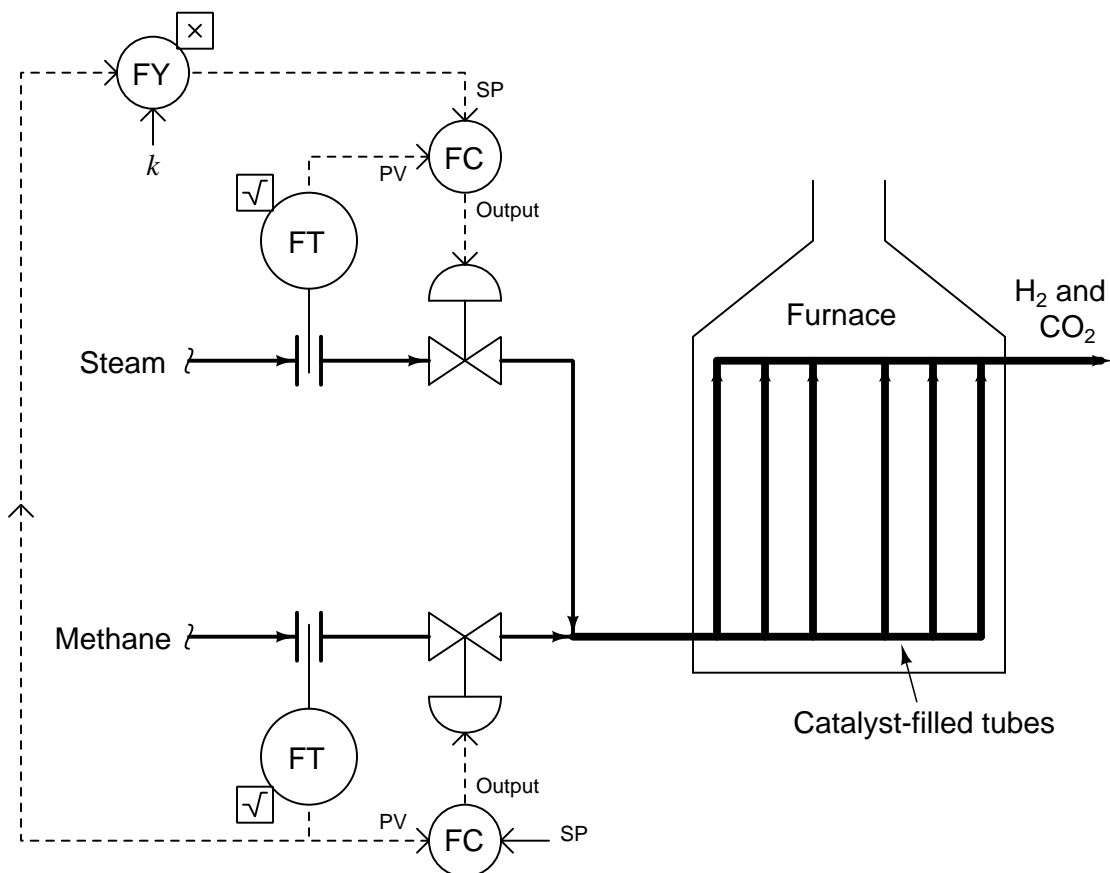
Another example of ratio control at work is in a process whereby hydrocarbon gases (usually methane) are converted into hydrogen gas and carbon dioxide gas. This is known as the *steam-hydrocarbon reforming process*, and it is one of the more popular ways of generating hydrogen gas for industrial use. The overall reaction for this process with methane gas (CH_4) and steam (H_2O) as the reactants is as follows⁵:



⁵The conversion from hydrocarbon and steam to hydrogen and carbon dioxide is typically a two-stage process: the first (*reforming*) stage produces hydrogen gas and carbon monoxide, while a second (*water-gas-shift*) stage adds more steam to convert the carbon monoxide into carbon dioxide with more hydrogen liberated. Both reactions are endothermic, with the reforming reaction being more endothermic than the water-gas-shift reaction.

This is an *endothermic chemical reaction*, which means a net input of energy is required to make it happen. Typically, the hydrocarbon gas and steam are mixed together in a heated environment in the presence of a catalyst (to reduce the activation energy requirements of the reaction). This usually takes the form of catalyst-packed metal tubes inside a gas-fired furnace. It is important to control the proportion of gas to steam flow into this process. Too much hydrocarbon gas, and the result will be “coking” (solid hydrocarbon deposits) inside the heated tubes and on the surface of the catalyst beads, decreasing the efficiency of the process over time. Too much steam and the result is wasted energy as unreacted steam simply passes through the heater tubes, absorbing heat and carrying it away from the catalyst where it would otherwise do useful work.

One way to achieve the proper ratio of hydrocarbon gas to steam flow is to install a normal flow control loop on one of these two reactant feed lines, then use that process variable (flow) signal as a setpoint to a flow controller installed on the other reactant feed line. This way, the second controller will maintain a proper balance of flow to proportionately match the flow rate of the other reactant. An example P&ID is shown here, where the methane gas flow rate establishes the setpoint for steam flow control:



Note how the methane gas flow transmitter signal goes both to the methane flow controller

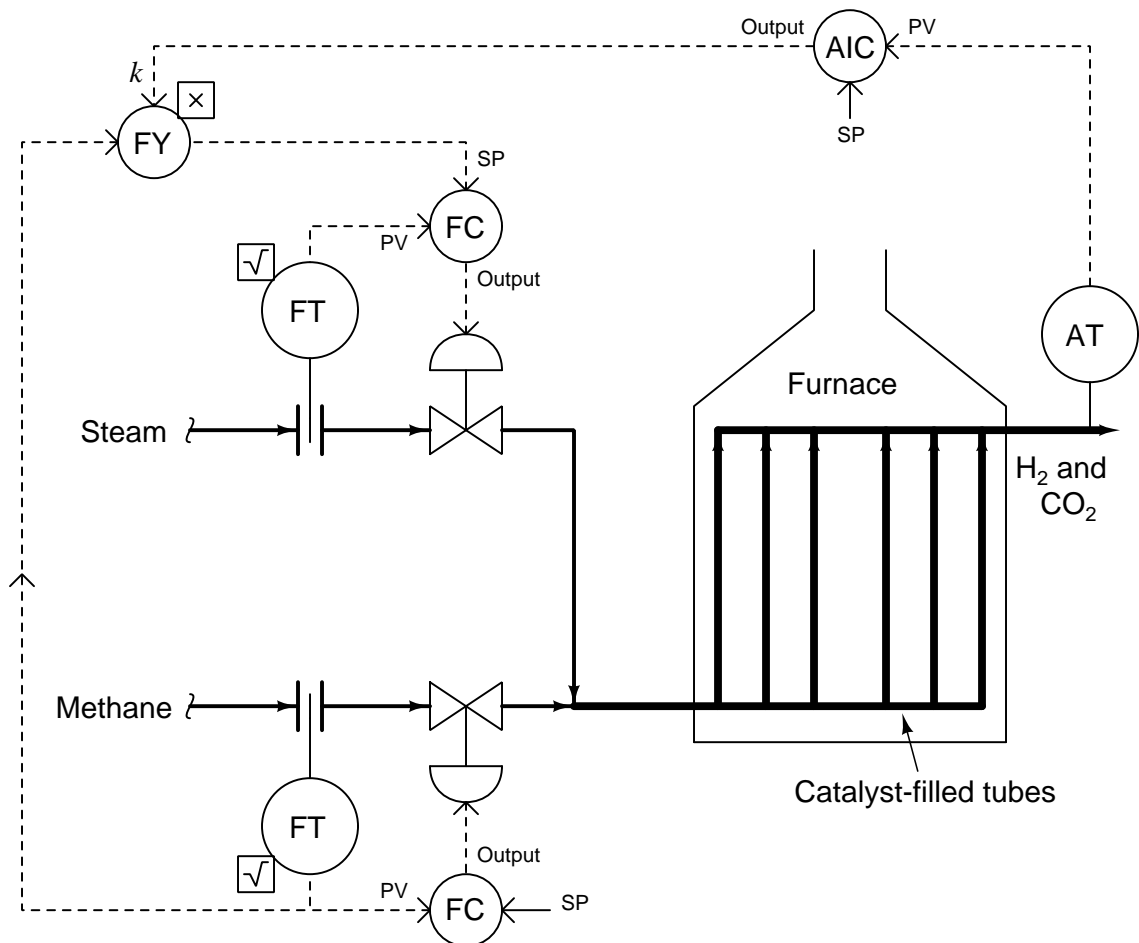
and to a *multiplying relay* that multiplies this signal by a constant value (k) before passing it on to the steam flow controller as a setpoint. This k value sets the *ratio* of steam flow to methane flow. Although this might appear to be a cascade control system at first glance, it is actually quite different. In a cascade system, the *output* of one controller becomes the setpoint for another. Here in a ratio control system, the *process variable* of one controller becomes the setpoint for another, such that two process variables remain in constant proportion (ratio) to one another.

If the two flow transmitters are compensated to measure mass flow, the ideal value of k should be set such that two molecules of steam vapor (H_2O) enter the reforming furnace for every one molecule of methane (CH_4). With a 2-to-1 molecular ratio of steam to methane (2 moles of steam per one mole of methane), this equates to a 9-to-4 mass flow ratio once the formula weights of steam and methane are calculated⁶. Thus, if the methane and gas flowmeters are calibrated for equal mass flow ranges, the ideal value for k should be $\frac{9}{4}$, or 2.25. Alternatively, the flow transmitter calibrations could be set in such a way that the ideal ratio is intrinsic to those transmitters' ranges (i.e. the methane flow transmitter has 2.25 times the mass flow range of the steam flow transmitter), with k set to an ideal value of 1. This way a 9:4 ratio of methane mass flow to steam mass flow will result in equal percentage output values from both flow transmitters. In practice, the value for k is set a bit higher than ideal, in order to ensure just a little excess steam to guard against coking inside the reaction heater tubes⁷.

⁶Steam has a formula weight of 18 amu per molecule, with two hydrogen atoms (1 amu each) and one oxygen atom (16 amu). Methane has a formula weight of 16 amu per molecule, with one carbon atom (12 amu) and four hydrogen atoms (1 amu each). If we wish to have a molecular ratio of 2:1, steam-to-methane, this makes a formula weight ratio of 36:16, or 9:4.

⁷It is quite common for industrial control systems to operate at ratios a little bit "skewed" from what is stoichiometrically ideal due to imperfect reaction efficiencies. Given the fact that no chemical reaction ever goes to 100% completion, a decision must be made as to which form of incompleteness is worse. In a steam-hydrocarbon reforming system, we must ask ourselves which is worse: excess (unreacted) steam at the outlet, or excess (unreacted) hydrocarbon at the outlet. Excess hydrocarbon content will "coke" the catalyst and heater tubes, which is very bad for the process over time. Excess steam merely results in a bit more operating energy loss, with no degradation to equipment life. The choice, then, is clear: it is better to operate this process "hydrocarbon-lean" (more steam than ideal) than "hydrocarbon-rich" (less steam than ideal).

We could add another layer of sophistication to this ratio control system by installing a gas analyzer at the outlet of the reaction furnace designed to measure the composition of the product stream. This analyzer's signal could be used to adjust the value of k so the ratio of steam to methane would automatically vary to ensure optimum production quality even if the feedstock composition (i.e. percentage concentration of methane in the hydrocarbon gas input) changes:




As we saw before, pure methane feed requires a 9-to-4 steam-to-methane mass flow ratio for the desired reaction to be stoichiometrically balanced. This mass ratio, however, is not balanced for hydrocarbons other than methane. Ethane (C_2H_6) processed in the same way requires a 12-to-5 steam-to-ethane mass flow ratio. Propane (C_3H_8) requires a 26-to-11 steam-to-propane mass flow ratio. If the hydrocarbon feed to the reforming furnace varies in composition, the steam flow ratio (k) must change accordingly for efficient reaction.

31.4 Relation control

A control strategy similar to ratio control is *relation* control. This is similar to ratio control in that a “wild” variable determines the setpoint for a captive variable, but with relation control the mathematical relationship between the wild and captive variables is one of addition (or subtraction) rather than multiplication (or division). In other words, a relation control system works to maintain a specific *difference* between wild and captive flow values, whereas a ratio control system works to maintain a specific *ratio* between wild and captive flow values.

An example of relation control appears here, where a temperature controller for a steam superheater on a boiler receives its setpoint from the biased output of a temperature transmitter sensing the temperature of saturated steam (that is, steam exactly at the boiling point of water) in the steam drum:



It is a basic principle of thermodynamics that the vapor emitted at the surface of a boiling liquid will be at the same temperature as that liquid. Furthermore, any heat lost from that vapor will cause at least some of that vapor to condense back into liquid. In order to ensure the vapor is “dry” (i.e. it may lose substantial heat energy without condensing), the vapor must be heated beyond the liquid’s boiling point at some later stage in the process.

Steam within the steam drum of a boiler is *saturated* steam: at the same temperature as the boiling water. Any heat lost from saturated steam causes at least some of it to immediately condense back into water. In order to ensure “dry” steam output from the boiler, the saturated steam taken from the steam drum must be further heated through a set of tubes called a *superheater*. The resulting “dry” steam is said to be *superheated*, and the difference between its temperature and the temperature of the boiling water (saturated steam) is called *superheat*.

This control system maintains a set amount of superheat by measuring the saturated steam’s temperature (within the steam drum), adding a “superheat setpoint” bias value to that signal, then passing the biased signal to the temperature indicating controller (TIC) where the superheated steam temperature is regulated by adding water⁸ to the superheated steam. With this system in place, the boiler operator may freely define how much superheat is desired, and the controller attempts to maintain the superheated steam at that much higher temperature than the saturated steam in the drum, over a wide range of saturated steam temperatures.


A ratio control system would not be appropriate here, since what we desire in this process is a controlled *offset* (rather than a controlled *ratio*) between two steam temperatures. The control strategy looks very much like a ratio control, except for the substitution of a summing function instead of a multiplying function.

⁸This mixing of superheated steam and cold water happens in a specially-designed device called a *desuperheater*. The basic concept is that the water will absorb heat from the superheated steam, turning that injected water completely into steam and also reducing the temperature of the superheated steam. The result is a greater volume of steam than before, at a reduced temperature. So long as some amount of superheat remains, the de-superheated steam will still be “dry” (above its condensing temperature). The desuperheater control merely adds the appropriate amount of water until it achieves the desired superheat value.

31.5 Feedforward control

“Feedforward” is a rather under-used control strategy capable of managing a great many types of process problems. It is based on the principle of *preemptive load counter-action*: that if all significant loads on a process variable are monitored, and their effects on that process variable are well-understood, a control system programmed to take appropriate action based on load changes will shield the process variable from any ill effect. That is to say, the feedforward control system uses data from load sensors to predict when an upset is about to occur, then *feeds that information forward to the final control element* to counteract the load change before it has an opportunity to affect the process variable. Feedback control systems are *reactive*, taking action after to changes in the process variable occur. Feedforward control systems are *proactive*, taking action before changes to the process variable can occur.

This photograph shows a kind of feedforward strategy employed by human operators running a *retort*: a steam-powered machine used to pressure-treat wooden beams at a milled lumber operation. The sign taped to this control panel reminds the operator to warn the maintenance department of an impending steam usage:




The story behind this sign is that a sudden demand in retort steam causes the entire facility’s steam supply pressure to sag if it happens at a time when the boiler is idling. Since the boiler’s pressure control system can only react to deviations in steam pressure from setpoint, the boiler pressure controller will not take any action to compensate for sudden demand until *after* it sees the steam pressure fall, at which point it may be too late to fully recover. If operators give the maintenance personnel advance notice of the steam demand, though, the boiler may be fired up for extra steam capacity and thus will be prepared for the extra demand when it comes. The upset avoided here is abnormally low steam header pressure, with the predictive load being the retort

operator's planned usage of steam. Crude as this solution might be, it illustrates the fundamental concept of feedforward control: information about a load change is "fed forward" to the final control element to preemptively stabilize the process variable.

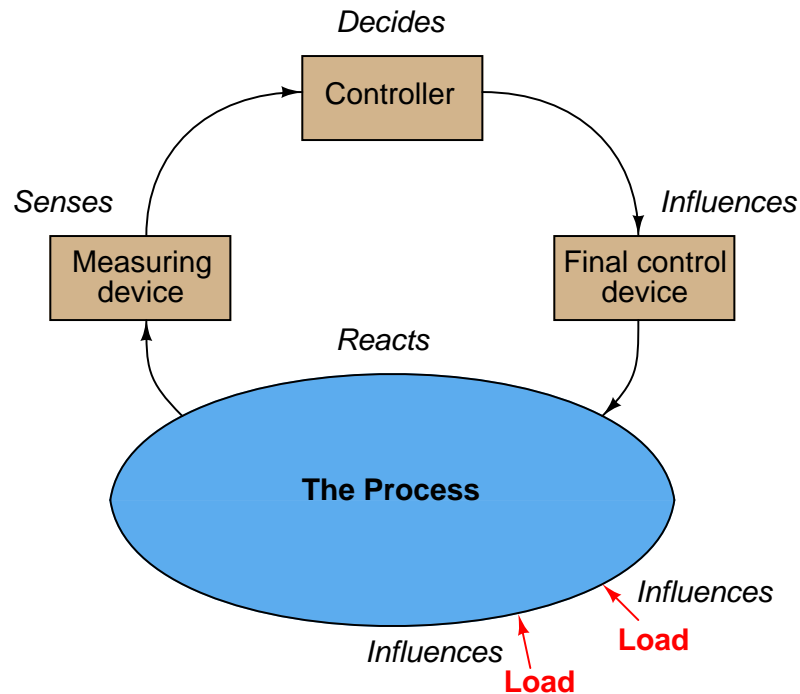
As the following section explains, perfect feedforward control action is nearly impossible to achieve. However, even imperfect feedforward action is often far better than none at all, and so this control strategy is quite valuable in process control applications challenged by frequent and/or large variations in load.

31.5.1 Load Compensation

Feedback control works on the principle of information from the outlet of a process being "fed back" to the input of that process for corrective action. A block diagram of feedback control looks like a loop:



The reason any control system is necessary at all⁹ to maintain a process variable at some stable value is the existence of something called a *load*. A “load” is a variable influencing a process that is not itself under direct control, and may be represented in the block diagram as an arrow entering the process, but not within the control loop:



For example, consider the problem of controlling the speed of an automobile. In this scenario, vehicle speed is the process variable being measured and controlled, while the final control device is the accelerator pedal controlling engine power output. If it were not for the existence of hills and valleys, head-winds and tail-winds, air temperature changes, road surface variations, and a host of other “load” variables affecting car speed, maintaining a constant speed would be as simple as holding the accelerator pedal at a constant position.


However, the presence of these “load” variables makes necessitates a human driver (or a *cruise control* system) continually adjusting engine power to maintain constant speed. Using the car’s measured speed as feedback, the driver (or cruise control) adjusts the accelerator pedal position as necessary based on whether or not the car’s speed matches the desired “setpoint” value.

An inherent weakness of any feedback control system is that it can never be *proactive*. The best any feedback control system can ever do is *react* to detected disturbances in the process variable. This makes deviations from setpoint inevitable, even if only for short periods of time. In the context of our automobile cruise control system, this means the car can never maintain a *perfectly* constant

⁹This statement is true only for self-regulating processes. Integrating and “runaway” processes require control systems to achieve stability even in the complete absence of any loads. However, since self-regulation typifies the vast majority of industrial processes, we may conclude that the fundamental purpose of most control systems is to counteract the effects of loads.

speed in the face of loads because the control system does not have the ability to anticipate loads (e.g. hills, wind gusts, changes in air temperature, changes in road surface, etc.). At best, all the feedback cruise control system can do is react to changes in speed it senses *after* some load has disturbed it.

Feedforward control addresses this weakness by taking a fundamentally different approach, basing final control decisions on the states of load variables rather than the process variable. In other words, a feedforward control system monitors the factor(s) influencing a process and decides how to compensate *ahead of time* before the process variable deviates from setpoint. If all loads are accurately measured, and the control algorithm realistic enough to predict process response for these known load values, the process variable (ideally) need not be measured at all:



As was the case with cascade control, feedforward control also has an analogue in workplace management. If you consider a supervisor to be the “controller” of a work group (issuing orders to his or her subordinates to accomplish important tasks), a feedforward system would be when someone informs the supervisor of an important change that will soon impact the work group. By having this information “fed forward” to the supervisor, the supervisor may then take *preemptive* measures to better manage this change before its effects are fully felt. If this predictive information is accurate, and the supervisor’s response appropriate, any negative impacts of the change will be minimized to the point where no reactive steps will be needed. Stated differently, good feedforward control action translates what would otherwise be a crisis into an insignificant event.

Returning to the cruise control application, a purely feedforward automobile cruise control system would be interfaced with topographical maps, real-time weather monitors, and road surface sensors to decide how much engine power was necessary at any given time to attain the desired speed¹⁰.

¹⁰The load variables I keep mentioning that influence a car’s speed constitute an incomplete list at best. Many


Assuming all relevant load variables are accounted for, the cruise control would be able to maintain constant speed regardless of conditions, and without the need to even monitor the car's speed.

This is the promise of feedforward control: a method of controlling a process variable so perfect in its predictive power that it eliminates the need to even measure that process variable. If you are skeptical of this feedforward principle and its ability to control a process variable without even measuring it, this is a good thing – you are thinking critically! In practice, it is nearly impossible to accurately account for *all* loads influencing a process and to both anticipate and counter-act their combined effects, and so *pure* feedforward control systems are rare¹¹. Instead, the feedforward principle finds use as a supplement to normal feedback control. To understand feedforward control better, however, we will consider its pure application before exploring how it may be combined with feedback control.

other variables come into play, such as fuel quality, engine tuning, and tire pressure, just to name a few. In order for a purely feedforward (i.e. no feedback monitoring of the process variable) control system to work, *every single load variable* must be accurately monitored and factored into the system's output signal. This is impractical or impossible for a great many applications, which is why we usually find feedforward control used in conjunction with feedback control, rather than feedforward control used alone.

¹¹In fact, the only pure feedforward control strategies I have ever seen have been in cases where the process variable was nearly impossible to measure and could only be inferred from other variables.

First, let us consider a liquid level control system on an open tank, where three different fluid ingredients (shown in the following P&ID simply as A, B, and C) are mixed to produce a final product. A level transmitter (LT) measures liquid level, while a level controller (LC) compares this level to a setpoint value, and outputs a signal calling for a certain amount of discharge flow. A cascaded (slave) flow controller (FC) senses outgoing flow via a flow transmitter (FT) and works to maintain whatever rate of flow is “asked” for by the level controller:




The level control system acts to keep liquid level constant in the vessel, ensuring adequate mixing of the three ingredients¹². Being a feedback level control system, it adjusts the discharge flow rate in response to measured changes in liquid level. Like all feedback control systems, this one is *reactive* in nature: it can only take corrective action *after* a deviation between process variable (level) and setpoint is detected. As a result, temporary deviations from setpoint are guaranteed to occur with

¹²If the liquid level drops too low, there will be insufficient *retention time* in the vessel for the fluids to mix before they exit the product line at the bottom.

this control system every time the combined flow rate of the three ingredients increases or decreases.

Let us now change the control system strategy from feedback to feedforward. It is clear what the loads are in this process: the three ingredient flows entering the vessel. If we measure and sum these three flow rates¹³, then use the total incoming flow signal as a setpoint for the discharge flow controller, the outlet flow should (ideally) match the inlet flow, resulting in a constant liquid level. Being a purely feedforward control system, there is no level transmitter (LT) any more, just flow transmitters measuring the three loads:



If all flow transmitter calibrations are perfect, the summing of flow rates flawless, and the flow controller's tuning robust, this level control system should control liquid level in the vessel by proactive effort ("thinking ahead") rather than reactive effort ("after the fact"). Any change

¹³The device or computer function performing the summation is shown in the P&ID as a bubble with "FY" as the label. The letter "F" denotes *Flow*, while the letter "Y" denotes a signal relay or transducer.

in the flow rate of ingredients A, B, and/or C is quickly matched by an equal adjustment to the discharge flow rate. So long as total volumetric flow out of the vessel is held equal to total volumetric flow into the vessel, the liquid level inside the vessel *cannot* change¹⁴.

If this feedforward strategy reminds you of ratio control, you are thinking correctly: the ingredient flow sum signal is the *wild variable*, and the discharge flow signal is the *captive variable*. The flow controller simply maintains the discharge flow rate at a 1:1 ratio with the (total) ingredient flow rate. In fact, pure feedforward control is a variation of 1:1 ratio control, except that the real process variable (tank level) is neither the wild (total incoming flow) nor the captive variable (discharge flow) in the process.


An interesting property of feedforward and ratio control systems alike is that they cannot generate oscillations as is the case with an over-tuned (excessive gain) feedback system. Since a feedforward system does not monitor the effects of its actions, it cannot react to something it did to the process, which is the root cause of feedback oscillation. While it is entirely possible for a feedforward control system to be configured with too much gain, the effect of this will be *overcompensation* for a load change rather than oscillation. In the case of the mixing tank feedforward level control process, improper instrument scaling and/or offsets will merely cause the discharge and inlet flows to mismatch, resulting in a liquid level that either continues to increase or decrease over time (“integrate”). However, no amount of mis-adjustment can cause this feedforward system to produce *oscillations* in the liquid level.

In reality, this pure feedforward control system is impractical even if all instrument calibrations and control calculations are perfect. There are still loads unaccounted for: evaporation of liquid from the vessel, for example, or the occasional pipe fitting leak. Furthermore, since the control system has no “knowledge” of the actual liquid level, it cannot make adjustments to that level. If an operator, for instance, desired to decrease the liquid level in order to reduce the residence time (also known as “retention time”)¹⁵, he or she would have to manually drain liquid out of the vessel, or temporarily place the discharge flow controller in “manual” mode and increase the flow there (then place back into “cascade” mode where it follows the remote setpoint signal again). The advantage of proactive control and minimum deviation from setpoint over time comes at a fairly high price of impracticality and inconvenience.

¹⁴Incidentally, this is a good example of an *integrating* mass-balance process, where the rate of process variable change over time is proportional to the imbalance of flow rates in and out of the process. Stated another way, total accumulated (or lost) mass in a mass-balance system such as this is the time-integral of the difference between incoming and outgoing mass flow rates: $\Delta m = \int_0^T (W_{in} - W_{out}) dt$.


¹⁵*Residence time* or *Retention time* is the average amount of time each liquid molecule spends inside the vessel. It is an important variable in chemical reaction processes, where adequate time must be given to the reactant molecules in order to ensure a complete reaction. It is also important for non-reactive mixing processes such as paint and food manufacturing, to ensure the ingredients are thoroughly mixed together and not stratified. For any given flow rate through a vessel, the residence time is directly proportional to the volume of liquid contained in that vessel: double the captive volume, and you double the residence time. For any given captive volume, the residence time is inversely proportional to the flow rate through the vessel: double the flow rate through the vessel, and you halve the residence time. In some mixing systems where residence time is critical to the thorough mixing of liquids, vessel level control may be coupled to measured flow rate, such that an increase in flow rate results in an increased level setpoint, thus maintaining a constant residence time despite changes in production rate.

For these reasons, feedforward control is most often found in conjunction with feedback control. To show how this would work in the liquid level control system, we will incorporate a level transmitter and level controller back into the system, the output of that level controller being summed with the feedforward flow signal (by the LY summing relay) before going to the cascaded setpoint input of the discharge flow controller:



This hybrid control strategy is sometimes called *feedforward with trim*. In this context, “trim” refers to the level controller’s (LC) output signal contributing to the discharge flow setpoint, helping to compensate for any unaccounted loads (evaporation, leaks) and provide for level setpoint changes. This “trim” signal should do very little of the control work in this system, the bulk of the liquid level stability coming from the feedforward signals provided by the incoming flow transmitters.

A very similar control strategy commonly used on large steam boilers for the precise control of steam drum water level goes by the name of *three-element feedwater control*. The following illustration shows an example of this control strategy implemented with pneumatic (3-15 PSI signal) instruments:



Such a control system is called “three-element” because it makes use of three process measurements:

- Feedwater flow rate
- Steam drum water level
- Steam flow rate


Feedwater flow is controlled by a dedicated flow controller (FIC), receiving a remote setpoint signal from a summing relay (LY). The summer receives two inputs: a steam flow signal and the

output signal (trim) from the level controller (LIC). The feedforward portion of this system (steam flow feeding forward to water flow) is intended to match the mass flow rates of water into the boiler with steam flow out of the boiler. If steam demand suddenly increases, this feedforward portion of the system immediately calls for a matching increase in water flow into the boiler, since every molecule of steam exiting the boiler must come from one molecule of water entering the boiler. The level controller and transmitter act as a feedback control loop, supplementing the feedforward signal to the cascaded water flow controller to make up for (“trim”) any shortcomings of the feedforward loop.

A three-element boiler feedwater control system is a good example of a feedforward strategy designed to ensure *mass balance*, defined as a state of equality between all incoming mass flow rates and all outgoing mass flow rates. The steam flow transmitter measures outgoing mass flow, its signal being used to adjust incoming water mass rate. Since mass cannot be created or destroyed (the Law of Mass Conservation), every unit of steam mass leaving the boiler must be accounted for as an equivalent unit of water mass entering the boiler. If the control system perfectly balances these mass flow rates, water level inside the boiler *cannot* change.

In processes where the process variable is affected by energy flow rates rather than mass, the balance maintained by a feedforward control system will be *energy balance* rather than mass balance. Like mass, energy cannot be created or destroyed (the Law of Energy Conservation), but must be accounted for. A feedforward control system monitoring all incoming energy flows into a process and adjusting the outgoing energy flow rate (or vice-versa) will ensure no energy is depleted from or accumulated within the process, thus ensuring the stability of the processes’ internal energy state.

An example of energy-balance feedforward control appears in this heat exchanger temperature control system:



The two transmitters on the incoming (cold oil) line measure oil temperature and oil flow rate, respectively. The first “summing” function subtracts the incoming oil temperature from the setpoint (desired) temperature, and then the difference of these two temperatures is then multiplied by the flow rate signal to produce a signal representing the *energy demand*¹⁶ of the incoming oil (i.e. how much energy will be required to elevate the oil flow’s temperature to setpoint). The “energy demand” signal is summed with the temperature controller’s output signal to set the steam valve position (adding energy to the process).


There do exist other loads in this process, such as ambient air temperature and chemical composition of the oil, but these variables are generally less influential on discharge temperature than feed temperature and flow rate. This illustrates a practical facet of feedforward control: although there may be a great many loads affecting our process variable, we must generally limit our application of feedforward to only the most dominant loads in order to limit control system cost. Simply put, we usually cannot justify the expense and complexity of a feedforward control system compensating for *every single load* in a system.

¹⁶Energy demand is an example of what is called an *inferred variable*: a physical quantity that we cannot measure directly but instead calculate from measurements made of other variables.

31.5.2 Proportioning feedforward action


Feedforward control works by directly modulating the manipulated variable in a control system according to changes sensed in the load(s). In order for feedforward to function optimally, it must adjust the manipulated variable in a manner that is proportionate to the need: no more, and no less. At this juncture it is appropriate to ask the question, “how do we know the amount of feedforward action that will be adequate for a process, and how do we adjust it if it is too much or too little?”

In processes where the feedforward control strategy attempts to achieve direct mass- or energy-balance, the question of adequate feedforward action is answered in the mathematics of measuring the incoming and outgoing flows. Consider the following mass-balance level control system where the combined sum of three inlet flows is routed to the setpoint of the exit flow control loop. In this diagram, the portions of the control strategy implemented as function blocks (algorithms in software) appear inside a yellow-colored bounded area, while all real physical instruments appear outside the yellow area:



If all flowmeters are calibrated in pounds per minute, then the feedforward signal will likewise be scaled in pounds per minute, and so will the setpoint be for flow control loop. In a digital control system, it is quite customary to scale each and every analog input signal with some real “engineering unit” of measurement, so that the signal will be treated as a physical quantity throughout as opposed to being treated as some anonymous percentage value. Not only is this consistent scaling a standard feature in digital control systems, but it also helps the implementation of this feedforward control strategy, because we desire the out-going mass flow rate to precisely match the (total) in-coming flow rate. So long as all flowmeters and their associated scaling factors are accurate, the feedforward control’s action *must* be exactly right: an increase of +5 pounds per minute in incoming flow rate will prompt an immediate increase of +5 pounds per minute in outgoing flow rate, simply by virtue of all these measured flows having been scaled in the same unit of measurement.

The situation is not as simple in systems where the feedforward control is not precisely balancing mass-flow or energy rates. By contrast, let us examine the following pH neutralization system equipped with feedforward control action. Here, the incoming liquid is alkaline (pH greater than 7), and the control system's job is to mix just enough acid reagent to "neutralize" the solution (bring the pH value down to 7):




Controlling the pH (acidity/alkalinity) of a liquid solution is challenging for many reasons, not the least of which being the need to have adequate mixing time for the reagent to react with the influent. This mixing time translates to *dead time* in the feedback control system. If the influent's pH suddenly changes for any reason, the feedback control system will be slow to alter the reagent flow rate due to this dead time, causing long-lasting deviations from setpoint. The goal of the feedforward signal (from the influent pH transmitter to the summer) is to preemptively adjust reagent flow rate according to how alkaline the incoming flow is, countering any sudden changes in influent pH so the feedback control system doesn't have to take (delayed) action.

Once again, it is appropriate to ask the question, "how do we know the amount of feedforward action that will be adequate for this process, and how do we adjust it if it is too much or too little?" It would be blind luck if the system happened to work perfectly as shown, with the influent pH transmitter's signal going straight to the summing function to be added to the pH controller's

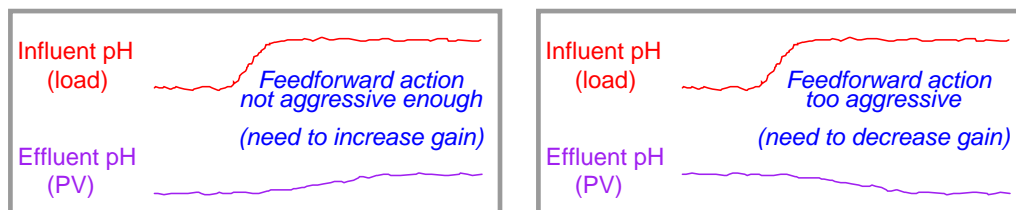
output signal. Certainly, an increase in influent pH would cause more acid to be added to the mix thanks to feedforward action, but it would likely add either too much or too little acid than it should. The scale of the influent pH transmitter does not match the scale of the signal sent to the control valve, and so we do not have a neat “pound-for-pound” balance of mass flow as we did in the case of the level control system.

A neat solution to this problem is to add another function block to the feedforward portion of the control system. This block takes the influent pH transmitter signal and skews it using multiplication and addition, using the familiar linear equation $y = mx + b$ (where y is the output signal of the function and x is the input signal; m and b being constants). This function block is typically called a *gain and bias* block:



The gain adjustment (m) in this function block serves to amplify or attenuate the feedforward signal's magnitude, while the bias adjustment (b) offsets it.

Determining practical values for these “feedforward tuning constants” is relatively easy. First, disable feedback control¹⁷ with the output value at or near 50%. Next, introduce load changes to the process while watching the process variable’s value after sufficient time has elapsed to see the effects of those load changes. Increase or decrease the gain value until step-changes in load cease to yield significant changes in the process variable. The following trends show what too much and too little feedforward gain would look like in this pH control system:



If the gain is properly set in the gain/bias function block, these load changes should have minimal effect on the process variable. Step-changes in influent pH should have little effect on the process variable after sufficient time has passed for the load change to have fully propagated through the process.

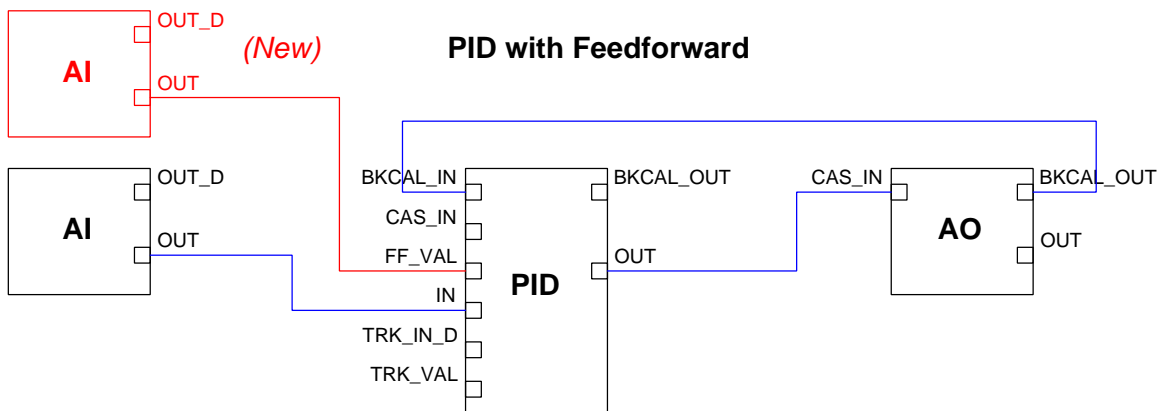
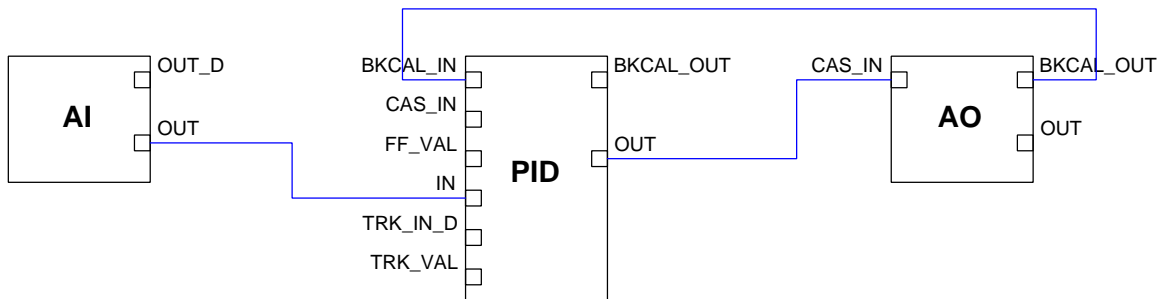
Once a good gain value has been found, change the bias value until the process variable approaches the normal setpoint value¹⁸.

¹⁷Most control systems’ feedforward function blocks are designed in such a way that both the feedback and the feedforward signal paths are disabled when the controller is placed into manual mode, in order to give the human operator 100% control over the final element (valve) in that mode. For the purpose of “tuning” the feedforward gain/bias function block, one must disable the feedback control *only* so feedforward action is still able to respond to load changes. If simply switching the feedback controller to manual mode is not an option (which it usually is not), one may achieve the equivalent result by setting the gain value of the feedback controller to zero and ensuring the PID equation is not the “parallel” type. If the PID equation is parallel, you will need to set all three terms (P, I, and D) at their minimum settings.

¹⁸This is why it was recommended to leave the feedback controller’s output at or near 50%. The goal is to have the feedforward action adjusted such that the feedback controller’s output is “neutral,” and has room to swing either direction if needed to provide necessary trim to the process.


Some control systems provide convenient methods of incorporating feedforward action. The standard FOUNDATION Fieldbus PID function block, for example, has its own dedicated signal input for feedforward, with feedforward gain as an existing parameter. The following illustration shows Fieldbus function blocks used to implement normal feedback (PID) control, and also PID with feedforward control:

Standard PID control algorithm with FF function blocks



As you can see, all that is required to augment a FOUNDATION Fieldbus control system with feedforward control action is the addition of one more analog input (AI) function block for the load transmitter, and one connecting line between that block and the PID block's "FF_Val" input.

Other control systems require a bit more programming to implement feedforward. For example, consider the standard “Factory Configured Option 101” function block program for basic PID feedback control in a Siemens model 353 panel-mounted loop controller, contrasted against the version necessary for implementing feedforward action:



Note the necessary addition of *two* “math” function blocks as well as the extra analog input block for receiving the feedforward (load variable) transmitter’s signal: MTH1 needed to add the load transmitter’s feedforward signal to the controller’s output signal, and MTH2 needed to ensure

output tracking still works properly between the PID and A/M function blocks when the human operator switches between automatic and manual modes. The amount of feedforward action is specified by the “gain” parameter for input B of *both* math function blocks, which means any adjustment to feedforward action must be manually entered in two different places!

31.6 Feedforward with dynamic compensation

As we have seen, feedforward control is a way to improve the stability of a feedback control system in the face of changing loads. Rather than rely on feedback to make corrective changes to a process only *after* some load change has driven the process variable away from setpoint, feedforward systems monitor the relevant load(s) and use that information to preemptively make stabilizing changes to the final control element such that the process variable will not be affected. In this way, the feedback loop’s role is to merely “trim” the process for factors lying outside the realm of the feedforward system.


At least, this is how feedforward control is *supposed* to work. One way feedforward controls commonly fail to live up to their promise is if the effects of load changes and of manipulated variable changes possess different time lags in their respective effects on the process variable. This is a problem in feedforward control systems because it means the corrective action called for in response to a change in load will not affect the process variable at the same time, or in the same way over time, as the load will. In order to correct this problem, we must intelligently insert time lags (or advancing time-based functions called *leads*) into the control system to equalize the time lags of load and feedforward correction. This is called *dynamic compensation*.

The following subsections will explore illustrative examples to make both the problem and the solution(s) clear.

A common area of confusion among students first approaching this topic is deciding where to place the dynamic compensation function in a feedforward control system. The answer to this question is surprisingly simple, although it may seem elusive at first glance. The key is found in the following principle: *the only time-dynamic we have the ability to alter with our control system is the dynamic of the final control element*. We cannot alter the time-dynamic of the load’s effect on the process variable, as that is strictly a function of process physics. Therefore, when we test a process employing feedforward control with an eye toward incorporating dynamic compensation, we must measure the time lag of the load’s effect on PV and also the time lag of our final control element’s effect on PV, then compare those two time lags. If the final control element’s time lag is shorter (quicker) than the load’s, then we must add a delay or lag to the feedforward signal so that the final control element’s preemptive action does not occur too soon. If the final control element’s time lag is longer (slower) than the load’s, then we must either find a way to alter the process itself to decrease the load’s time lag, or add a “lead” function to the feedforward signal in order to advance the final control element’s response and thereby ensure the preemptive action does not occur too late. *Remember, all we can do with dynamic compensation is alter how the final control element responds (i.e. how slowly or quickly the preemptive action of feedforward occurs). The load’s effect on the process variable is fixed by the physics of the process and therefore lies beyond our direct control.*

31.6.1 Dead time compensation

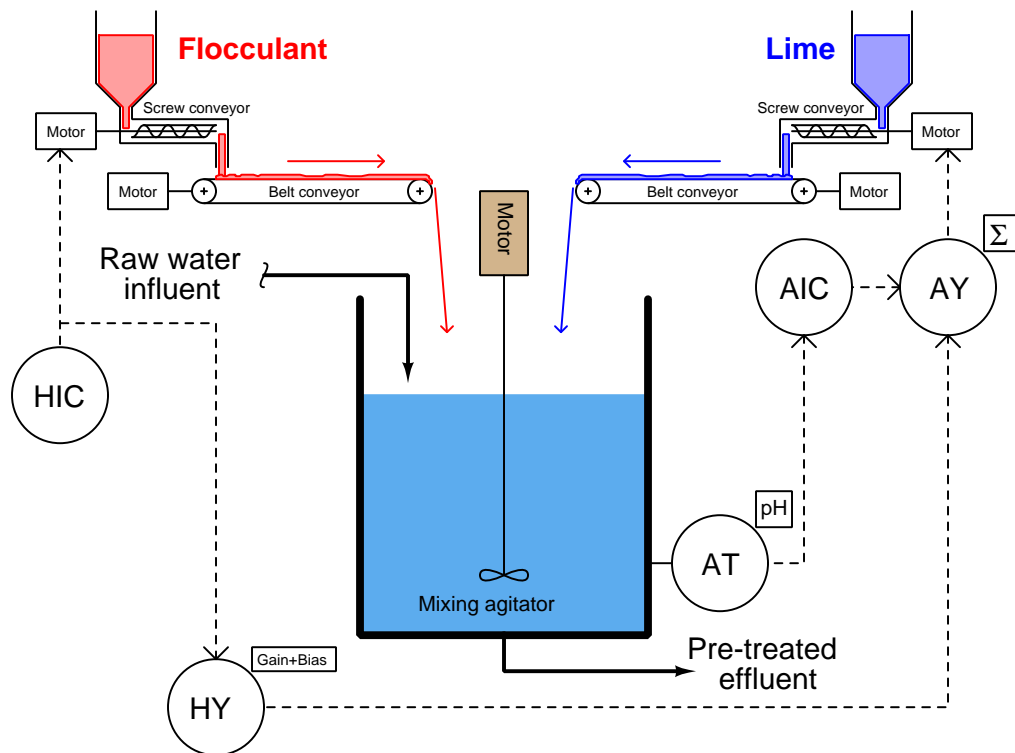
Examine the following control system P&ID showing the addition of *flocculant* (a chemical compound used in water treatment to help suspended solids clump together for easier removal by filtering and/or gravity clarification) and *lime* for pH balance. Flocculant is necessary to expedite the removal of impurities from the water, but some flocculation compounds have the unfortunate effect of decreasing the pH value of the water (turning it more acidic). If the water's pH value is too low, the flocculant ironically loses its ability to function. Thus, lime (an alkaline substance – high pH value) must be added to the water to counter-act the flocculant's effect on pH to ensure efficient flocculation. Both substances are powders in this water pre-treatment system, metered by variable-speed screw conveyors and carried to the mixing tank by belt-style conveyors:



The control system shown in this P&ID consists of a pH analyzer (AT) transmitting a signal to a pH indicating controller (AIC), adjusting the speed of the lime screw conveyor. The flocculant screw conveyor speed is manually set by a *hand indicating controller* (HIC) – sometimes known as a *manual loading station* – adjusted when necessary by experienced water treatment operators who periodically monitor the effectiveness of flocculation in the system.

This simple feedback control system will work fine in steady-state conditions, but if the operator suddenly changes flocculant flow rate into the mixing vessel, there will be a temporary deviation of pH from setpoint before the pH controller is able to find the correct lime flow rate into the vessel to compensate for the change in flocculant flow. In other words, flocculant feed rate into the mixing tank is a *load* for which the pH control loop must compensate.

Dynamic response could be greatly improved with the addition of feedforward control to this system:




Here, the hand controller's signal gets added to the pH controller's output signal to directly influence lime feed rate in addition to acting as a control signal to the flocculant screw conveyor motor drive. If an operator changes the flocculant feed rate, the lime feed rate will immediately adjust to compensate, *before* any change in pH value takes place in the water. Ideally, the pH controller need only make minor "trim" adjustments to lime feed rate, while the feedforward signal does most of the work in maintaining a steady pH value. The proper proportioning and offset between flocculant and lime feed rates is established in the gain/bias function, which is "tuned"¹⁹ to ensure the feedforward signal does not over- or under-react, calling for too much (or too little) lime to compensate.

Even if all components in the feedforward system have been calibrated and configured properly, however, a potential problem still lurks in this system which can cause the pH value to temporarily deviate from setpoint following flocculant feed rate changes. This problem is the *transport delay* – otherwise known as *dead time* – inherent to the two belt conveyors transporting both flocculant and lime powder from their respective screw conveyors to the mixing vessel. If the rotational speed of a

¹⁹Tuning this gain/bias block is done with the pH controller in manual mode with its output at 50%. The gain value is adjusted such that step-changes in flocculant feed rate have little long-term effect on pH. The bias value is adjusted until the pH approaches setpoint (even with the pH controller in manual mode).

screw conveyor changes, the flow rate of powder exiting that screw conveyor will immediately and proportionately change. However, the belt conveyor imposes a time delay before the new powder feed rate enters the mixing vessel. In other words, the water in the vessel will not “see” the effects of a change in flocculant or lime feed rate until after the *belt conveyor’s* time delay has elapsed. This is not a problem if the dead times of both belt conveyors are exactly equal, since this means any compensatory change in lime feed rate initiated by the feedforward system will reach the water at exactly the same time the new flocculant rate reaches the water. So long as flocculant and lime feed rates are precisely balanced with one another at the point in time they reach the mixing vessel, pH should remain stable. But what if their arrival times are not coordinated – what will happen to pH then?


Let us engage in a “thought experiment” to explore the consequences of the flocculant conveyor belt moving much slower than, and/or being much longer than, the lime conveyor belt. Suppose the flocculant belt imposed a dead time of 60 seconds on flocculant powder making it to the vessel, while the lime belt only delayed lime powder transit by 5 seconds from screw conveyor to mixing tank. This would mean changes in flocculant flow (set by the hand controller) would compensate with changes in lime flow *55 seconds too soon*. Now imagine the human operator making a sudden increase to the flocculant powder feed rate. The lime feed rate would immediately increase thanks to the efforts of the feedforward system. However, since the increased flow rate of lime powder will reach the mixing vessel 55 seconds before the increased flow rate of flocculant powder, the effect will be a temporary increase in pH value beginning about 5 seconds after the operator’s change, and then a settling of pH value back to setpoint²⁰, as shown in this timing diagram:



The obvious solution to this problem is to mechanically alter the belt conveyor systems for equal transport times of flocculant and lime powders. If this is impractical, we may achieve a similar result by incorporating another signal relay (or digital function block) inserting dead time into the feedforward control system. In other words, we can modify the control system in such a way to emulate what would be impractical to modify in the process itself.

²⁰This “thought experiment” assumes no compensating action on the part of the feedback pH controller for the sake of simplicity. However, even if we include the pH controller’s efforts, the problem does not go away. As pH rises due to the premature addition of extra lime, the controller will try to reduce the lime feed rate. This will initially reduce the degree to which pH deviates from setpoint, but then the reverse problem will occur when the increased flocculant enters the vessel 55 seconds later. Now, the pH will drop below setpoint, and the feedback controller will have to ramp up lime addition (to the amount it was before the additional lime reached the vessel) to achieve setpoint.


This new function will add a dead time of 55 seconds to the feedforward signal before it enters the summer, thus delaying the lime feed rate's response to feedforward effect by just the right amount of time such that any lime feed rate changes called for by feedforward action will arrive at the vessel *simultaneously* with the changed flocculant feed rate:



Adding time-based functions to a control system in order to equalize inherently unequal time delays in the physical process is called *dynamic compensation*. It is important to note that dynamic compensation cannot make physically unequal time lags equal – all it does is modify the feedforward system so that signal's effect arrives at the right time to compensate the load. In simple terms, we can only use dynamic compensation to modify the FCE's behavior, not the process behavior. Here, there is absolutely nothing the feedforward system can do to speed up the slower flocculant belt, so instead we chose to slow down the feedforward manipulation of lime flow to make it match the flocculant flow.

Note how the feedback pH controller's loop was purposely spared the effects of the added dead time function, by placing the function outside of that controller's feedback loop. This is important, as dead time in any form is the bane of feedback control. The more dead time within a feedback loop, the easier that loop will tend to oscillate. By strategically placing the dead time function before the summing relay rather than after (between the summer and the lime screw conveyor motor drive), the feedback control system still achieves minimum response time while only the feedforward signal gets delayed.

Let us now consider the same flocculant and lime powder control system, this time with transport delays reversed between the two belt conveyors. If the flocculant conveyor belt is now the fast one (5 seconds dead time) and the lime belt slow (60 seconds), the effects of flocculant feed rate changes will be reversed. An increase in flocculant powder feed rate to the vessel will result in a drop in pH beginning 5 seconds after the HIC setting change, followed by a rise in pH value after the additional lime feed rate finally reaches the vessel:



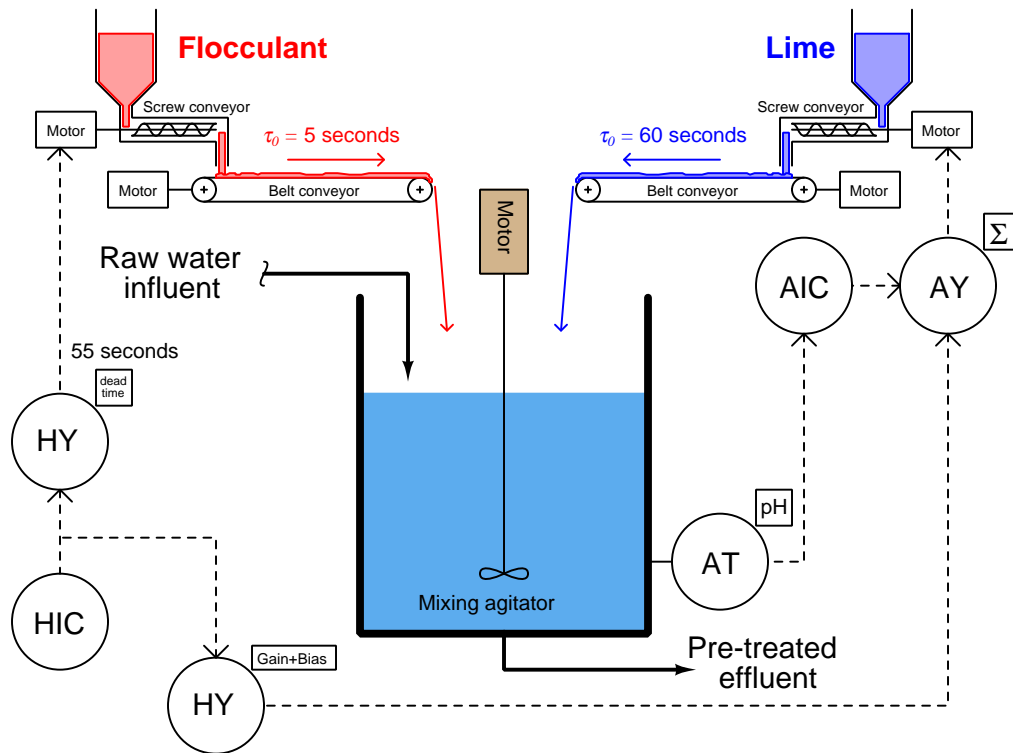
What is happening here is that the feedforward signal's effect is coming too late. In a perfect world our lime flow rate into the mixing vessel should change at the same time that the flocculant flow rate changes. Instead, our lime flow rate's necessary increase happens 55 seconds too late to prevent the pH from deviating.

It would be possible to compensate for the difference in conveyor belt transport times using a special relay in the same location of the feedforward signal path as before, if only there was such a thing as a relay that could *predict the future exactly 55 seconds in advance!*²¹. Since no such device exists (or ever will exist), we must apply dynamic compensation elsewhere in the feedforward control system.

If a time delay is the only type of compensation function at our disposal, then the only thing we can delay in this system to make the two dead times equal is the flocculation feed rate. Thus, we should place a 55-second dead time relay in the signal path between the hand indicating controller (HIC) and the flocculant screw conveyor motor drive.

²¹Let me know if you are ever able to invent such a thing. I'll even pay your transportation costs to Stockholm, Sweden so you can collect your Nobel prize. Of course, I will demand to see the prize before buying tickets for your travel, but with your time-travel device that should not be a problem for you.

This diagram shows the proper placement of the dead time function:




With this dead time relay in place, any change in flocculation feed rate initiated by a human operator will immediately adjust the feed rate of lime powder, but delay an adjustment to flocculant powder feed rate by 55 seconds, so the two powders' feed rate changes arrive at the mixing vessel simultaneously.

Bear in mind that this solution will really only work in a system like this where the major load happens to be controlled by the human operator. In most processes the load variable is not under anyone's control, making it difficult if not impossible to purposely delay its action. If the load has a shorter dead time than the compensation, there is usually little we can do within the control strategy to equalize those effects for better dynamic stability. Typically, the best solution is to *physically alter the process* (e.g. slow down the flocculant belt conveyor's speed) so that the load and compensation dead times are closer to being equal.

31.6.2 Lag time compensation

Process time delays characterized by pure transport delay (dead time) are less common in industry than other forms of time delays, most notably *lag times*²². A simple “lag” time is the characteristic exhibited by a low-pass RC filter circuit, where a step-change in input voltage results in an output voltage asymptotically rising to the new voltage value over time:




The *time constant* (τ) of such a system – be it an RC circuit or some other physical process – is the time required for the output to move 63.2% of the way to its final value ($1 - e^{-1}$). For an RC circuit such as the one shown, $\tau = RC$ (assuming $R_{load} \gg R$ so the load resistance will have negligible effect on timing).

Lag times differ fundamentally from dead times. With a dead time, the effect is simply time-delayed by a finite amount from the cause, like an echo. With a lag time, the effect begins at the exact same time as the cause, but does not follow the same rapid change over time as the cause. Like dead times in a feedforward system, it is quite possible (and in fact usually the case) for loads and final control variables to have differing lag times regarding their respective effects on the process variable. This presents another form of the same problem we saw in the two-conveyor water pre-treatment system, where an attempt at feedforward control was not completely successful because the corrective feedforward action did not occur with the same amount of time delay as the load.

²²For a more detailed discussion of lag times and their meaning, see section 30.1.5 beginning on page 2429.


To illustrate, we will analyze a heat exchanger used to pre-heat fuel oil before being sent to a combustion furnace. Hot steam is the heating fluid used to pre-heat the oil in the heat exchanger. As steam gives up its thermal energy to the oil through the walls of the heat exchanger tubes, it undergoes a phase change to liquid form (water), where it exits the shell of the exchanger as “condensate” ready to be re-boiled back into steam.

A simple feedback control system regulates steam flow to the heat exchanger, maintaining the discharge temperature of the oil at a constant setpoint value:



Once again, it should come as no surprise to us that the outlet temperature will suffer temporary deviations from setpoint if load conditions happen to change. The feedback control system may be able to *eventually* bring the exiting oil’s temperature back to setpoint, but it cannot begin corrective action until *after* a load has driven the oil temperature off setpoint. What we need for improved control is *feedforward* action in addition to feedback action. This way, the control system can take corrective action in response to load changes *before* the process variable gets affected.

Suppose we know that the dominant load in this system is oil flow rate²³, caused by changes in demand at the combustion furnace where this oil is being used as fuel. Adapting this control system to include feedforward is as simple as installing an oil flow transmitter, a gain/bias function, and a summing function block:



With feedforward control action in place, the steam flow rate will immediately change with oil flow rate, preemptively compensating for the increased or decreased heat demand of the oil. In other words, the feedforward system attempts to maintain *energy balance* in the process, with the goal of stabilizing the outlet temperature:

There is a problem of time delay in this system, however: a change in oil flow rate has a *faster* effect on outlet temperature than a proportional change in steam flow rate. This is due to the relative masses impacting the temperature of each fluid. The oil's temperature is primarily coupled to the temperature of the tubes, whereas the steam's temperature is coupled to both the tubes and the shell of the heat exchanger. So, the steam has a greater mass to heat than the oil has to cool, giving the steam a larger thermal time constant than the oil.


For the sake of illustration, we will assume transport delays are short enough to ignore²⁴, so we

²³Knowing this allows us to avoid measuring the incoming cold oil temperature and just measure incoming cold oil flow rate as the feedforward variable. If the incoming oil's temperature were known to vary substantially over time, we would be forced to measure it as well as flow rate, combining the two variables together to calculate the *energy demand* and use this inferred variable as the feedforward variable.

²⁴Transport delay (dead time) in heat exchanger systems can be a thorny problem to overcome, as they tend to change with flow rate! For reasons of simplicity in our illustration, we will treat this process as if it only possessed lag times, not dead times.


are only dealing with different *lag* times between the oil flow's effect on temperature and the steam flow's effect on temperature.

This is what would happen to the heated oil temperature if steam flow were held constant and oil flow were suddenly increased:




Increased oil flow convects heat away from the steam at a faster rate than before, resulting in decreased oil temperature. This drop in temperature is fairly quick, and is self-regulating.

By contrast, this is what would happen to the heated oil temperature if oil flow were held constant and steam flow were suddenly increased:



Increased steam flow convects heat into the oil at a faster rate than before, resulting in increased oil temperature. This rise in temperature is also self-regulating, but much slower than the temperature change resulting from a proportional adjustment in oil flow. In other words, the time constant (τ) of the process with regard to steam flow changes is greater than the time constant of the process with regard to oil flow changes ($\tau_{steam} > \tau_{oil}$).


If we superimpose these two effects, as will be the case when the feedforward system is working (without the benefit of feedback “trim” control), what we will see when oil flow suddenly increases is a “fight” between the cooling effect of the increased oil flow and the heating effect of the increased steam flow. However, it will not be a fair fight: the oil flow’s effect will temporarily win over the steam’s effect because of the oil’s faster time constant. Another way of stating this is to say the feedforward action *temporarily under-compensates* for the change in load. The result will be a momentary dip in outlet temperature before the system achieves equilibrium again:



The solution to this problem is not unlike the solution we applied to the water treatment system: we must somehow equalize these two lag times so their superimposed effects will directly cancel, resulting in an undisturbed process variable. An approximate solution for equalizing two different lag times is to cascade two lags together in order to emulate one larger lag time²⁵. This may be done by inserting a lag time relay or function block in the feedforward system.

When we look at our P&ID, though, a problem is immediately evident. The lag time we need to slow down is the lag time of the oil flow’s effect on temperature. In this system, oil flow is a wild variable, not something we have the ability to control (or delay at will). Our feedforward control system can only manipulate the steam valve position in response to oil flow, not influence oil flow in order to give the steam time to “catch up.”


If we cannot slow down the time constant inherent to the wild variable (oil flow), then the best we can do is speed up the time constant of the variable we do have influence over (steam flow). The solution is to insert something called a *lead function* into the feedforward signal driving the steam valve. A “lead” is the mathematical inverse of a lag. If a lag is modeled by an RC low-pass filter circuit, then a “lead” is modeled by an RC high-pass filter circuit:



²⁵Technically, two cascaded lag times is not the same as one large lag time, no matter the time constant values. Two first-order lags in series with one another create a *second-order lag*, which is a different effect. However imperfect as the added lag solution is, it is still better than nothing at all!

Being mathematical inverses of each other, a lead function should perfectly cancel a lag function when the output of one is fed to the input of the other, and when the time constants of each are equal. If the time constants of lead and lag are not equal, their cascaded effect will be a partial cancellation. In our heat exchanger control application, this is what we need to do: partially cancel the steam valve's slow time constant so it will be more equal with the oil flow's time constant. Therefore, we need to insert a lead function into the feedforward signal path.

A lead function will take the form of either a physical signal relay or (more likely with modern technology) a function block executed inside a digital control system. The proper place for the lead function is between the oil flow transmitter and the summation function:




Now, when the oil flow rate to this heat exchanger suddenly increases, the lead function will add a “surge” to the feedforward signal before it goes to the summing function, quickly opening the steam valve further than usual and sending a surge of steam to the exchanger to help overcome the naturally sluggish response of the oil temperature to changes in steam flow. The feedforward action won’t be perfect with this lead function added, but it will be substantially better than if there was no dynamic compensation added to the feedforward signal.

31.6.3 Lead/Lag and dead time function blocks


The addition of dynamic compensation in a feedforward control system may require a lag function, a lead function, and/or a dead time function, depending on the nature of the time delay differences between the relevant process load and the system's corrective action. Modern control systems provide all these functions as digital *function blocks*. In the past, these functions could only be implemented in the form of individual instruments with these time characteristics, called *relays*. As we have already seen, lead and lag functions may be rather easily implemented as simple RC filter circuits. Pneumatic equivalents also exist, which were the only practical solution in the days of pneumatic transmitters and controllers. Dead time is notoriously difficult to emulate using analog components of any kind, and so it was common to use lag-time elements (sometimes more than one connected in series) to provide an approximation of dead time.

With digital computer technology, all these dynamic compensation functions are easy to implement and readily available in a control system. Some single-loop controllers even have these capabilities programmed within, ready to use when needed.

A dead time function block is most easily implemented using the concept of a *first-in, first-out shift register*, sometimes called a *FIFO*. With this concept, successive values of the input variable are stored in a series of registers (memory), their progression to the output delayed by a certain amount of time:




Lead and lag functions are also implemented digitally in modern controllers and control systems, but they are actually easier to comprehend in their analog (RC circuit) forms. The most common way lead and lag functions are found in modern control systems is in combination as the so-called *lead/lag function*, merging both lead and lag characteristics in a single function block (or relay):




Each parallel RC subcircuit represents a time constant (τ), one for lead and one for lag. The overall behavior of the network is determined by the relative magnitudes of these two time constants. Which ever time constant is larger, determines the overall characteristic of the network.

If the two time constant values are equal to each other ($\tau_{lead} = \tau_{lag}$), then the circuit performs no dynamic compensation at all, simply passing the input signal to the output with no change except for some attenuation:




A square wave signal entering this network will exit the network as a square wave. If the input signal is sinusoidal, the output will also be sinusoidal and in-phase with the input.

If the lag time constant exceeds the lead time constant ($\tau_{lag} > \tau_{lead}$), then the overall behavior of the circuit will be to introduce a first-order lag to the voltage signal:



A square wave signal entering the network will exit the network as a sawtooth-shaped wave. A sinusoidal input will emerge sinusoidal, but with a lagging phase shift. This, in fact, is where the *lag* function gets its name: from the negative (“lagging”) phase shift it imparts to a sinusoidal input.


Conversely, if the lead time constant exceeds the lag time constant ($\tau_{lead} > \tau_{lag}$), then the overall behavior of the circuit will be to introduce a first-order lead to the voltage signal (a step-change voltage input will cause the output to “spike” and then settle to a steady-state value):




A square wave signal entering the network will exit the network with sharp transients on each leading edge. A sinusoidal input will emerge sinusoidal, but with a leading phase shift. Not surprisingly, this is where the *lead* function gets its name: from the positive (“leading”) phase shift it imparts to a sinusoidal input.

This exact form of lead/lag circuit finds application in a context far removed from process control: compensation for coaxial cable capacitance in a $\times 10$ oscilloscope probe. Such probes are used to extend the voltage measurement range of standard oscilloscopes, and/or to increase the impedance of the instrument for minimal loading effect on sensitive electronic circuits. Using a $\times 10$ probe, an oscilloscope will display a waveform that is $\frac{1}{10}$ the amplitude of the actual signal, and present ten times the normal impedance to the circuit under test.



If a $9\text{ M}\Omega$ resistor is connected in series with a standard oscilloscope input (having an input impedance of $1\text{ M}\Omega$) to create a 10:1 voltage division ratio, problems will result from the cable capacitance connecting the probe to the oscilloscope input. What should display as a square-wave input instead looks “rounded” by the effect of capacitance in the coaxial cable and at the oscilloscope input:




The reason for this is signal distortion the combined effect of the $9\text{ M}\Omega$ resistor and the cable's natural capacitance forming an RC network:



A simple solution to this problem is to build the 10:1 probe with a variable capacitor connected in parallel across the $9\text{ M}\Omega$ resistor. The combination of the $9\text{ M}\Omega$ resistor and this capacitor creates a lead network to cancel out the effects of the lag caused by the cable capacitance and $1\text{ M}\Omega$ oscilloscope impedance in parallel. When the capacitor is properly adjusted, the oscilloscope will accurately show the shape of any waveform at the probe tip, including square waves:




If we re-draw this compensated probe circuit to show the resistor pair and the capacitor pair both working as 10:1 voltage dividers, it becomes clearer to see how the two divider circuits work in parallel with each other to provide the same 10:1 division ratio *only* if the component ratios are properly proportioned:



Compensation is effective because
 $(9R)(C) = (R)(9C)$
 $\tau_1 = \tau_2$

The series resistor pair forms an obvious 10:1 division ratio, with the smaller of the two resistors being in parallel with the oscilloscope input. The upper resistor, being 9 times larger in resistance value, drops 90% of the voltage applied to the probe tip. The series capacitor pair forms a less obvious 10:1 division ratio, with the cable capacitance being the larger of the two. Recall that the *reactance* of a capacitor to an AC voltage signal is inversely related to capacitance: a larger capacitor will have fewer ohms of reactance for any given frequency according to the formula $X_C = \frac{1}{2\pi fC}$. Thus, the capacitive voltage divider network has the same 10:1 division ratio as the resistive voltage divider network, even though the capacitance ratios may look “backward” at first glance.

If the compensation capacitor is adjusted to an excessive value, the probe will *overcompensate* for lag (too much lead), resulting in a “spiked” waveform on the oscilloscope display with a perfect square-wave input:



With the probe's compensating capacitor exhibiting an excessive amount of capacitance, the capacitive voltage divider network has a voltage division ratio *less than* 10:1. This is why the waveform “spikes” on the leading edges: the capacitive divider dominates the network's response in the short term, producing a voltage pulse at the oscilloscope input greater than it should be (divided by some ratio less than 10). Soon after the leading edge of the square wave passes, the capacitors' effects will wane, leaving the resistors to establish the voltage division ratio on their own. Since the two resistors have the proper 10:1 ratio, this causes the oscilloscope's signal to “settle” to its proper value over time. Thus, the waveform “spikes” too far at each leading edge and then decays to its proper amplitude over time.

While undesirable in the context of oscilloscope probes, this is precisely the effect we desire in a process control *lead* function. The purpose of a lead/lag function is to provide a signal gain that begins at some initial value, then “settles” at another value over time. This way, sudden changes in the feedforward signal will either be amplified or attenuated for a short duration to compensate for lags in other parts of the control system, while the steady-state gain of the feedforward loop remains at some other value necessary for long-term stability. For a lag function, the initial gain is less than the final gain; for a lead function, the initial gain exceeds the final gain. If the lead and lag time constants are set equal to each other, the initial and final gains will likewise be equal, with the function exhibiting a constant gain at all times.

Although lead-lag functions for process control systems may be constructed from analog electronic components, modern systems implement the function arithmetically using digital microprocessors. A typical time-domain equation describing a digital lead/lag function block's output response (y) to an input step-change from zero (0) to magnitude x over time (t) is as follows:





$$y = x \left(1 + \frac{\tau_{lead} - \tau_{lag}}{\tau_{lag}} e^{-\frac{t}{\tau_{lag}}} \right)$$

As you can see, if the two time constants are set equal to each other ($\tau_{lead} = \tau_{lag}$), the second term inside the parentheses will have a value of zero at all times, reducing the equation to $y = x$. If the lead time constant exceeds the lag time constant ($\tau_{lead} > \tau_{lag}$), then the fraction will begin with a positive value and decay to zero over time, giving us the “spike” response we expect from a lead function. Conversely, if the lag time constant exceeds the lead ($\tau_{lag} > \tau_{lead}$), the fraction will begin with a negative value at time = 0 (the beginning of the step-change) and decay to zero over time, giving us the “sawtooth” response we expect from a lag function.

It should also be evident from an examination of this equation that the “decay” time of the lead/lag function is set by the lag time constant (τ_{lag}). Even if we just need the function to produce a “lead” response, we must still properly set τ_{lag} in order for the lead response to decay at the correct rate for our control system. The intensity of the lead function (i.e. how far it “spikes” when presented with a step-change in input signal) varies with the ratio $\frac{\tau_{lead}}{\tau_{lag}}$, but the duration of the “settling” following that step-change is entirely set by τ_{lag} .

To summarize the behavior of a lead/lag function:


- If $\tau_{lead} = \tau_{lag}$, the lead/lag function will simply pass the input signal through to the output (no lead or lag action at all)
- If $\tau_{lead} = 0$, the lead/lag function will provide a pure lag response with a final gain of unity and a time constant of τ_{lag}
- If $\tau_{lead} = 2(\tau_{lag})$, the lead/lag function will provide a lead response with an initial gain of 2, a final gain of unity, and a time constant of τ_{lag}



31.7 Limit, Selector, and Override controls


Another category of control strategies involves the use of signal relays or function blocks with the ability to switch between different signal values, or re-direct signals to new pathways. Such functions are useful when we need a control system to choose between multiple signals of differing value in order to make the best control decisions.

The “building blocks” of such control strategies are special relays (or function blocks in a digital control system) shown here:




High-select functions output whichever input signal has the *greatest* value. *Low-select* functions do just the opposite: output whichever input signal has the *least* value. “Greater-than” and “Less than” symbols mark these two selector functions, respectively, and each type may be equipped to receive more than two input signals.

Sometimes you will see these relays represented in P&IDs simply by an inequality sign in the middle of the large bubble, rather than off to the side in a square. You should bear in mind that the location of the input lines has no relationship at all to the direction of the inequality symbol – e.g., it is not as though a high-select relay looks for the input on the left side to be greater than the input on the right. Note the examples shown below, complete with sample signal values:



High-limit and *low-limit* functions are similar to high- and low-select functions, but they only receive one input each, and the limit value is a parameter programmed into the function rather than received from another source. The purpose of these functions is to place a set limit on how high or how low a signal value is allowed to go before being passed on to another portion of the control system. If the signal value lies within the limit imposed by the function, the input signal value is simply passed on to the output with no modification.

Like the select functions, limit functions may appear in diagrams with nothing more than the limit symbol inside the bubble, rather than being drawn in a box off to the side:




Rate limit functions place a maximum rate-of-change limit on the input signal, such that the output signal will follow the input signal precisely until and unless the input signal's rate-of-change over time ($\frac{dx}{dt}$) exceeds the pre-configured limit value. In that case, the relay still produces a ramping output value, but the rate of that ramp remains fixed at the limit $\frac{dx}{dt}$ value no matter how fast the input keeps changing. After the output value "catches up" with the input value, the function once again will output a value precisely matching the input unless the input begins to rise or fall at too fast a rate again.

31.7.1 Limit controls


A common application for select and limit functions is in *cascade* control strategies, where the output of one controller becomes the setpoint for another. It is entirely possible for the primary (master) controller to call for a setpoint that is unreasonable or unsafe for the secondary (slave) to attain. If this possibility exists, it is wise to place a limit function between the two controllers to limit the cascaded setpoint signal.

In the following example, a cascade control system regulates the temperature of molten metal in a furnace, the output of the master (metal temperature) controller becoming the setpoint of the slave (air temperature) controller. A high limit function limits the maximum value this cascaded setpoint can attain, thereby protecting the refractory brick of the furnace from being exposed to excessive air temperatures:



It should be noted that although the different functions are drawn as separate bubbles in the P&ID, it is possible for multiple functions to exist within one physical control device. In this example, it is possible to find a controller able to perform the functions of both PID control blocks (master and slave) and the high limit function as well. It is also possible to use a distributed technology such as FOUNDATION Fieldbus to place all control functions inside field instruments, so only three field instruments exist in the loop: the air temperature transmitter, the metal temperature transmitter, and the control valve (with a Fieldbus positioner).

This same control strategy could have been implemented using a low select function block rather than a high limit:



Here, the low-select function selects whichever signal value is lesser: the setpoint value sent by the master temperature controller, or the maximum air temperature limit value sent by the hand indicating controller (HIC – sometimes referred to as a *manual loading station*).


An advantage of this latter approach over the former might be ease of limit value changes. With a pre-configured limit value residing in a high-limit function, it might be that only qualified maintenance people have access to changing that value. If the decision of the operations department is to have the air temperature limit value easily adjusted by anyone, the latter control strategy's use of a manual loading station would be better suited²⁶.

Another detail to note in this system is the possibility of *integral windup* in the master controller in the event that the high setpoint limit takes effect. Once the high-limit (or low-select) function secures the slave controller's remote setpoint at a fixed value, the master controller's output is no


²⁶I generally suggest keeping such limit values inaccessible to low-level operations personnel. This is especially true in cases such as this where the presence of a high temperature setpoint limit is intended for the longevity of the equipment. There is a strong tendency in manufacturing environments to “push the limits” of production beyond values considered safe or expedient by the engineers who designed the equipment. Limits are there for a reason, and should not be altered except by people with full understanding of and full responsibility over the consequences!

longer controlling anything: it has become decoupled from the process. If, when in this state of affairs, the metal temperature is still below setpoint, the master controller's integral action will "wind up" the output value over time with absolutely no effect, since the slave controller is no longer following its output signal. If and when the metal temperature reaches setpoint, the master controller's output will likely be saturated at 100% due to the time it spent winding up. This will cause the metal temperature to overshoot setpoint, as a positive error will be required for the master controller's integral action to wind back down from saturation.

A relatively easy solution to this problem is to configure the master controller to stop integral action when the high limit relay engages. This is easiest to do if the master PID and high limit functions both reside in the same physical controller. Many digital limit function blocks generate a bit representing the state of that block (whether it is passing the input signal to the output or limiting the signal at the pre-configured value), and some PID function blocks have a boolean input used to disable integral action. If this is the case with the function blocks comprising the high-limit control strategy, it may be implemented like this:




Another method used to prevent integral windup is to make use of the *feedback* input available on some PID function blocks. This is an input used to calculate the integral term of the PID equation. In the days of pneumatic PID controllers, this option used to be called *external reset*. Normally connected to the output of the PID block, if connected to the output of the high-limit function it will let the controller know whether or not any attempt to wind up the output is having an effect. If the output has been de-selected by the high-limit block, integral windup will cease:



Limit control strategies implemented in FOUNDATION Fieldbus instruments use the same principle, except that the concept of a “feedback” signal sending information backwards up the function block chain is an aggressively-applied design philosophy throughout the FOUNDATION Fieldbus standard. Nearly every function block in the Fieldbus suite provides a “back calculation” output, and nearly every function block accepts a “back calculation” input from a downstream block. The “Control Selector” (CS) function block specified in the FOUNDATION Fieldbus standard provides the limiting function we need between the master and slave controllers. The BKCAL.OUT signal of this selector block connects to the master controller’s BKCAL.IN input, making the master controller aware of its selection status. If ever the Control Selector function block de-selects the master controller’s output, the controller will immediately know to halt integral action:

*Mitigating master controller integral
wind-up in a FOUNDATION Fieldbus
high-limit control strategy*




This same “back calculation” philosophy – whereby the PID algorithm is aware of how another function is limiting or over-riding its output – is also found in some programmable logic controller (PLC) programming conventions. The Allen-Bradley Logix5000 series of PLCs, for example, provides a *tieback* variable to force the PID function’s output to track the overriding function. When the “tieback” variable is properly used, it allows the PID function to bumplessly transition from the “in-control” state to the “overridden” state.

31.7.2 Selector controls

In the broadest sense, a “selector” control strategy is where one signal is selected from multiple signals in a system to perform a measurement control function. In the context of this book and this chapter, I will use the term “selector” to reference automatic selection among multiple *measurement* or *setpoint* signals. Selection between multiple *controller output* signals will be explored in the next subsection, under the term “override” control.

Perhaps one of the simplest examples of a selector control strategy is where we must select a process variable signal from multiple transmitters. For example, consider this chemical reactor, where the control system must throttle the flow of coolant to keep the *hottest* measured temperature at setpoint, since the reaction happens to be exothermic (heat-releasing)²⁷:




The high-select relay (TY-24) sends only the highest temperature signal from the three transmitters to the controller. The other two temperature transmitter signals are simply ignored.

Another use of selector relays (or function blocks) is for the determination of a *median* process measurement. This sort of strategy is often used on triple-redundant measurement systems, where three transmitters are installed to measure the exact same process variable, providing a valid measurement even in the event of transmitter failure.

²⁷Only the coolant flow control instruments and piping are shown in this diagram, for simplicity. In a real P&ID, there would be many more pipes, valves, and other apparatus shown surrounding this process vessel.


The median select function may be implemented one of two ways using high- and low-select function blocks:

Two ways of obtaining a median signal value from three redundant inputs




The left-hand selector strategy selects the highest value from each pair of signals (A and B, B and C, A and C), then selects the lowest value of those three primary selections. The right-hand strategy is exactly opposite – first selecting the lowest value from each input pair, then selecting the highest of those values – but it still accomplishes the same function. Either strategy outputs the *middle* value of the three input signals²⁸.

Although either of these methods of obtaining a median measurement requires four signal selector functions, it is quite common to find function blocks available in control systems ready to perform the median select function all in a single block. The median-select function is so common to redundant sensor control systems that many control system manufacturers provide it as a standard function unto itself:



²⁸In order to understand how this works, I advise you try a “thought experiment” for each function block network whereby you arbitrarily assign three different numerical values for A, B, and C, then see for yourself which of those three values becomes the output value.

This is certainly true in the FOUNDATION Fieldbus standard, where two standardized function blocks are capable of this function, the CS (Control Selector) and the ISEL (Input Selector) blocks:



Of these two Fieldbus function blocks, the latter (ISEL) is expressly designed for selecting transmitter signals, whereas the former (CS) is best suited for selecting controller outputs with its “back calculation” facilities designed to modify the response of all de-selected controllers. Using the terminology of this book section, the ISEL function block is best suited for *selector* strategies, while the CS function block is ideal for *limit* and *override* strategies (discussed in the next section).

If receiving three “good” inputs, the ISEL function block will output the middle (median) value

of the three. If one of the inputs carries a “bad” status²⁹, the ISEL block outputs the averaged value of the remaining two (good) inputs. Note how this function block also possesses individual “disable” inputs, giving external boolean (on/off) signals the ability to disable any one of the transmitter inputs to this block. Thus, the ISEL function block may be configured to de-select a particular transmitter input based on some programmed condition other than internal diagnostics.


If receiving four “good” inputs, the ISEL function block normally outputs the average value of the two middle (median) signal values. If one of the four inputs becomes “bad” is disabled, the block behaves as a normal three-input median select.

A general design principle for redundant transmitters is that you *never* install exactly two transmitters to measure the same process variable. Instead, you should install three (minimum). The problem with having two transmitters is a lack of information for “voting” if the two transmitters happen to disagree. In a three-transmitter system, the function blocks may select the median signal value, or average the “best 2 out of 3.” If there are just two transmitters installed, and they do not substantially agree with one another, it is anyone’s guess which one should be trusted³⁰.

²⁹In FOUNDATION Fieldbus, each and every signal path not only carries the signal value, but also a “status” flag declaring it to be “Good,” “Bad,” or “Uncertain.” This status value gets propagated down the entire chain of connected function blocks, to alert dependent blocks of a possible signal integrity problem if one were to occur.

³⁰This principle holds true even for systems with no function blocks “voting” between the redundant transmitters. Perhaps the installation consists of two transmitters with remote indications for a human operator to view. If the two displays substantially disagree, which one should the operator trust? A set of *three* indicators would be much better, providing the operator with enough information to make an intelligent decision on which display(s) to trust.

A classic example of selectors in industrial control systems is that of a *cross-limited ratio control* strategy for air/fuel mixture applications. Before we explore the use of selector functions in such a strategy, we will begin by analyzing a simplified version of that strategy, where we control air and fuel flows to a precise ratio with no selector action at all:




Here, the fuel flow controller receives its setpoint directly from the firing command signal, which may originate from a human operator’s manual control or from the output of a temperature controller regulating temperature of the combustion-heated process. The air flow controller receives its setpoint from the fuel flow transmitter, with the calibrations of the air and fuel flow transmitters being appropriate to establish the proper air:fuel flow ratio when the transmitters register equally. From the perspective of the air flow controller, fuel flow is the *wild* flow while air flow is the *captive* flow.

There is a problem with this control system that may not be evident upon first inspection: the air:fuel ratio will tend to vary as the firing command signal increases or decreases in value over time. This is true even if the controllers are well-tuned and the air:fuel ratio remains well-controlled under steady-state conditions. The reason for this is linked to the roles of “wild” and “captive” flows, fuel and air flow respectively. Since the air flow controller receives its setpoint from the fuel flow transmitter, changes in air flow will always *lag* behind changes in fuel flow.

This sort of problem can be difficult to understand because it involves changes in multiple variables over time. A useful problem-solving technique to apply here is a “thought experiment,” coupled with a time-based graph to display the results. Our thought experiment consists of imagining what would happen if the firing command signal were to *suddenly* jump in value, then sketching the results on a graph.

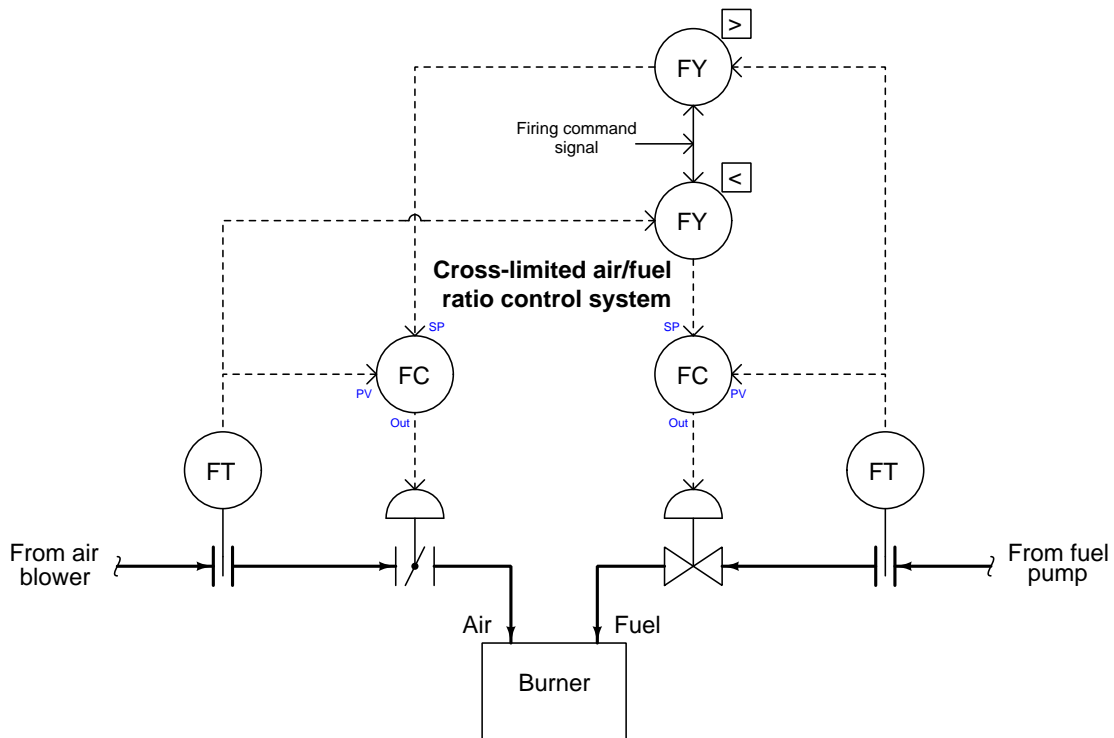
If the firing command signal suddenly increases, the fuel flow controller responds by opening up the fuel valve, which after a slight delay results in increased fuel flow to the burner. This increased fuel flow signal gets sent to the setpoint input of the air flow controller, which in turn opens up the air valve to increase air flow proportionately. If the firing command signal suddenly decreased, the same changes in flow would occur in reverse direction but in the same chronological sequence, since the fuel flow change still “leads” the subsequent air flow change:



Inevitable delays in controller response, valve response, and flow transmitter response conspire to upset the air:fuel ratio during the times immediately following a step-change in firing command signal. When the firing command steps up, the fuel flow increases before the air flow, resulting in a short time when the burner runs “rich” (too much fuel, not enough air). When the firing command steps down, the fuel flow decreases before the air flow, resulting in a short time when the burner runs “lean” (too much air, not enough fuel). The scenario problem is dangerous because it may result in an explosion if an accumulation of unburnt fuel collects in a pocket of the combustion chamber and then ignites. The second scenario is generally not a problem unless the flame burns *so* lean that it risks blowing out.


The solution to this vexing problem is to re-configure the control scheme so that the air flow controller “takes the lead” whenever the firing command signal rises, and that the fuel flow controller “takes the lead” whenever the firing command signal falls. This way, any upsets in air:fuel ratio resulting from changes in firing command will always err on the side of a lean burn rather than a rich burn.

We may implement precisely this strategy by adding some signal selector functions to our ratio control system. The ratio is now *cross-limited*, where both measured flow rates serve as limiting variables to each other to ensure the air:fuel ratio can never be too rich:



Now, the air flow controller receives its setpoint either directly from the firing command signal or from the fuel flow transmitter, whichever signal is *greater* in value. The fuel flow controller receives its setpoint either directly from the firing command signal or from the air flow transmitter, whichever signal is *least* in value. Thus, the air flow controller “takes the lead” and makes fuel flow the “captive” variable when the firing command signal rises. Conversely, the fuel flow controller “takes the lead” and makes air flow the “captive” variable when the firing command signal falls. Instead of having the roles of “wild” and “captive” flows permanently assigned, these roles switch depending on which way the firing command signal changes.

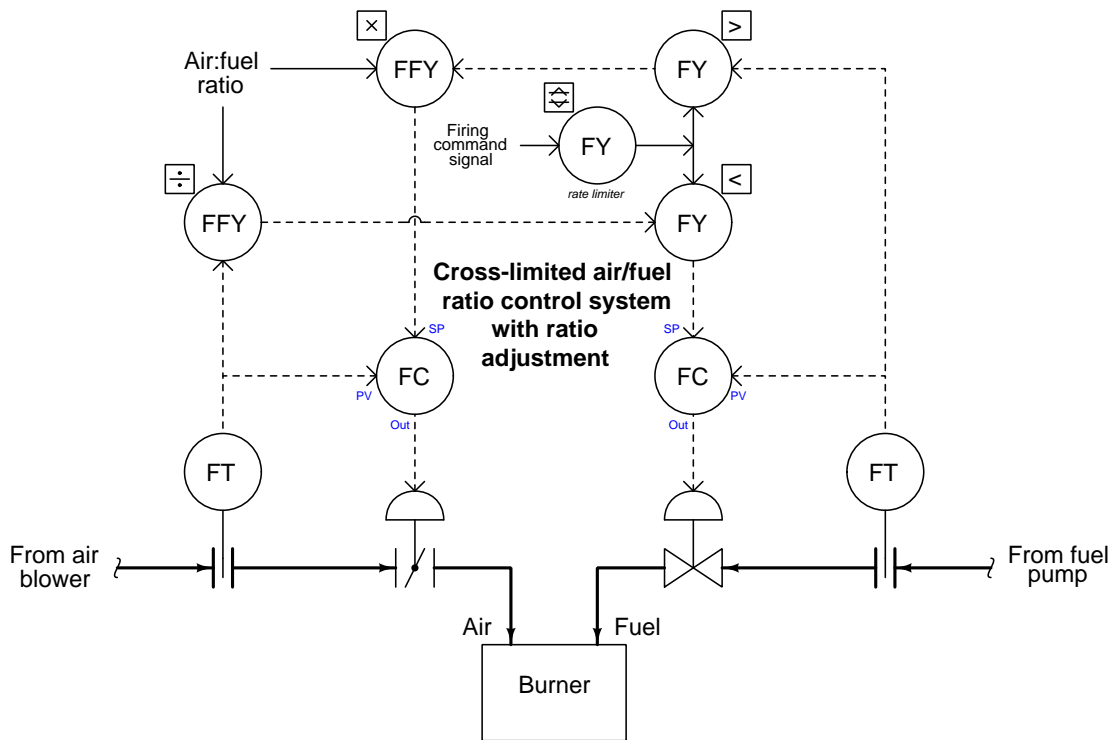
Examining the response of this cross-limited system to sudden changes in firing command signal, we see how the air flow controller takes the lead whenever the firing rate signal increases, and how the fuel flow controller takes the lead whenever the firing rate signal decreases:



In both transient scenarios, the mixture runs lean (safe) rather than rich (dangerous). Of course, care must be taken to ensure the firing rate signal never steps up or down so quickly that the flame runs lean enough to blow out (i.e. the mixture becomes much too lean during a transient "step-change" of the firing rate signal). If this is a problem, we may fix it by installing *rate-limiting* functions in the firing command signal path, so that the firing command signal can never rise or fall too rapidly.

A realistic cross-limited ratio control system also incorporates a means to adjust the air:fuel ratio without having to re-range the air and/or fuel flow transmitters. Such ratio adjustment may be achieved by the insertion of a “multiplying” function between one of the selectors and a controller setpoint, plus a “dividing” function to return that scaled flow to a normalized value for cross-limiting.

The complete control strategy looks something like this, complete with cross-limiting of air and fuel flows, rate-limiting of the firing command signal, and adjustable air:fuel ratio:




31.7.3 Override controls

An “override” control strategy involves a selection between two or more controller *output* signals, where only one controller at a time gets the opportunity to exert control over a process. All other “de-selected” controllers are thus *overridden* by the selected controller.

The general concept of override control is easily understood by appeal to a human example. Cargo truck drivers must monitor and make control decisions on a wide number of variables, including diesel engine operating parameters and road rules. A truck driver needs to keep a close watch on the exhaust gas temperature of the truck engine: a leading indicator of impending engine damage (if exhaust temperature exceeds a pre-determined limit established by the engine manufacturer). The same truck driver must also drive as fast as the law will allow on any given road in order to minimize shipping time and thereby maximize the amount of cargo transported over long periods of time. These two goals may become mutually exclusive when hauling heavy cargo loads up steep inclines, such as when ascending a mountain pass. The goal of avoiding engine damage necessarily overrides the goal of maintaining legal road speed in such conditions.

Imagine a diesel truck driver maintaining the legal speed limit on a highway, occasionally glancing at the EGT (Exhaust Gas Temperature) indicator in the instrument panel. Under normal operating conditions, the EGT should be well below the danger threshold for the engine. However, after pulling a full load up a mountain pass and noticing the EGT approach the high operating limit, the truck driver makes the decision to regulate the engine’s power based on EGT rather than road speed. In other words, the legal speed limit is no longer the “setpoint” to control to, and EGT now is.

If we were to model the truck driver’s decision-making processes in industrial instrumentation terms, it would look something like this:

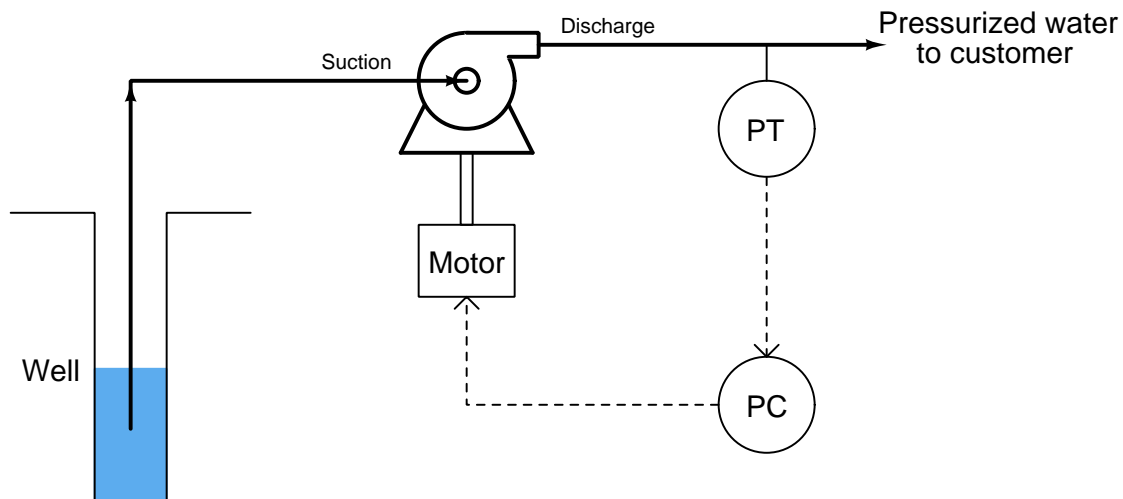


Which ever control decision calls for the least engine power output, “wins the vote” to control the engine’s power.

As is the case with limit and selector control strategies, a “select” function is used to choose one signal from multiple signals. The difference here is that the signals being selected are both *controller outputs* rather than transmitter (measurement) or setpoint signals. Both controllers are still active, but only one at a time will have any actual control over the process.

This model maps well to the truck driver analogy. Despite having “overridden” the goal of maintaining legal road speed in favor of maintaining a safe engine exhaust temperature, the driver is still thinking about road speed. In fact, if the driver happens to be behind schedule, you can be absolutely sure the goal of maintaining the highway speed limit has not been forgotten! In fact, the driver may become impatient as the long incline wears on, eager to make up lost time as soon as the opportunity allows. This is a potential problem for all override control systems: making sure the de-selected controller does not “wind up” (with integral action still active) while it has no control over the process.

An municipal example of override control is seen in this water pumping system, where a water pump is driven by a variable-speed³¹ electric motor to draw water from a well and provide constant water pressure to a customer:




Incidentally, this is an excellent application for a variable-speed motor as the final control element rather than a control valve. Reducing pump speed in low-flow conditions will save a lot of energy over time compared to the energy that would be wasted by a constant-speed pump and control valve.

A potential problem with this system is the pump running “dry” if the water level in the well gets too low, as might happen during summer months when rainfall is low and customer demand is high. If the pump runs for too long with no water passing through it, the seals will become damaged. This will necessitate a complete shut-down and costly rebuild of the pump, right at the time customers need it the most.

³¹In most applications this takes the form of an AC induction motor receiving power from a *Variable Frequency Drive* or *VFD*. Since the rotational speed of an induction motor is a function of frequency, the VFD achieves motor speed control by electronically converting the fixed-frequency line power into variable-frequency power to drive the motor.


One solution to this problem would be to install a level switch in the well, sensing water level and shutting off the electric motor driving the pump if the water level ever gets too low:



This may be considered a kind of “override” strategy, because the low-level switch over-rides the pressure controller’s command for the pump to turn. It is also a crude solution to the problem, for while it protects the pump from damage, it does so at the cost of completely shutting off water to customers. One way to describe this control strategy would be to call it a *hard override* system, suggesting the uncompromising action it will take to protect the pump.

A better solution to the dilemma would be to have the pump merely slow down as the well water level approaches a low-level condition. This way at least the pump could be kept running (and some amount of pressure maintained), decreasing demand on the well while maintaining curtailed service to customers and still protecting the pump from dry-running. This would be termed a *soft override* system.


We may create just such a control strategy by replacing the well water level switch with a level transmitter, connecting the level transmitter to a level controller, and using a low-select relay or function block to select the lowest-valued output between the pressure and level controllers. The level controller's setpoint will be set at some low level above the acceptable limit for continuous pump operation:



If ever the well's water level goes below this setpoint, the level controller will command the pump to slow down, even if the pressure controller is calling for a higher speed. The level controller will have *overridden* the pressure controller, prioritizing pump longevity over customer demand.

Bear in mind that the concept of a low-level switch completely shutting off the pump is not an entirely bad idea. In fact, it might be prudent to integrate such a "hard" shutdown control in the override control system, just in case something goes wrong with the level controller (e.g. an improperly adjusted setpoint or poor tuning) or the low-select function.

With two layers of safety control for the pump, this system provides both a “soft constraint” providing moderated action and a “hard constraint” providing aggressive action to protect the pump from dry operation:



In order that these two levels of pump protection work in the proper order, the level controller's (LC) setpoint needs to be set to a higher value than the low level alarm's (LAL) trip point.

A very important consideration for any override control strategy is how to manage integral windup. Any time a controller with any integral (reset) action at all is de-selected by the selector function, the integral term of the controller will have the tendency to wind up (or wind down) over time. With the output of that controller de-coupled from the final control element, it can have no effect on the process variable. Thus, integral control action – the purpose of which being to constantly drive the output signal in the direction necessary to achieve equality between process variable and setpoint – will work in vain to eliminate an error it cannot influence. If and when control is handed back to that controller, the integral action will have to spend time “winding” the other way to un-do what it did while it was de-selected.

Thus, override controls demand some form of integral windup limits that engage when a controller is de-selected. Methods of accomplishing this function are discussed in an earlier section on limit controls (section 31.7.1 beginning on page 2568).


31.8 Techniques for analyzing control strategies

Control strategies such as cascade, ratio, feedforward, and those containing limit and selector functions can be quite daunting to analyze, especially for students new to the subject. As a teacher, I have seen first-hand where students tend to get confused on these topics, and have seen how certain problem-solving techniques work well to overcome these conceptual barriers. This section explores some of these techniques and the reasons why they work.

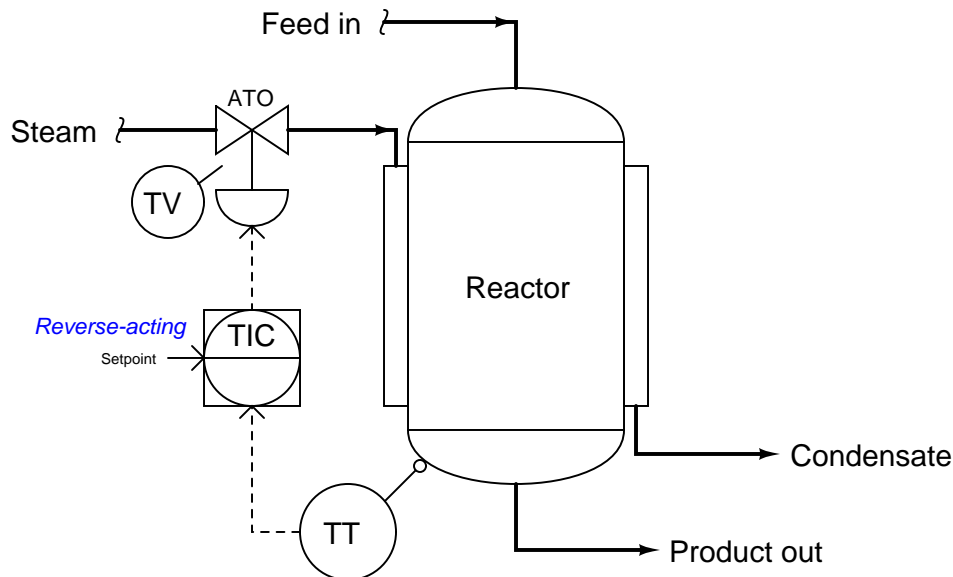
31.8.1 Explicitly denoting controller actions

The direction of action for a loop controller – either *direct* or *reverse* – at first seems like a very simple concept. It certainly is fundamental to the comprehension of any control strategy containing PID loop controllers, but this seemingly simple concept harbors an easy-to-overlook fact causing much confusion for students as they begin to analyze any control strategy where a loop controller receives a remote setpoint signal from some other device, most notably in cascade and ratio control strategies.

A *direct-acting* loop controller is defined as one where the output signal increases as the process variable signal increases. A *reverse-acting* controller is defined as one where the output signal decreases as the process variable signal increases. Both types of action are shown here:



Let us apply this concept to a realistic application, in this case the control of temperature in a steam-heated chemical reactor vessel:




As the reactor vessel's temperature increases, we need the temperature controller (TIC) to reduce the amount of hot steam entering the jacket in order to stabilize that temperature. Since the steam control valve is air-to-open (ATO), this means we need the controller to output a *decreasing* signal as the process variable (temperature) signal increases. This, by definition, is a *reverse-acting* controller. This example also showcases the utility of the problem-solving technique known as a "thought experiment," whereby we imagine a certain condition changing (in this case, the reactor temperature increasing) and then we mentally model the desired response of the system (in this case, closing the steam valve) in order to determine the necessary controller action.

So far, this example poses no confusion. But suppose we were to perform another thought experiment, this time supposing the *setpoint* signal increases rather than the reactor temperature increases. How will the controller respond now?

Many students will conclude that the controller's output signal will once again decrease, because we have determined this controller's action to be *reverse*, and "reverse" implies the output will go the opposite direction as the input. However, this is not the case: the controller output will actually *increase* if its setpoint signal is increased. This, in fact, is precisely how any reverse-acting controller should respond to an increase in setpoint.


The reason for this is evident if we take a close look at the characteristic equation for a reverse-acting proportional controller. Note how the gain is multiplied by the difference between setpoint and process variable. Note how the process variable has a negative sign in front of it, while setpoint does not.



An increase in process variable (PV) causes the quantity inside the parentheses to become more negative, or less positive, causing the output to decrease toward 0%. Conversely, an increase in setpoint (SP) causes the quantity inside the parentheses to become more positive, causing the output to increase toward 100%. This is precisely how any loop controller should respond: with the setpoint having the opposite effect of the process variable, because those two quantities are always being *subtracted* from one another in the proportional controller's equation.

Where students get confused is the single label of either “direct” or “reverse” describing a controller's action. We define a controller as being either “direct-acting” or “reverse-acting” based on how it responds to changes in process variable, but it is easy to overlook the fact that the controller's setpoint input must necessarily have the *opposite* effect. What we really need is a way to more clearly denote the respective actions of a controller's two inputs than a single word.


Thankfully, such a convention already exists in the field of electronics³², where we must denote the “actions” of an operational amplifier’s two inputs. In the case of an opamp, one input has a direct effect on the output (i.e. a change in signal at that input drives the output the same direction) while the other has a reverse effect on the output (i.e. a change in signal at that input drives the output in the opposite direction). Instead of calling these inputs “direct” and “reverse”, however, they are conventionally denoted as *noninverting* and *inverting*, respectively. If we draw a proportional controller as though it were an opamp, we may clearly denote the actions of both inputs in this manner:



I strongly recommend students label the loop controllers in any complex control strategy in the same manner, with “+” and “-” labels next to the PV and SP inputs for each controller, in order to unambiguously represent the effects of each signal on a controller’s output. This will be far more informative, and far less confusing, than merely labeling each controller with the word “direct” or “reverse”.

³²Some differential pressure transmitter manufacturers, such as Bailey, apply the same convention to denote the actions of a DP transmitter’s two pressure ports: using a “+” label to represent direct action (i.e. increasing pressure at this port drives the output signal up) and a “-” symbol to represent reverse action (i.e. increasing pressure at this port drives the output signal down).


Let us return to our example of the steam-heated reactor to apply this technique, labeling the reverse-acting controller's process variable input with a “-” symbol and its setpoint input with a “+” symbol:



With these labels in place we can see clearly how an increase in temperature going into the “-” (inverting) input of the temperature controller will drive the valve signal down, counter-acting the change in temperature and thereby stabilizing it. Likewise, we can see clearly how an increase in setpoint going into the “+” (noninverting) input of the temperature controller will drive the valve signal up, sending more steam to the reactor to achieve a greater temperature.

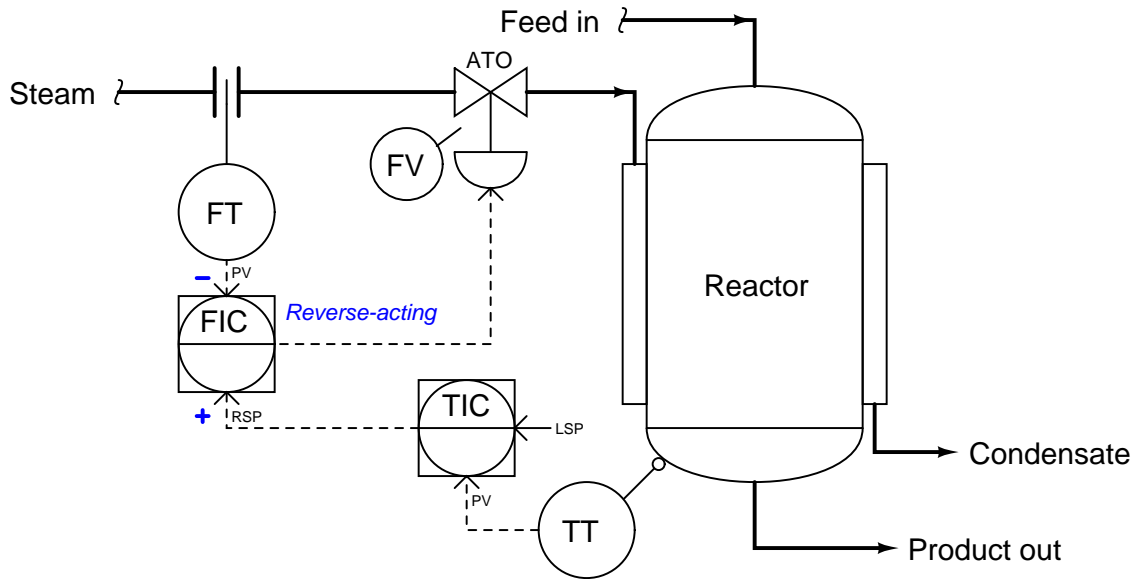
While this technique of labeling the PV and SP inputs of a loop controller as though it were an operational amplifier is helpful in single-loop controller systems, it is incredibly valuable when analyzing more complex control strategies where the setpoint to a controller is a live signal rather than a static value set by a human operator. In fact, it is for this very reason that many students do not begin to have trouble with this concept until they begin to study cascade control, where one controller provides a live (“remote”) setpoint value to another controller. Up until that point in their study, they never rarely had to consider the effects of a setpoint change on a control system because the setpoint value for a single-loop controller is usually static.

Let us modify our steam-heated reactor control system to include a cascade strategy, where the temperature controller drives a setpoint signal to a “slave” steam flow controller:




In order to determine the proper actions for each controller in this system, it is wise to begin with the slave controller (FIC), since the master controller (TIC) depends on the slave controller being properly configured in order to do its job properly. Just as we would first tune the slave controller in a cascade system prior to tuning the master controller, we should first determine the correct action for the slave controller prior to determining the correct action for the master controller.

Once again we may apply a “thought experiment” to this system in order to choose the appropriate slave controller action. If we imagine the steam flow rate suddenly increasing, we know we need the control valve to close off in order to counter-act this change. Since the valve is still air-to-open, this requires a decrease in the output signal from the FIC. Thus, the FIC must be reverse-acting. We shall denote this with a “-” label next to the process variable (PV) input, and a “+” label next to the remote setpoint (RSP) input:



Now that we know the slave controller must be reverse-acting, we may choose the action of the master controller. Applying another “thought experiment” to this system, we may imagine the reactor temperature suddenly increasing. If this were to happen, we know we would need the control valve to close off in order to counter-act this change: sending less steam to a reactor that is getting too hot. Since the valve is air-to-open, this requires a decrease in the output signal from the FIC. Following the signal path backwards from the control valve to the FIC to the TIC, we can see that the TIC must output a decreasing signal to the FIC, calling for less steam flow. A decreasing output signal at the TIC enters the FIC’s noninverting (“+”) input, causing the FIC output signal to also decrease. Thus, we need the TIC to be reverse-acting as well. We shall denote this with a “-” label next to the process variable (PV) input, and a “+” label next to the local setpoint (LSP) input:




With these unambiguous labels in place at each controller’s inputs, we are well-prepared to qualitatively analyze the response of this cascade control system to process upsets, to instrument failure scenarios, or to any other change. No longer will we be led astray by the singular label of “reverse-acting”, but instead will properly recognize the different directions of action associated with each input signal to each controller.

31.8.2 Determining the design purpose of override controls

Override control strategies are a source of much confusion for students first learning the concept. Perhaps the most fundamental question students find difficult to answer when faced with an override strategy is how to determine the intended purpose for that strategy if no explanation is given.

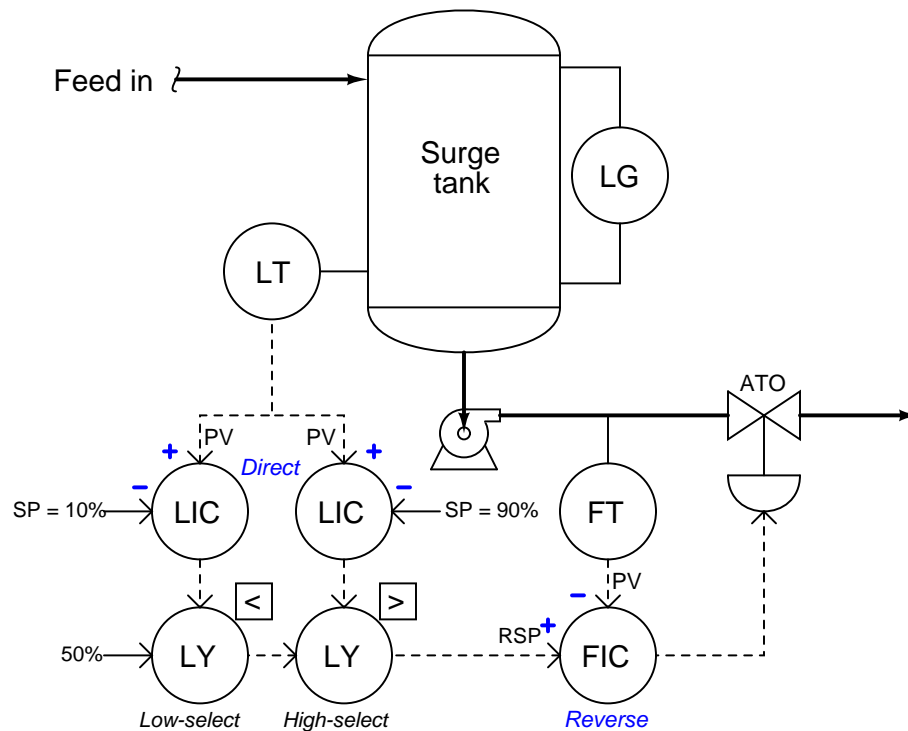
Take for example this surge tank level/flow control system. While it may be obvious that the flow controller is there for the purpose of regulating flow out of the tank, it is not so clear what the two level controllers are doing, or what purposes are served by the two selector functions:



A good starting point in our analysis is to first determine the proper directions of action for each controller. This is wise because the selector functions perform their tasks based on the relative values of the controller output signals: controllers become selected or de-selected on the basis of their output signals being greater or less than some other signal. Therefore, before we may be able to determine the purpose of a selector function, we must know how the loop controller feeding that selector function is supposed to react to process conditions. Once we have determined each controller's proper action, we may then interpret each selector's function in light of what process conditions will cause a particular controller to become selected.

When choosing the proper action for each controller, we must consider each controller in this system – one at a time – as though it were the one being selected. In other words, we may give ourselves license to ignore the selector functions and just concentrate for the time being on how each controller needs to act in order to do its job when selected. Looking at the system from this perspective, we see that each level controller (when selected) acts as a master to the flow (slave) controller. Thus, what we have here is a cascade level/flow control system, with two master controllers selected on the basis of their output signals.

The flow controller (FIC) needs to be reverse-acting, because in order to counter-act an increase in flow rate it must close off the valve (i.e. decreasing output with increasing input = reverse action). Each level controller needs to be direct-acting, because in order to counter-act an increase in level it must call for more flow exiting the tank (i.e. increasing output with increasing input = direct action). Denoting these actions using “+” and “-” labels at each PV and SP input line:



Only now are we prepared to analyze the purpose of each selector function. Let's begin with the low-select first. It selects the lowest of two values, either a fixed value of 50% or the output of the level controller with the 10% setpoint. Since we know this level controller is direct-acting, we may conclude that it will be selected if it sees a *low level* at its PV input. More specifically, it will be surely be selected if the measured tank level drops significantly below the setpoint value of 10%. Thus, we may conclude that the purpose of this level controller is to take over control if the tank level reaches or drops below the 10% mark.

Next, let us analyze the purpose of the other level controller (connected to the high-select function). Since the high-select function will select this level controller only if its output signal

exceeds the signal passed on by the low-select function, and we know that this controller is direct-acting, we may conclude that it will be selected if it sees a *high level* at its PV input. More specifically, it will surely be selected if the measured tank level rises significantly above the setpoint value of 90%. Thus, we may conclude that the purpose of this level controller is to take over control if the tank level reaches or exceeds the 90% mark.

If neither level controller is selected, the signal that gets passed on to the flow controller as a remote setpoint is the 50% fixed signal entering on the left-hand side of the low-select function. Thus, the flow controller tries to maintain a steady flow rate of 50% in the event neither level controller is selected.

Putting all these pieces together, we may conclude that the purpose of this surge tank control system is to maintain as steady a flow rate as possible out of the tank (and on to some other process), letting the liquid level inside the tank rise and fall significantly before any action is taken to change the flow rate. Only if the level drops below 10% will the flow rate be reduced, and only if the level rises above 90% will the flow rate be increased. Otherwise, the flow rate will hold steady at 50%.

To summarize, the recommended technique for analyzing the purpose of an override control system is as follows:

- First, determine the necessary actions of each controller (assume the selector functions are absent, and each controller gets its turn controlling the process).
- Identify the type of selection (high or low) implemented by each selector function.
- Based on the type of selection and the action of the controller, identify what process condition will cause that controller to become selected. *This is the condition the controller exists to regulate.*

31.9 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Negative feedback:** when the output of a system is degeneratively fed back to the input of that same system, the result is decreased (overall) gain and greater stability. Relevant to loop controller action: in order for a control system to be stable, the feedback must be negative.
- **Time constant:** (τ), defined as the amount of time it takes a system to change 63.2% of the way from where it began to where it will eventually stabilize. The system will be within 1% of its final value after 5 time constants' worth of time has passed (5τ). Relevant to process control loops, where natural lags contribute to time constants, usually of multiple order.

References

Austin, George T., *Shreve's Chemical Process Industries*, McGraw-Hill Book Company, New York, NY, 1984.

Dorf, Richard C., *Modern Control Systems*, Fifth Edition, Addison-Wesley Publishing Company, Reading, MA, 1989.

Eckman, Donald P., *Automatic Process Control*, John Wiley & Sons, Inc., New York, NY, 1958.

“FoundationTM Fieldbus Blocks”, document 00809-0100-4783, Rev BA, Rosemount, Inc., Chanhassen, MN, 2000.

“Function Blocks Instruction Manual”, document FBLOC-FFME, Smar Equipamentos Ind. Ltda., Sertãozinho, Brazil, 2005.

Lavigne, John R., *An Introduction To Paper Industry Instrumentation*, Miller Freeman Publications, Inc., San Francisco, CA, 1972.

Lavigne, John R., *Instrumentation Applications for the Pulp and Paper Industry*, The Foxboro Company, Foxboro, MA, 1979.

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Control Volume II*, Third Edition, CRC Press, Boca Raton, FL, 1999.

Mollenkamp, Robert A., *Introduction to Automatic Process Control*, Instrument Society of America, Research Triangle Park, NC, 1984.

Palm, William J., *Control Systems Engineering*, John Wiley & Sons, Inc., New York, NY, 1986.

Shinskey, Francis G., *Energy Conservation through Control*, Academic Press, New York, NY, 1978.

Shinskey, Francis G., *Process-Control Systems – Application / Design / Adjustment*, Second Edition, McGraw-Hill Book Company, New York, NY, 1979.

Chapter 32

Process safety and instrumentation

This chapter discusses instrumentation issues related to industrial process safety. Instrumentation safety may be broadly divided into two categories: how instruments themselves may pose a safety hazard (electrical signals possibly igniting hazardous atmospheres), and how instruments and control systems may be configured to detect unsafe process conditions and automatically shut an unsafe process down.

In either case, the intent of this chapter is to help define and teach how to mitigate hazards encountered in certain instrumented processes. I purposely use the word “mitigate” rather than “eliminate” because the complete elimination of all risk is an impossibility. Despite our best efforts and intentions, no one can absolutely eliminate all dangers from industrial processes¹. What we can do, though, is *significantly* reduce those risks to the point they begin to approach the low level of “background” risks we all face in daily life, and that is no small achievement.

An important philosophy to follow in the safe design is something called *defense-in-depth*. This is the principle of using multiple layers² of protection, in case one or more of those layers fail. Applying defense-in-depth to process design means regarding each and every safety tool and technique as part of a multi-faceted strategy, rather than as a set of mutually-exclusive alternatives.

To give a brief example of defense-in-depth applied to overpressure protection in a fluid processing system, that system might defend against excessive fluid pressure using all of the following techniques:

- A pressure-control system with an operator-adjusted setpoint
- High-pressure alarms to force operator attention
- A safety shutdown system triggered by abnormally high pressure
- Temperature control systems (both regulatory and safety shutdown) to prevent excessive temperature from helping to create excessive fluid pressure

¹For that matter, it is impossible to eliminate all danger from *life in general*. Every thing you do (or don't do) involves some level of risk. The question really should be, “how much risk is there in a given action, and how much risk am I willing to tolerate?” To illustrate, there does exist a non-zero probability that something you will read in this book is so shocking it will cause you to suffer a heart attack. However, the odds of you walking away from this book and never reading it again over concern of epiphany-induced cardiac arrest are just as slim.

²Also humorously referred to as the “belt *and* suspenders” school of engineering.

- Pressure-relief valves which automatically open to vent high pressure
- Pressure vessels built with “frangible³” tops designed to burst in the safest manner possible
- Locating the process far away from anything (or anyone) that might be harmed by an overpressure event

Any one of these techniques will work to reduce the risk posed by excessive fluid pressure in the system, but all of them used together will provide greater risk reduction than any one used alone.

32.1 Classified areas and electrical safety measures

Any physical location in an industrial facility harboring the potential of explosion due to the presence of flammable process matter suspended in the air is called a *hazardous* or *classified* location. In this context, the label “hazardous” specifically refers to the hazard of explosion, not of other health or safety hazards⁴.

³Frangible roofs are a common design applied to liquid storage tanks harboring the potential for overpressure, such as sulfuric acid storage tanks which may generate accumulations of explosive hydrogen gas. Having the roof seam rupture from overpressure is a far less destructive event than having a side seam or floor seam rupture and consequently spill large volumes of acid. This technique of mitigating overpressure risk does not work to reduce pressure in the system, but it does reduce the risk of damage caused by overpressure in the system.

⁴Chemical corrosiveness, biohazardous substances, poisonous materials, and radiation are all examples of other types of industrial hazards not covered by the label “hazardous” in this context. This is not to understate the danger of these other hazards, but merely to focus our attention on the specific hazard of explosions and how to build instrument systems that will not trigger explosions due to electrical spark.

32.1.1 Classified area taxonomy

In the United States, the National Electrical Code (NEC) published by the National Fire Protection Association (NFPA) defines different categories of “classified” industrial areas and prescribes safe electrical system design practices for those areas. Article 500 of the NEC categorizes classified areas into a system of *Classes* and *Divisions*. Articles 505 and 506⁵ of the NEC provide alternative categorizations for classified areas based on *Zones* that is more closely aligned with European safety standards.

The Class and Division taxonomy defines classified areas in terms of hazard type and hazard probability. Each “Class” contains (or may contain) different types of potentially explosive substances: Class I is for gases or vapors, Class II is for combustible dusts, and Class III is for flammable fibers. The three-fold class designation is roughly scaled on the size of the flammable particles, with Class I being the smallest (gas or vapor molecules) and Class III being the largest (fibers of solid matter). Each “Division” ranks a classified area according to the likelihood of explosive gases, dusts, or fibers being present. Division 1 areas are those where explosive concentrations can or do exist under normal operating conditions. Division 2 areas are those where explosive concentrations only exist infrequently or under abnormal conditions⁶.

The “Zone” method of area classifications defined in Article 505 of the National Electrical Code applies to Class I (explosive gas or vapor) applications, but the three-fold Zone ranks (0, 1, and 2) are analogous to Divisions in their rating of explosive concentration probabilities. Zone 0 defines areas where explosive concentrations are continually present or normally present for long periods of time. Zone 1 defines areas where those concentrations may be present under normal operating conditions, but not as frequently as Zone 0. Zone 2 defines areas where explosive concentrations are unlikely under normal operating conditions, and when present do not exist for substantial periods of time. This three-fold Zone taxonomy may be thought of as expansion on the two-fold Division system, where Zones 0 and 1 are sub-categories of Division 1 areas, and Zone 2 is nearly equivalent to a Division 2 area⁷. A similar three-zone taxonomy for Class II and Class III applications is defined in Article 506 of the National Electrical Code, the zone ranks for these dust and fiber hazards numbered 20, 21, and 22 (and having analogous meanings to zones 0, 1, and 2 for Class I applications).

An example of a classified area common to most peoples’ experience is a vehicle refueling station. Being a (potentially) explosive *vapor*, the hazard in question here is deemed Class I. The Division rating varies with proximity to the fume source. For an upward-discharging vent pipe from an underground gasoline storage tank, the area is rated as Division 1 within 900 millimeters (3 feet) from the vent hole. Between 3 feet and 5 feet away from the vent, the area is rated as Division 2. In relation to an outdoor fuel pump (dispenser), the space internal to the pump enclosure is rated Division 1, and any space up to 18 inches from grade level and up to 20 feet away (horizontally) from the pump is rated Division 2.

⁵Article 506 is a new addition to the NEC as of 2008. Prior to that, the only “zone”-based categories were those specified in Article 505.

⁶The final authority on Class and Division definitions is the National Electrical Code itself. The definitions presented here, especially with regard to Divisions, may not be precise enough for many applications. Article 500 of the NEC is quite specific for each Class and Division combination, and should be referred to for detailed information in any particular application.

⁷Once again, the final authority on this is the National Electrical Code, in this case Article 505. My descriptions of Zones and Divisions are for general information only, and may not be specific or detailed enough for many applications.

Within Class I and Class II (but not Class III), the National Electrical Code further sub-divides hazards according to explosive properties called *Groups*. Each group is defined either according to a substance type, or according to specific ignition criteria. Ignition criteria listed in the National Electrical Code (Article 500) include the *maximum experimental safe gap* (MESG) and the *minimum ignition current ratio* (MICR). The MESG is based on a test where two hollow hemispheres separated by a small gap enclose both an explosive air/fuel mixture and an ignition source. Tests are performed with this apparatus to determine the maximum gap width between the hemispheres that will not permit the excursion of flame from an explosion within the hemispheres triggered by the ignition source. The MICR is the ratio of electrical ignition current for an explosive air/fuel mixture compared to an optimum mixture of methane and air. The smaller of either these two values, the more dangerous the explosive substance is.

Class I substances are grouped according to their respective MESG and MICR values, with typical gas types given for each group:

Group	Typical substance	Safe gap	Ignition current
A	Acetylene		
B	Hydrogen	$\text{MESG} \leq 0.45 \text{ mm}$	$\text{MICR} \leq 0.40$
C	Ethylene	$0.45 \text{ mm} < \text{MESG} \leq 0.75 \text{ mm}$	$0.40 < \text{MICR} \leq 0.80$
D	Propane	$0.75 \text{ mm} < \text{MESG}$	$0.80 < \text{MICR}$

Class II substances are grouped according to material type:


Group	Substances
E	Metal dusts
F	Carbon-based dusts
G	Other dusts (wood, grain, flour, plastic, etc.)

Just to make things confusing, the Class/Zone system described in NEC Article 505 uses a completely different lettering order to describe gas and vapor groups (at the time of this writing there is no grouping of dust or fiber types for the zone system described in Article 506 of the NEC):

Group	Typical substance(s)	Safe gap	Ignition current
IIC	Acetylene, Hydrogen	$\text{MESG} \leq 0.50 \text{ mm}$	$\text{MICR} \leq 0.45$
IIB	Ethylene	$0.50 \text{ mm} < \text{MESG} \leq 0.90 \text{ mm}$	$0.45 < \text{MICR} \leq 0.80$
IIA	Acetone, Propane	$0.90 \text{ mm} < \text{MESG}$	$0.80 < \text{MICR}$

32.1.2 Explosive limits

In order to have combustion (an explosion being a particularly aggressive form of combustion), certain basic criteria must be satisfied: a proper *oxidizer/fuel ratio*, sufficient *energy* for ignition, and the *potential for a self-sustaining chemical reaction* (i.e. the absence of any chemical inhibitors). We may show these criteria in the form of a *fire triangle*⁸, the concept being that removing any of these three critical elements renders a fire (or explosion) impossible:



The fire triangle serves as a qualitative guide for *preventing* fires and explosions, but it does not give sufficient information to tell us if the necessary conditions exist to *support* a fire or explosion. In order for a fire or explosion to occur, we need to have an adequate mixture of fuel and oxidizer in the correct proportions, and a source of ignition energy exceeding a certain minimum threshold.


Suppose we had a laboratory test chamber filled with a mixture of acetone vapor (70% by volume) and air at room temperature, with an electrical spark gap providing convenient ignition. No matter how energetic the spark, this mixture would not explode, because there is too *rich* a mixture of acetone (i.e. too much acetone mixed with not enough air). Every time the spark gap discharges, its energy would surely cause some acetone molecules to combust with available oxygen molecules. However, since the air is so dilute in this rich acetone mixture, those scarce oxygen molecules are depleted fast enough that the flame temperature quickly falls off and is no longer hot enough to trigger the remaining oxygen molecules to combust with the plentiful acetone molecules.

The same problem occurs if the acetone/air mixture is too *lean* (not enough acetone and too much air). This is what would happen if we diluted the acetone vapors to a volumetric concentration of only 0.5% inside the test chamber: any spark at the gap would indeed cause some acetone molecules to combust, but there would be too few available to support expansive combustion across the rest of the chamber.

We could also have an acetone/air mixture in the chamber ideal for combustion (about 9.5% acetone by volume) and still not have an explosion if the spark's energy were insufficient. Most combustion reactions require a certain minimum level of *activation energy* to overcome the potential barrier before molecular bonding between fuel atoms and oxidizer atoms occurs. Stated differently, many combustion reactions are not *spontaneous* at room temperature and at atmospheric pressure – they need a bit of “help” to initiate.

⁸Traditionally, the three elements of a “fire triangle” were fuel, oxidizer, and ignition source. However, this model fails to account for fuels not requiring oxygen as well as cases where a chemical inhibitor prevents a self-sustaining reaction even in the presence of fuel, oxidizer, and ignition source.

All the necessary conditions for an explosion (assuming no chemical inhibitors are present) may be quantified and plotted as an *ignition curve* for any particular fuel and oxidizer combination. This next graph shows an ignition curve for an hypothetical fuel gas mixed with air:



Note how any point in the chart lying *above* the curve is “dangerous,” while any point *below* the curve is “safe.” The three critical values on this graph are the *Lower Explosive Limit* (LEL), the *Upper Explosive Limit* (UEL), and the *Minimum Ignition Energy* (MIE). These critical values differ for every type of fuel and oxidizer combination, change with ambient temperature and pressure, and may be rendered irrelevant in the presence of a catalyst (a chemical substance that works to promote a reaction without itself being consumed by the reaction). Most ignition curves are published with the assumed conditions of air as the oxidizer, at room temperature and atmospheric pressure, with no catalyst(s) present.

Some substances are so reactive that their minimum ignition energy (MIE) levels are well below the thermal energy of ambient air temperatures. Such fuels will *auto-ignite* the moment they come into contact with air, which effectively means one cannot prevent a fire or explosion by eliminating sources of flame or sparks. When dealing with such substances, the only means for preventing fires and explosions lies with maintaining fuel/air ratios outside of the danger zone (i.e. below the LEL or above the UEL), or by using a chemical inhibitor to prevent a self-sustaining reaction.

The greater the difference in LEL and UEL values, the greater “explosive potential” a fuel gas or vapor presents (all other factors being equal), because it means the fuel may explode over a wider range of mixture conditions. It is instructive to research the LEL and UEL values for many common substances, just to see how “explosive” they are relative to each other:

Substance	LEL (% volume)	UEL (% volume)
Acetylene	2.5%	100%
Acetone	2.5%	12.8%
Butane	1.5%	8.5%
Carbon disulfide	1.3%	50%
Carbon monoxide	12.5%	74%
Ether	1.9%	36%
Ethylene oxide	2.6%	100%
Gasoline	1.4%	7.6%
Kerosene	0.7%	5%
Hydrazine	2.9%	98%
Hydrogen	4.0%	74.2%
Methane	4.4%	17%
Propane	2.1%	9.5%

Note how both acetylene and ethylene oxide have UEL values of 100%. This means it is possible for these gases to explode *even when there is no oxidizer present*. Some other chemical substances exhibit this same property (n-propyl nitrate being another example), where the lack of an oxidizer does not prevent an explosion. With these substances in high concentration, our only practical hope of avoiding explosion is to eliminate the possibility of an ignition source in its presence. Some substances have UEL values so high that the elimination of oxidizers is only an uncertain guard against combustion: hydrazine being one example with a UEL of 98%, and diborane being another example with a UEL of 88%.

32.1.3 Protective measures

Different strategies exist to help prevent electrical devices from triggering fires or explosions in classified areas. These strategies may be broadly divided four ways:

- **Contain the explosion:** enclose the device inside a very strong box that contains any explosion generated by the device so as to not trigger a larger explosion outside the box. This strategy may be viewed as eliminating the “ignition” component of the fire triangle, from the perspective of the atmosphere outside the explosion-proof enclosure (ensuring the explosion inside the enclosure does not ignite a larger explosion outside).
- **Shield the device:** enclose the electrical device inside a suitable box or shelter, then purge that enclosure with clean air (or a pure gas) that prevents an explosive mixture from forming inside the enclosure. This strategy works by eliminating the “proper fuel/oxidizer ratio” component of the fire triangle: by eliminating fuel (if purged by air), or by eliminating oxidizer (if purged by fuel gas), or by eliminating both (if purged by an inert gas).
- **Encapsulated design:** manufacture the device so that it is self-enclosing. In other words, build the device in such a way that any spark-producing elements are sealed air-tight within the device from any explosive atmosphere. This strategy works by eliminating the “ignition” component of the fire triangle (from the perspective of outside the device) or by eliminating the “proper fuel/oxidizer ratio” component (from the perspective of inside the device).
- **Limit total circuit energy:** design the circuit such that there is insufficient energy to trigger an explosion, even in the event of an electrical fault. This strategy works by eliminating the “ignition” component of the fire triangle.

It should be noted that any one of these strategies, correctly and thoroughly applied, is sufficient to mitigate the risk of fire and explosion. For this reason you will seldom see more than one of these strategies simultaneously applied (e.g. an explosion-proof enclosure housing a circuit with insufficient energy to trigger an explosion).

A common example of the first strategy is to use extremely rugged metal *explosion-proof* (NEMA 7 or NEMA 8) enclosures instead of the more common sheet-metal or fiberglass enclosures to house electrical equipment. Two photographs of explosion-proof electrical enclosures reveal their unusually rugged construction:




Note the abundance of bolts securing the covers of these enclosures! This is necessary in order to withstand the enormous forces generated by the pressure of an explosion developing inside the enclosure. Note also how most of the bolts have been removed from the door of the right-hand enclosure. This is an unsafe and very unfortunate occurrence at many industrial facilities, where technicians leave just a few bolts securing the cover of an explosion-proof enclosure because it is so time-consuming to remove all of them to gain access inside the enclosure for maintenance work. Such practices negate the safety of the explosion-proof enclosure, rendering it just as dangerous as a sheet metal enclosure in a classified area.

Explosion-proof enclosures are designed in such a way that high-pressure gases resulting from an explosion within the enclosure must pass through small gaps (either holes in vent devices, and/or the gap formed by a bulging door forced away from the enclosure box) en route to exiting the enclosure. As hot gases pass through these tight metal gaps, they are forced to cool to the point where they will not ignite explosive gases outside the enclosure, thus preventing the original explosion inside the enclosure from triggering a far more violent event. This is the same phenomenon measured in determinations of MESH (Maximum Experimental Safe Gap) for an explosive air/fuel mixture. With an explosion-proof enclosure, all gaps are designed to be less than the MESH for the mixtures in question.

A similar strategy involves the use of a non-flammable *purge gas* pressurizing an ordinary electrical enclosure such that explosive atmospheres are prevented from entering the enclosure. Ordinary compressed air may be used as the purge gas, so long as provisions are made to ensure

the air compressor supplying the compressed air is in a non-classified area where explosive gases will never be drawn into the compressed air system.

Devices may be encapsulated in such a way that explosive atmospheres cannot penetrate the device to reach anything generating sufficient spark or heat. *Hermetically sealed* devices are an example of this protective strategy, where the structure of the device has been made completely fluid-tight by fusion joints of its casing. Mercury tilt-switches are good examples of such electrical devices, where a small quantity of liquid mercury is hermetically sealed inside a glass tube. No outside gases, vapors, dusts, or fibers can ever reach the spark generated when the mercury comes into contact (or breaks contact with) the electrodes:




The ultimate method for ensuring instrument circuit safety in classified areas is to intentionally limit the amount of energy available within a circuit such that it *cannot* generate enough heat or spark to ignite an explosive atmosphere, even in the event of an electrical fault within the circuit. Article 504 of the National Electrical Code specifies standards for this method. Any system meeting these requirements is called an *intrinsically safe* or *I.S.* system. The word “intrinsic” implies that the safety is a natural property of the circuit, since it lacks even the ability to produce an explosion-triggering spark⁹.

One way to underscore the meaning of intrinsic safety is to contrast it against a different concept that has the appearance of similarity. Article 500 of the National Electrical Code defines *nonincendive equipment* as devices incapable of igniting a hazardous atmosphere *under normal operating conditions*. However, the standard for nonincendive devices or circuits does not guarantee what will happen under *abnormal* conditions, such as an open- or short-circuit in the wiring. So, a “nonincendive” circuit may very well pose an explosion hazard, whereas an “intrinsically safe” circuit will not because the intrinsically safe circuit simply does not possess enough energy to trigger an explosion under *any* electrical fault condition. As a result, nonincendive circuits are not approved in Class I or Class II Division 1 locations whereas intrinsically safe circuits are approved for all hazardous locations.

⁹To illustrate this concept in a different context, consider my own personal history of automobiles. For many years I drove an ugly and inexpensive truck which I joked had “intrinsic theft protection:” it was so ugly, no one would ever want to steal it. Due to this “intrinsic” property of my vehicle, I had no need to invest in an alarm system or any other protective measure to deter theft. Similarly, the components of an intrinsically safe system need not be located in explosion-proof or purged enclosures because the intrinsic energy limitation of the system is protection enough.

Most modern 4 to 20 mA analog signal instruments may be used as part of intrinsically safe circuits so long as they are connected to control equipment through suitable *safety barrier* interfaces, the purpose of which is to limit the amount of voltage and current available at the field device to low enough levels that an explosion-triggering spark is impossible even under fault conditions (e.g. a short-circuit in the field instrument or wiring). A simple intrinsic safety barrier circuit made from passive components is shown in the following diagram¹⁰:



In normal operation, the 4-20 mA field instrument possesses insufficient terminal voltage and insufficient loop current to pose any threat of hazardous atmosphere ignition. The series resistance of the barrier circuit is low enough that the 4-20 mA signal will be unaffected by its presence. As far as the receiving instrument (indicator or controller) is “concerned,” the safety barrier might as well not exist.


If a short-circuit develops in the field instrument, the series resistance of the barrier circuit will limit fault current to a value low enough not to pose a threat in the hazardous area. If something fails in the receiving instrument to cause a much greater power supply voltage to develop at its terminals, the zener diode inside the barrier will break down and provide a shunt path for fault current that bypasses the field instrument (and may possibly blow the fuse in the barrier). Thus, the intrinsic safety barrier circuit provides protection against overcurrent *and* overvoltage faults, so that neither type of fault will result in enough electrical energy available at the field device to ignite an explosive atmosphere.

Note that a barrier device such as this *must* be present in the 4-20 mA analog circuit in order for the circuit to be intrinsically safe. The “intrinsic” safety rating of the circuit depends on this barrier, not on the integrity of the field device or of the receiving device. Without this barrier in place, the instrument circuit is not intrinsically safe, even though the *normal* operating voltage and current parameters of the field and receiving devices are well within the parameters of safety for

¹⁰Real passive barriers often used redundant zener diodes connected in parallel to ensure protection against excessive voltage even in the event of a zener diode failing open.

classified areas. It is the barrier and the barrier alone which guarantees those voltage and current levels will remain within safe limits in the event of *abnormal* circuit conditions such as a field wiring short or a faulty loop power supply.

More sophisticated *active* barrier devices are manufactured which provide electrical isolation from ground in the instrument wiring, thus eliminating the need for a safety ground connection at the barrier device.



In the example shown here, transformers¹¹ are used to electrically isolate the analog current signal so that there is no path for DC fault current between the field instrument and the receiving instrument, ground or no ground.

Safety barrier circuits fundamentally limit the amount of power deliverable to a field device from a power supply located in the safe area. Barrier circuits cannot, however, ensure safety for field devices capable of *generating* their own electrical energy. In order for such devices to be considered intrinsically safe, their natural abilities for generating voltage, current, and power must fall below limits defined in NEC article 504. Sensors such as pH electrodes, thermocouples, and photovoltaic light detectors are examples of such field devices, and are called *simple apparatus* by the NEC. The qualifications for a generating device to be a “simple apparatus” is that it cannot generate more than 1.5 volts of voltage, and more than 100 milliamps of current, and more than 25 milliwatts of power. If a device’s ability to generate electricity exceeds these limits, the device is not a “simple apparatus” and therefore its circuit is not intrinsically safe.

An example of a generating field device exceeding these limits is a tachogenerator: a small DC generator used to measure the speed of rotating equipment by outputting a DC voltage proportional to speed (typically over a 0-10 volt range). An alternative to a tachogenerator for measuring machine

¹¹Of course, transformers cannot be used to pass DC signals of any kind, which is why chopper/converter circuits are used before and after the signal transformer to convert each DC current signal into a form of chopped (AC) signal that *can* be fed through the transformer. This way, the *information* carried by each 4-20 mA DC current signal passes through the barrier, but electrical fault current cannot.

speed is an *optical encoder*, using a slotted wheel to chop a light beam (from an LED), generating a pulsed electrical signal of sufficiently low intensity to qualify as a simple apparatus.

Passive (non-generating) field devices may also be classified as “simple apparatus” if they do not dissipate more than 1.3 watts of power. Examples of passive, simple apparatus include switches, LED indicator lamps, and RTD (Resistive Temperature Detector) sensors. Even devices with internal inductance and/or capacitance may be deemed “simple apparatus” if their stored energy capacity is insufficient to pose a hazard.

In addition to the use of barrier devices to create an intrinsically safe circuit, the National Electrical Code (NEC) article 504 specifies certain wiring practices different from normal control circuits. The conductors of an intrinsically safe circuit (i.e. conductors on the “field” side of a barrier) must be separated from the conductors of the non-intrinsically safe circuit (i.e. conductors on the “supply” side of the barrier) by at least 50 millimeters, which is approximately 2 inches. Conductors must be secured prior to terminals in such a way that they cannot come into contact with non-intrinsically safe conductors if the terminal becomes loose. Also, the color *light blue* may be used to identify intrinsically safe conductors, raceways, cable trays, and junction boxes so long as that color is not used for any other wiring in the system.

32.2 Concepts of probability

While the term “probability” may evoke images of imprecision, probability is in fact an exact mathematical concept: *the ratio a specific outcome to total possible outcomes* where 1 (100%) represents certainty and 0 (0%) represents impossibility. A probability value between 1 and 0 describes an outcome that occurs some of the time but not all of the time. Reliability – which is the expression of how likely a device or a system is to function as intended – is based on the mathematics of probability. Therefore, a rudimentary understanding of probability mathematics is necessary to grasp what reliability means.

Before we delve too deeply into discussions of reliability, some definition of terms is in order. We have defined “reliability” to mean the probability of a device or system functioning as designed, which is a good general definition but sometimes not specific enough for our needs. There are usually a variety of different ways in which a device or system can fail, and these different failure modes usually have different probability values. Let’s take for example a fire alarm system triggered by a manual pushbutton switch: the intended function of such a system is to activate an alarm whenever the switch is pressed. If we wish to express the reliability of this system, we must first carefully define what we mean by “failure”. One way in which this simple fire alarm system could fail is if it remained silent when the pushbutton switch was pressed (i.e. *not* alerting people when it should have). Another, completely different, way this simple system could fail is by accidentally sounding the alarm when no one pressed the switch (i.e. alerting people when it had no reason to, otherwise known as a “false alarm”). If we discuss the “reliability” of this fire alarm system, we may need to differentiate between these two different kinds of unreliable system behaviors.

The electrical power industry has an interest in ensuring the safe delivery of electrical power to loads, both ensuring maximum service to customers while simultaneously shutting power off as quickly as possible in the event of dangerous system faults. A complex system of fuses, circuit breakers, and protective relays work together to ensure the flow of power remains uninterrupted as long as safely possible. These protective devices must shut off power when they sense dangerous conditions, but they must also refrain from needlessly shutting off power when there is no danger. Like our fire alarm system which must alert people when needed yet not sound false alarms, electrical protective systems serve two different needs. In order to avoid confusion when quantifying the reliability of electrical protective systems to function as designed, the power industry consistently uses the following terms:

- **Dependability:** The probability a protective system will shut off power when needed
- **Security:** The probability a protective system will allow power to flow when there is no danger
- **Reliability:** A combination of dependability and security

For the sake of clarity I will use these same terms when discussing the reliability of any instrument or control systems. “Dependability” is how reliably a device or system will take appropriate action when it is actively called to do so – in other words, the degree to which we may *depend* on this device or system to do its job when activated. “Security” is how reliably a device or system refrains from taking action when no action should be taken – in other words, the degree to which we may feel *secure* it won’t needlessly trigger a system shutdown or generate a false alarm. If there is no need to differentiate, the term “reliability” will be used as a general description of how probable a device or system will do what it was designed to do.

The following matrix should help clarify the meanings of these three terms, defined in terms of what the protective component or system does under various conditions:

	Reliable	Unreliable
Ordinary condition (context of <i>security</i>)	No action taken (secure) S	Shut-down (trip)! (unsecure) \bar{S}
Emergency condition (context of <i>dependability</i>)	Shut-down (trip)! (dependable) D	No action taken (undependable) \bar{D}

In summary: a protective function that does not trip when it doesn't need to is *secure*; a protective function that trips when it needs to is *dependable*; a protective system that does both is *reliable*.

The Boolean variables used to symbolize dependability (D), security (S), undependability (\bar{D}), and unsecurity (\bar{S}) tell us something about the relationships between those four quantities. A bar appearing over a Boolean variable represents the *complement* of that variable. For example, security (S) and unsecurity (\bar{S}) are complementary to each other: if we happen to know the probability that a device or system will be secure, then we may calculate with assurance the probability that it is unsecure. A fire alarm system that is 99.3% secure (i.e. 99.3% of the time it generates no false alarms) must generate false alarms the other 0.7% of the time in order to account for all possible system responses 100% of the time no fires occur. If that same fire alarm system is 99.8% dependable (i.e. it alerts people to the presence of a real fire 99.8% of the time), then we may conclude it will fail to report 0.02% of real fire incidents in order to account for all possible responses during 100% of fire incidents.

However, it should be clearly understood that there is no such simple relationship between security (S) and dependability (D) because these two measures refer to completely different conditions and (potentially) different modes of failure. The specific faults causing a fire alarm system to generate a false alarm (an example of an *unsecure* outcome, \bar{S}) are quite different from the faults disabling that same fire alarm system in the event of a real fire (an example of an *undependable* outcome, \bar{D}). Through the application of redundant components and clever system design we may augment dependability and/or security (sometimes improving one at the expense of the other), but it should be understood that these are really two fundamentally different probability measures and as such are not necessarily related.

32.2.1 Mathematical probability

Probability may be defined as a ratio of specific outcomes to total (possible) outcomes. If you were to flip a coin, there are really only two possibilities¹² for how that coin may land: face-up (“heads”) or face-down (“tails”). The probability of a coin falling “tails” is thus one-half ($\frac{1}{2}$), since “tails” is one specific outcome out of two total possibilities. Calculating the probability (P) is a matter of setting up a ratio of outcomes:

$$P(\text{“tails”}) = \frac{\text{“tails”}}{\text{“heads”} + \text{“tails”}} = \frac{1}{2} = 0.5$$


This may be shown graphically by displaying all possible outcomes for the coin’s landing (“heads” or “tails”), with the one specific outcome we’re interested in (“tails”) highlighted for emphasis:



The probability of the coin landing “heads” is of course exactly the same, because “heads” is also *one* specific outcome out of *two* total possibilities.

If we were to roll a six-sided die, the probability of that die landing on any particular side (let’s arbitrarily choose the “four” side) is one out of six, because we’re looking at one specific outcome out of six total possibilities:


$$P(\text{“four”}) = \frac{\text{“four”}}{\text{“one”} + \text{“two”} + \text{“three”} + \text{“four”} + \text{“five”} + \text{“six”}} = \frac{1}{6} = 0.1\overline{6}$$



¹²To be honest, the coin could also land on its edge, which is a third possibility. However, that third possibility is so remote as to be negligible in the presence of the other two. Strictly speaking, $P(\text{“heads”}) + P(\text{“tails”}) + P(\text{“edge”}) = 1$.

If we were to roll the same six-sided die, the probability of that die landing on an even-numbered side (2, 4, or 6) is three out of six, because we're looking at three specific outcomes out of six total possibilities:

$$P(\text{even}) = \frac{\text{"two"} + \text{"four"} + \text{"six"}}{\text{"one"} + \text{"two"} + \text{"three"} + \text{"four"} + \text{"five"} + \text{"six"}} = \frac{3}{6} = 0.5$$



As a ratio of specific outcomes to total possible outcomes, the probability of any event will always be a number ranging in value from 0 to 1, inclusive. This value may be expressed as a fraction ($\frac{1}{2}$), as a per unit value (0.5), as a percentage (50%), or as a verbal statement (e.g. “three out of six”). A probability value of zero (0) means a specific event is impossible, while a probability of one (1) means a specific event is guaranteed to occur.

Probability values realistically apply only to large samples. A coin tossed ten times may very well fail to land “heads” exactly five times and land “tails” exactly five times. For that matter, it may fail to land on each side exactly 500000 times out of a million tosses. However, so long as the coin and the coin-tossing method are *fair* (i.e. not biased in any way), the experimental results will approach¹³ the ideal probability value as the number of trials approaches infinity. Ideal probability values become less and less certain as the number of trials decreases, and are completely useless for singular (non-repeating) events.

A familiar application of probability values is the forecasting of meteorological events such as rainfall. When a weather forecast service provides a rainfall prediction of 65% for a particular day, it means that out of a large number of days sampled in the past having similar measured conditions (cloud cover, barometric pressure, temperature and dew point, etc.), 65% of those days experienced rainfall. This past history gives us some idea of how likely rainfall will be for any present situation, based on similarity of measured conditions.

Like all probability values, forecasts of rainfall are more meaningful with greater samples. If we wish to know how many days with measured conditions similar to those of the forecast day will experience rainfall over the *next ten years*, the forecast probability value of 65% will be quite accurate. However, if we wish to know whether or not rain will fall on any particular (single) day having those same conditions, the value of 65% tells us very little. So it is with all measurements of probability: precise for large samples, ambiguous for small samples, and virtually meaningless for singular conditions¹⁴.

In the field of instrumentation – and more specifically the field of *safety* instrumented systems – probability is useful for the mitigation of hazards based on equipment failures where the probability of failure for specific pieces of equipment is known from mass production of that equipment and years of data gathered describing the reliability of the equipment. If we have data showing the probabilities

¹³In his excellent book, *Reliability Theory and Practice*, Igor Bazovsky describes the relationship between true probability (P) calculated from ideal values and estimated probability (\hat{P}) calculated from experimental trials as a limit function: $P = \lim_{N \rightarrow \infty} \hat{P}$, where N is the number of trials.

¹⁴Most people can recall instances where a weather forecast proved to be completely false: a prediction for rainfall resulting in a completely dry day, or vice-versa. In such cases, one is tempted to blame the weather service for poor forecasting, but in reality it has more to do with the nature of probability, specifically the meaninglessness of probability calculations in predicting singular events.

of failure for different pieces of equipment, we may use this data to calculate the probability of failure for the system as a whole. Furthermore, we may apply certain mathematical laws of probability to calculate system reliability for different equipment configurations, and therefore minimize the probability of system failure by optimizing those configurations.

As with weather predictions, predictions of system reliability (or conversely, of system failure) become more accurate as the sample size grows larger. Given an accurate probabilistic model of system reliability, a system (or a set of systems) with enough individual components, and a sufficiently long time-frame, an organization may accurately predict the number of system failures and the cost of those failures (or alternatively, the cost of minimizing those failures through preventive maintenance). However, no probabilistic model can accurately predict which component in a large system will fail at any specific point in time.


The ultimate purpose, then, in probability calculations for process systems and automation is to optimize the safety and availability of large systems over many years of time. Calculations of reliability, while useful to the technician in understanding the nature of system failures and how to minimize them, are actually more valuable (more meaningful) at the enterprise level.

32.2.2 Laws of probability

Probability mathematics bears an interesting similarity to Boolean algebra in that probability values (like Boolean values) range between zero (0) and one (1). The difference, of course, is that while Boolean variables may *only* have values equal to zero or one, probability variables range continuously between those limits. Given this similarity, we may apply standard Boolean operations such as NOT, AND, and OR to probabilities. These Boolean operations lead us to our first “laws” of probability for combination events.

The logical “NOT” function

For instance, if we know the probability of rolling a “four” on a six-sided die is $\frac{1}{6}$, then we may safely say the probability of *not* rolling a “four” is $\frac{5}{6}$, the complement of $\frac{1}{6}$. The common “inverter” logic symbol is shown here representing the complementation function, turning a probability of rolling a “four” into the probability of *not* rolling a “four”:



Symbolically, we may express this as a sum of probabilities equal to one:

$$P(\text{total}) = P(\text{“one”}) + P(\text{“two”}) + P(\text{“three”}) + P(\text{“four”}) + P(\text{“five”}) + P(\text{“six”}) = 1$$

$$P(\text{total}) = \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} + \frac{1}{6} = 1$$

$$P(\text{total}) = P(\text{“four”}) + P(\text{not “four”}) = \frac{1}{6} + \frac{5}{6} = 1$$

$$P(\text{“four”}) = 1 - P(\text{not “four”}) = 1 - \frac{5}{6} = \frac{1}{6}$$


We may state this as a general “law” of complementation for any event (A):

$$P(A) = 1 - \overline{P}(A)$$

Complements of probability values find frequent use in reliability engineering. If we know the probability value for the failure of a component (i.e. how likely it is to fail when called upon to function – termed the *Probability of Failure on Demand*, or *PF_D* – which is a measure of that component’s *undependability*), then we know the *dependability* value (i.e. how likely it is to function on demand) will be the mathematical complement. To illustrate, consider a device with a PFD value of $\frac{1}{100000}$. Such a device could be said to have a dependability value of $\frac{99999}{100000}$, or 99.999%, since $1 - \frac{1}{100000} = \frac{99999}{100000}$.

The logical “AND” function

The AND function regards probabilities of two or more coincidental events (i.e. where the outcome of interest only happens if two or more events happen together, or in a specific sequence). Another example using a die is the probability of rolling a “four” on the first toss, then rolling a “one” on the second toss. It should be intuitively obvious that the probability of rolling this specific combination of values will be less (i.e. less likely) than rolling either of those values in a single toss. The shaded field of possibilities (36 in all) demonstrate the unlikelihood of this sequential combination of values compared to the unlikelihood of either value on either toss:



As you can see, there is only one outcome matching the specific criteria out of 36 total possible outcomes. This yields a probability value of one-in-thirty six ($\frac{1}{36}$) for the specified combination, which is the *product* of the individual probabilities. This, then, is our second law of probability:

$$P(A \text{ and } B) = P(A) \times P(B)$$

A practical application of this would be the calculation of failure probability for a double-block valve assembly, designed to positively stop the flow of a dangerous process fluid. Double-block valves are used to provide increased assurance of shut-off, since the shutting of *either* block valve is sufficient in itself to stop fluid flow. The probability of failure for a double-block valve assembly – “failure” defined as not being able to stop fluid flow when needed – is the product of each valve’s un-dependability (i.e. probability of failing in the open position when commanded to shut off):



With these two valves in service, the probability of neither valve successfully shutting off flow (i.e. *both* valve 1 *and* valve 2 failing on demand; remaining open when they should shut) is the product of their individual failure probabilities.

$$P(\text{assembly fail}) = P(\text{valve 1 fail open}) \times P(\text{valve 2 fail open})$$

$$P(\text{assembly fail}) = 0.0002 \times 0.0003$$

$$P(\text{assembly fail}) = 0.00000006 = 6 \times 10^{-8}$$

An extremely important assumption in performing such an AND calculation is that the probabilities of failure for each valve are completely unrelated. For instance, if the failure probabilities of both valve 1 and valve 2 were largely based on the possibility of a certain residue accumulating inside the valve mechanism (causing the mechanism to freeze in the open position), and *both* valves were equally susceptible to this residue accumulation, there would be virtually no advantage to having double block valves. If said residue were to accumulate in the piping, it would affect both valves practically the same. Thus, the failure of one valve due to this effect would virtually ensure the failure of the other valve as well. *The probability of simultaneous or sequential events being the product of the individual events’ probabilities is true if and only if the events in question are completely independent.*

We may illustrate the same caveat with the sequential rolling of a die. Our previous calculation showed the probability of rolling a “four” on the first toss and a “one” on the second toss to be $\frac{1}{6} \times \frac{1}{6}$, or $\frac{1}{36}$. However, if the person throwing the die is extremely consistent in their throwing technique and the way they orient the die after each throw, such that rolling a “four” on one toss makes it very likely to roll a “one” on the next toss, the sequential events of a “four” followed by a “one” would be far more likely than if the two events were completely random and independent. The probability calculation of $\frac{1}{6} \times \frac{1}{6} = \frac{1}{36}$ holds true only if all the throws’ results are completely unrelated to each other.

Another, similar application of the Boolean AND function to probability is the calculation of system reliability (R) based on the individual reliability values of components necessary for the system's function. If we know the reliability values for several essential¹⁵ system components, and we also know those reliability values are based on independent (unrelated) failure modes, the overall system reliability will be the product (Boolean AND) of those component reliabilities. This mathematical expression is known as *Lusser's product law of reliabilities*:

$$R_{system} = R_1 \times R_2 \times R_3 \times \cdots \times R_n$$

As simple as this law is, it is surprisingly unintuitive. Lusser's Law tells us that any system depending on the performance of several essential components will be *less* reliable than the least-reliable of those components. This is akin to saying that a chain will be *weaker* than its weakest link!

To give an illustrative example, suppose a complex system depended on the reliable operation of six key components in order to function, with the individual reliabilities of those six components being 91%, 92%, 96%, 95%, 93%, and 92%, respectively. Given individual component reliabilities all greater than 90%, one might be inclined to think the overall reliability would be quite good. However, following Lusser's Law we find the reliability of this system (as a whole) is only 65.3% because $0.91 \times 0.92 \times 0.96 \times 0.95 \times 0.93 \times 0.92 = 0.653$.


In his excellent text *Reliability Theory and Practice*, author Igor Bazovsky recounts the German V1 missile project during World War Two, and how early assumptions of system reliability were grossly inaccurate¹⁶. Once these faulty assumptions of reliability were corrected, development of the V1 missile resulted in greatly increased reliability until a system reliability of 75% (three out of four) was achieved.

¹⁵Here, "essential" means the system will fail if *any* of these identified components fails. Thus, Lusser's Law implies a logical "AND" relationship between the components' reliability values and the overall system reliability.

¹⁶According to Bazovsky (pp. 275-276), the first reliability principle adopted by the design team was that the system could be no more reliable than its least-reliable (weakest) component. While this is technically true, the mistake was to assume that the system would be *as reliable* as its weakest component (i.e. the "chain" would be exactly as strong as its weakest link). This proved to be too optimistic, as the system would still fail due to the failure of "stronger" components even when the "weaker" components happened to survive. After noting the influence of "stronger" components' unreliabilities on overall system reliability, engineers somehow reached the bizarre conclusion that system reliability was equal to the mathematical *average* of the components' reliabilities. Not surprisingly, this proved even less accurate than the "weakest link" principle. Finally, the designers were assisted by the mathematician Erich Pieruschka, who helped formulate Lusser's Law.

The logical “OR” function

The OR function regards probabilities of two or more redundant events (i.e. where the outcome of interest happens if any one of the events happen). Another example using a die is the probability of rolling a “four” on the first toss *or* on the second toss. It should be intuitively obvious that the probability of rolling a “four” on either toss will be more probable (i.e. more likely) than rolling a “four” on a single toss. The shaded field of possibilities (36 in all) demonstrate the likelihood of this either/or result compared to the likelihood of either value on either toss:




As you can see, there are eleven outcomes matching the specific criteria out of 36 total possible outcomes (the outcome with two “four” rolls counts as a single trial matching the stated criteria, just as all the other trials containing only one “four” roll count as single trials). This yields a probability value of eleven-in-thirty six ($\frac{11}{36}$) for the specified combination. This result may defy your intuition, if you assumed the OR function would be the simple *sum* of individual probabilities ($\frac{1}{6} + \frac{1}{6} = \frac{2}{6}$ or $\frac{1}{3}$), as opposed to the AND function’s *product* of probabilities ($\frac{1}{6} \times \frac{1}{6} = \frac{1}{36}$). In truth, there is an application of the OR function where the probability is the simple sum, but that will come later in this presentation.


As with the logical “AND” function, the logical “OR” function assumes the events in question are independent from each other. That is to say, the events lack a common cause, and are not contingent upon one another in any way.

For now, a way to understand why we get a probability value of $\frac{11}{36}$ for our OR function with two $\frac{1}{6}$ input probabilities is to derive the OR function from other functions whose probability laws we already know with certainty. From Boolean algebra, DeMorgan's Theorem tells us an OR function is equivalent to an AND function with all inputs and outputs inverted ($A + B = \overline{\overline{A} \overline{B}}$):

(Equivalent logic functions)



We already know the complement (inversion) of a probability is the value of that probability subtracted from one ($\overline{P} = 1 - P$). This gives us a way to symbolically express the DeMorgan's Theorem definition of an OR function in terms of an AND function with three inversions:



Knowing that $\overline{P}(A) = 1 - P(A)$ and $\overline{P}(B) = 1 - P(B)$, we may substitute these inversions into the triple-inverted AND function to arrive at an expression for the OR function in simple terms of $P(A)$ and $P(B)$:

$$P(A \text{ or } B) = \overline{\overline{P}(A) \times \overline{P}(B)}$$

$$P(A \text{ or } B) = \overline{(1 - P(A))(1 - P(B))}$$

$$P(A \text{ or } B) = 1 - [(1 - P(A))(1 - P(B))]$$

Distributing terms on the right side of the equation:

$$P(A \text{ or } B) = 1 - [1 - P(B) - P(A) + P(A)P(B)]$$

$$P(A \text{ or } B) = P(B) + P(A) - P(A)P(B)$$

This, then, is our third law of probability:

$$P(A \text{ or } B) = P(B) + P(A) - P(A) \times P(B)$$

Inserting our example probabilities of $\frac{1}{6}$ for both $P(A)$ and $P(B)$, we obtain the following probability for the OR function:

$$P(A \text{ or } B) = \frac{1}{6} + \frac{1}{6} - \left(\frac{1}{6}\right) \left(\frac{1}{6}\right)$$


$$P(A \text{ or } B) = \frac{2}{6} - \left(\frac{1}{36}\right)$$

$$P(A \text{ or } B) = \frac{12}{36} - \frac{1}{36}$$

$$P(A \text{ or } B) = \frac{11}{36}$$

This confirms our previous conclusion of there being an $\frac{11}{36}$ probability of rolling a “four” on the first or second rolls of a die.

We may return to our example of a double-block valve assembly for a practical application of OR probability. When illustrating the AND probability function, we focused on the probability of both block valves failing to shut off when needed, since both valve 1 *and* valve 2 would have to fail open in order for the double-block assembly to fail in shutting off flow. Now, we will focus on the probability of *either* block valve failing to open when needed. While the AND scenario was an exploration of the system’s un-dependability (i.e. the probability it might fail to stop a dangerous condition), this scenario is an exploration of the system’s *un-security* (i.e. the probability it might fail to resume normal operation).




Each block valve is designed to be able to shut off flow independently, so that the flow of (potentially) dangerous process fluid will be halted if *either or both* valves shut off. The probability that process fluid flow may be impeded by the failure of either valve to open is thus a simple (non-exclusive) OR function:

$$P(\text{assembly fail}) = P(\text{valve 1 fail shut}) + P(\text{valve 2 fail shut}) - P(\text{valve 1 fail shut}) \times P(\text{valve 2 fail shut})$$

$$P(\text{assembly fail}) = 0.0003 + 0.0001 - (0.0003 \times 0.0001)$$

$$P(\text{assembly fail}) = 0.0003997 = 3.9997 \times 10^{-4}$$


A similar application of the OR function is seen when we are dealing with *exclusive* events. For instance, we could calculate the probability of rolling either a “three” or a “four” in a single toss of a die. Unlike the previous example where we had two opportunities to roll a “four,” and two sequential rolls of “four” counted as a single successful trial, here we know with certainty that the die cannot land on “three” *and* “four” in the same roll. Therefore, the exclusive OR probability (XOR) is much simpler to determine than a regular OR function:



This is the only type of scenario where the function probability is the simple sum of the input probabilities. In cases where the input probabilities are mutually exclusive (i.e. they *cannot* occur simultaneously or in a specific sequence), the probability of one *or* the other happening is the sum of the individual probabilities. This leads us to our fourth probability law:

$$P(A \text{ exclusively or } B) = P(A) + P(B)$$

A practical example of the exclusive-or (XOR) probability function may be found in the failure analysis of a single block valve. If we consider the probability this valve may fail in either condition (stuck open or stuck shut), and we have data on the probabilities of the valve failing open and failing shut, we may use the XOR function to model the system's general unreliability¹⁷. We know that the exclusive-or function is the appropriate one to use here because the two "input" scenarios (failing open versus failing shut) *absolutely cannot* occur at the same time:



$$P(\text{valve fail}) = P(\text{valve fail open}) + P(\text{valve fail shut})$$

$$P(\text{valve fail}) = 0.0002 + 0.0003$$

$$P(\text{valve fail}) = 0.0005 = 5 \times 10^{-4}$$

If the intended safety function of this block valve is to shut off the flow of fluid if a dangerous condition is detected, then the probability of this valve's failure to shut when needed is a measure of its *undependability*. Conversely, the probability of this valve's failure to open under normal (safe) operating conditions is a measure of its *unsecurity*. The XOR of the valve's undependability and its insecurity therefore represents its *unreliability*. The complement of this value will be the valve's *reliability*: $1 - 0.0005 = 0.9995$. This reliability value tells us we can expect the valve to operate as it's called to 99.95% of the time, and we should expect 5 failures out of every 10,000 calls for action.

¹⁷Here we have an example where dependability and security are lumped together into one "reliability" quantity.

Summary of probability laws

The complement (inversion) of a probability:

$$P(A) = 1 - \overline{P}(A)$$

The probability of coincidental events (where both must happen either simultaneously or in specific sequence) for the result of interest to occur:

$$P(A \text{ and } B) = P(A) \times P(B)$$

The probability of redundant events (where either or both may happen) for the result of interest to occur:

$$P(A \text{ or } B) = P(B) + P(A) - P(A) \times P(B)$$

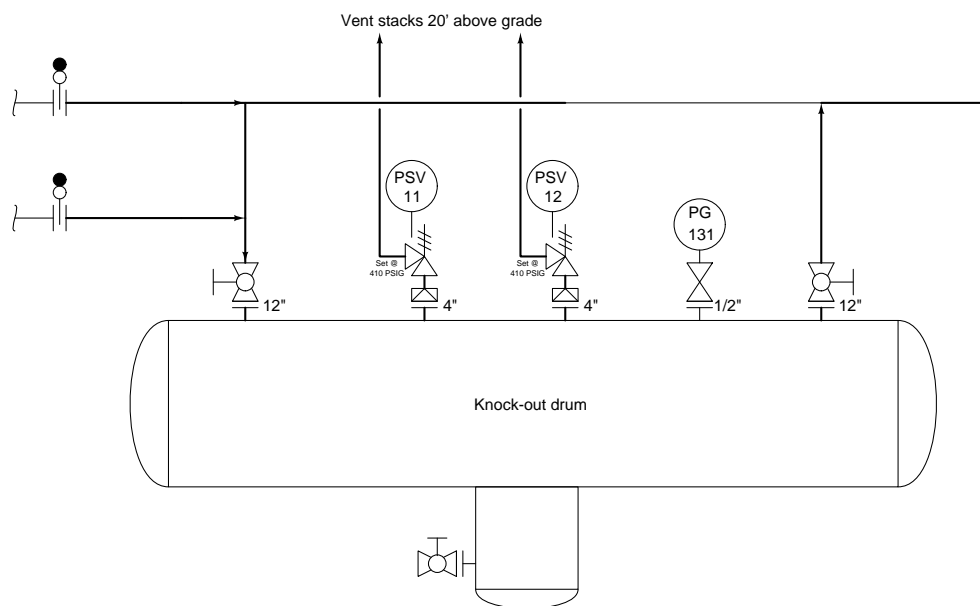
The probability of exclusively redundant events (where either may happen, but not simultaneously or in specific sequence) for the result of interest to occur:

$$P(A \text{ exclusively or } B \text{ exclusively}) = P(A) + P(B)$$

32.2.3 Applying probability laws to real systems

The relatively simple concepts of AND and OR Boolean functions become surprisingly complicated when applying them to real-life measures of component reliability, mainly because reliability is measured in multiple ways. As we have already seen, *dependability* (D) and *security* (S) are related concepts in that they both describe the probability of a system or system component functioning properly, but defy simple correlation because they imply different failure modes. “Dependability” for any safety-related system or component is the probability that it will perform its safety function when called upon in an emergency. “Security” by contrast is the probability that the system or component in question will maintain normal operation when there is no emergency.

To illustrate, we will examine the overpressure protection features of a “knock-out drum” used to collect small amounts of liquid entrained in a gas stream. This particular vessel is equipped with two pressure-safety valves (PSV-11 and PSV-12) designed to open and vent gas to atmosphere in the event of an overpressure condition (over 410 PSIG):



Suppose each of these pressure safety valves has a dependability (D) rating of 0.9992, which means each one has a 99.92% chance of opening up to relieve excess pressure when a high-pressure condition exists. Let us also suppose each of these PSVs has a security (S)¹⁸ rating of 0.995, which means each one has a 99.5% chance of remaining in the shut condition when no overpressure condition exists. Furthermore, assume each of the two pressure safety valves individually has a high enough flow capacity to adequately vent the vessel during an overpressure condition.

¹⁸An easy way to remember what each of these terms mean in the context of a protective system is to associate D (Dependability) with a *dangerous* scenario and S (Security) with a *safe* scenario: D expresses what the system or component will do when a dangerous condition presents itself to the protective system and it needs to act; S expresses what the system or component will do when conditions are safe and there is no need to act.

How might we calculate the overall dependability and security ratings of this dual-PSV overpressure protection system? Clearly, we must use Boolean functions to combine the two valves' D ratings into a D_{system} rating, and likewise with the two valves' S ratings, but which logical function should we use to calculate each measure of reliability? The choice between AND and OR functions may not be obvious at first inspection.

One way to analyze logical functions is in terms of what state (0 or 1) at any input will *guarantee* a certain output state. For an AND function, any 0 state in guarantees a 0 state out. For an OR function, any 1 state in guarantees a 1 state out. These facts are useful when selecting logical functions for a variety of purposes, and they will serve us well in this application of probability values too.

A useful problem-solving technique for this application is called *limiting cases*, where we take some quantity to its extreme limits in an effort to simplify the problem at hand. To begin, we will assume that one of the two pressure safety valves in this system has a D rating of 1, which means it is perfectly reliable when called to open by a high-pressure condition. A D rating of 1 is a “limiting case” of the pressure safety valve’s dependability: a perfectly dependable PSV. If this were true, would it guarantee the whole overpressure protection system is dependable, or not? Since we know each valve is sized large enough to protect the vessel on its own (without need of the second PSV opening), then the answer to this question is “yes”: a single PSV with a D rating of 1 guarantees a D_{system} rating of 1. All we need is for one of these PSVs to vent when it senses a high-pressure condition to protect the vessel from overpressure damage. Therefore, the proper Boolean function to calculate D_{system} from the valves’ individual D ratings is the OR function, because given the choices of AND and OR only the OR function guarantees a certain output state with any “1” input. Calculating system dependability using both valves’ D ratings:

$$\begin{array}{l}
 D_{PSV11} = 0.9992 \\
 D_{PSV12} = 0.9992
 \end{array}
 \begin{array}{c}
 \text{---} \\
 \text{---}
 \end{array}
 \begin{array}{c}
 \diagup \\
 \diagdown
 \end{array}
 \begin{array}{c}
 \text{---} \\
 \text{---}
 \end{array}
 \begin{array}{l}
 D_{system} = 0.99999936
 \end{array}$$

The numerical results shown here should make sense: in an overpressure protection system where we only need one of the two valves to vent gas during an overpressure condition, having two valves increases the probability that the vessel will be adequately protected.

Now we will apply this same problem-solving strategy to the system’s *security* (S). Taking the high limiting-case value of either PSV’s S rating, we ask ourselves the question “Does any one perfectly secure PSV ($S = 1$) make the system secure?” In other words, if one of these valves was guaranteed not to vent when no overpressure condition exists, would that mean the entire system was guaranteed not to vent when it didn’t need to? The answer here is “no”, since the presence of *two* pressure safety valves increases the chance of unnecessary leakage. This tells us we cannot use the OR function for security, because a perfectly secure PSV ($S = 1$) does *not* guarantee a perfectly secure system.

At this point we may conclude that the proper Boolean function for system security in this application is the AND, by process of elimination. However, we may also consider a different limiting-case scenario to verify this conclusion. Let us suppose one of the pressure safety valves failed in

such a way that it had *zero* security, meaning there was no chance at all it would remain shut when no overpressure condition existed (i.e. a security rating of $S = 0$ means it is guaranteed to vent when it shouldn't). Would one PSV in this state guarantee a certain system security state? We see here that this is true: any one PSV with an S rating of zero means the system as a whole has a zero S rating as well, because all it takes is one PSV to unnecessarily vent to make the system as a whole unnecessarily vent. Since we know the Boolean AND function guarantees a zero output for any zero input, this is the function we should use when calculating system security. Calculating system security using both valves' S ratings::

$$\begin{array}{l} S_{PSV11} = 0.995 \\ S_{PSV12} = 0.995 \end{array} \quad \text{AND} \quad S_{system} = 0.990025$$

These numerical results should make sense as well: in an overpressure protection system where a leak in one valve is enough to constitute a problem, the presence of multiple valves is a liability and therefore reduces the over-all security.

It is worth noting that a simple change in parameters may strongly impact our reliability calculations. In this scenario we were told each pressure safety valve was sized large enough to adequately vent the vessel on its own, without the help of the other PSV, in the event of an overpressure condition. What if the PSVs were undersized, and *both* of them would be required to vent in order to protect the vessel from overpressure damage? How would this alteration impact our reliability calculations?

It should be obvious that this change will have no effect whatsoever on the system's security, because it still takes just one PSV to leak in order to make the whole system insecure. However, dependability will definitely be affected by this change because now a single PSV with a $D = 1$ rating is not enough to guarantee a protected system. With undersized PSVs, *both* valves must be dependable in order to guarantee dependable overpressure protection. Conversely, if only one of the PSVs fails in such a way as to be completely undependable ($D = 0$, meaning the valve is guaranteed to fail in the shut condition when faced with high pressure), it makes the whole system undependable because the other valve on its own is not enough to adequately vent the excess gas. From this analysis we can see that the proper Boolean function for dependability will now be AND, because any zero into an AND function guarantees a 0 output. Re-calculating dependability for undersized PSVs:

$$\begin{array}{l} D_{PSV11} = 0.9992 \\ D_{PSV12} = 0.9992 \end{array} \quad \text{AND} \quad D_{system} = 0.99840064$$

32.3 Practical measures of reliability

In reliability engineering, it is important to be able to quantify the reliability (or conversely, the probability of failure) for common components, and for systems comprised of those components. As such, special terms and mathematical models have been developed to describe probability as it applies to component and system reliability.

32.3.1 Failure rate and MTBF

Perhaps the first and most fundamental measure of (un)reliability is the *failure rate* of a component or system of components, symbolized by the Greek letter lambda (λ). The definition of “failure rate” for a group of components undergoing reliability tests is the instantaneous rate of failures per number of surviving components:

$$\lambda = \frac{dN_f}{dt} \frac{1}{N_s} \quad \text{or} \quad \lambda = \frac{dN_f}{dt} \frac{1}{N_s}$$

Where,

λ = Failure rate

N_f = Number of components failed during testing period

N_s = Number of components surviving during testing period

t = Time

The unit of measurement for failure rate (λ) is inverted time units (e.g. “per hour” or “per year”). An alternative expression for failure rate sometimes seen in reliability literature is the acronym *FIT* (“Failures In Time”), in units of 10^{-9} failures per hour. Using a unit with a built-in multiplier such as 10^{-9} makes it easier for human beings to manage the very small λ values normally associated with high-reliability industrial components and systems.

Failure rate may also be applied to discrete-switching (on/off) components and systems of discrete-switching components on the basis of the number of on/off cycles rather than clock time. In such cases, we define failure rate in terms of cycles (c) instead of in terms of minutes, hours, or any other measure of time (t):


$$\lambda = \frac{dN_f}{dc} \frac{1}{N_s} \quad \text{or} \quad \lambda = \frac{dN_f}{dc} \frac{1}{N_s}$$

One of the conceptual difficulties inherent to the definition of lambda (λ) is that it is fundamentally a *rate* of failure over time. This is why the calculus notation $\frac{dN_f}{dt}$ is used to define lambda: a “derivative” in calculus always expresses a rate of change. However, a failure *rate* is not the same thing as the number of devices failed in a test, nor is it the same thing as the probability of failure for one or more of those devices. Failure rate (λ) has more in common with the *time constant* of an resistor-capacitor circuit (τ) than anything else.

An illustrative example is helpful here: if we were to test a large batch of identical components for proper operation over some extended period of time with no maintenance or other intervention, the number of failed components in that batch would gradually accumulate while the number of surviving components in the batch would gradually decline. The reason for this is obvious: every component that fails remains failed (with no repair), leaving one fewer surviving component to function. If we limit the duration of this test to a time-span much shorter than the expected lifetime of the components, any failures that occur during the test must be due to random causes (“Acts of God”) rather than component wear-out.

This scenario is analogous to another random process: rolling a large set of dice, counting any “1” roll as a “fail” and any other rolled number as a “survive.” Imagine rolling the whole batch of dice at once, setting aside any dice landing on “1” aside (counting them as “failed” components in the batch), then only rolling the *remaining* dice the next time. If we maintain this protocol – setting aside “failed” dice after each roll and only continuing to roll “surviving” dice the next time – we will find ourselves rolling fewer and fewer “surviving” dice in each successive roll of the batch. Even though each of the six-sided die has a fixed failure probability of $\frac{1}{6}$, the population of “failed” dice keeps growing over time while the population of “surviving” dice keeps dwindling over time.

Not only does the number of surviving components in such a test dwindle over time, but that number dwindles at an ever-decreasing rate. Likewise with the number of failures: the number of components failing (dice coming up “1”) is greatest at first, but then tapers off after the population of surviving components gets smaller and smaller. Plotted over time, the graph looks something like this:



Rapid changes in the failed and surviving component populations occurs at the start of the test when there is the greatest number of functioning components “in play.” As components fail due to random events, the smaller and smaller number of surviving components results in a slower approach for both curves, simply because there are fewer surviving components remaining to fail.

These curves are precisely identical to those seen in RC (resistor-capacitor) charging circuits, with voltage and current tracing complementary paths: one climbing to 100% and the other falling to 0%, but both of them doing so at ever-decreasing rates. Despite the asymptotic approach of both curves, however, we can describe their approaches in an RC circuit with a constant value τ , otherwise known as the *time constant* for the RC circuit. Failure rate (λ) plays a similar role in describing the failed/surviving curves of a batch of tested components:

$$N_{surviving} = N_o e^{-\lambda t} \quad N_{failed} = N_o (1 - e^{-\lambda t})$$

Where,

$N_{surviving}$ = Number of components surviving at time t

N_{failed} = Number of components failed at time t

N_o = Total number of components in test batch

e = Euler's constant (≈ 2.71828)

λ = Failure rate (assumed to be a constant during the useful life period)

Following these formulae, we see that 63.2% of the components will fail (36.8% will survive) when $\lambda t = 1$ (i.e. after one "time constant" has elapsed).

Unfortunately, this definition for lambda doesn't make much intuitive sense. There is a way, however, to model failure rate in a way that not only makes more immediate sense, but is also more realistic to industrial applications. Imagine a different testing protocol where we maintain a constant sample quantity of components over the entire testing period by immediately replacing each failed device with a working substitute as soon as it fails. Now, the number of functioning devices under test will remain constant rather than declining as components fail. Imagine counting the number of "fails" (dice falling on a "1") for each batch roll, and then rolling *all* the dice in each successive trial rather than setting aside the "failed" dice and only rolling those remaining. If we did this, we would expect a constant fraction ($\frac{1}{6}$) of the six-sided dice to "fail" with each and every roll. The number of failures per roll divided by the total number of dice would be the failure rate (lambda, λ) for these dice. We do not see a curve over time because we do not let the failed components remain failed, and thus we see a constant number of failures with each period of time (with each group-roll).

We may mathematically express this using a different formula:

$$\lambda = \frac{N_f}{t N_o} \quad \text{or} \quad \lambda = \frac{N_f}{t} \frac{1}{N_o}$$

Where,

λ = Failure rate


N_f = Number of components failed during testing period

N_o = Number of components under test (maintained constant) during testing period by immediate replacement of failed components

t = Time

An alternative way of expressing the failure rate for a component or system is the reciprocal of lambda ($\frac{1}{\lambda}$), otherwise known as *Mean Time Between Failures* (MTBF). If the component or system in question is repairable, the expression *Mean Time To Failure* (MTTF) is often used instead¹⁹. Whereas failure rate (λ) is measured in reciprocal units of time (e.g. “per hour” or “per year”), MTBF is simply expressed in units of time (e.g. “hours” or “years”).

For non-maintained tests where the number of failed components accumulates over time (and the number of survivors dwindles), MTBF is precisely equivalent to “time constant” in an RC circuit: MTBF is the amount of time it will take for 63.2% of the components to fail due to random causes, leaving 36.8% of the component surviving. For maintained tests where the number of functioning components remains constant due to swift repairs or replacement of failed components, MTBF (or MTTF) is the amount of time it will take for the total number of tested components to fail²⁰.



It should be noted that these definitions for lambda and MTBF are idealized, and do not necessarily represent all the complexity we see in real-life applications. The task of calculating

¹⁹Since most high-quality industrial devices and systems are repairable for most faults, MTBF and MTTF are interchangeable terms.

²⁰This does not mean the amount of time for *all* components to fail, but rather the amount of time to log a total number of failures equal to the total number of components tested. Some of those failures may be multiple for single components, while some other components in the batch might never fail within the MTBF time.

lambda or MTBF for any real component sample can be quite complex, involving statistical techniques well beyond the scope of instrument technician work.

Simple calculation example: transistor failure rate

Problem: Suppose a semiconductor manufacturer creates a microprocessor “chip” containing 2500000 transistors, each of which is virtually identical to the next in terms of ruggedness and exposure to degrading factors such as heat. The architecture of this microprocessor is such that there is enough redundancy to allow continued operation despite the failure of some of its transistors. This integrated circuit is continuously tested for a period of 1000 days (24000 hours), after which the circuit is examined to count the number of failed transistors. This testing period is well within the useful life of the microprocessor chip, so we know none of the failures will be due to wear-out, but rather to random causes.

Supposing several tests are run on identical chips, with an average of 3.4 transistors failing per 1000-day test. Calculate the failure rate (λ) and the MTBF for these transistors.

Solution: The testing scenario is one where failed components are not replaced, which means both the number of failed transistors and the number of surviving transistors changes over time like voltage and current in an RC charging circuit. Thus, we must calculate lambda by solving for it in the exponential formula.

Using the appropriate formula, relating number of failed components to the total number of components:

$$N_{failed} = N_o (1 - e^{-\lambda t})$$

$$3.4 = 2500000 (1 - e^{-24000\lambda})$$

$$1.36 \times 10^{-6} = 1 - e^{-24000\lambda}$$

$$e^{-24000\lambda} = 1 - 1.36 \times 10^{-6}$$

$$-24000\lambda = \ln(1 - 1.36 \times 10^{-6})$$

$$-24000\lambda = -1.360000925 \times 10^{-6}$$

$$\lambda = 5.66667 \times 10^{-11} \text{ per hour} = 0.0566667 \text{ FIT}$$

Failure rate may be expressed in units of “per hour,” “Failures In Time” (FIT, which means failures per 10^9 hours), or “per year” (pa).

$$\text{MTBF} = \frac{1}{\lambda} = 1.7647 \times 10^{10} \text{ hours} = 2.0145 \times 10^6 \text{ years}$$

Recall that Mean Time Between Failures (MTBF) is essentially the “time constant” for this decaying collection of transistors inside each microprocessor chip.

Simple calculation example: control valve failure rate

Problem: Suppose a control valve manufacturer produces a large number of valves, which are then sold to customers and used in comparable process applications. After a period of 5 years, data is collected on the number of failures these valves experienced. Five years is well within the useful life of these control valves, so we know none of the failures will be due to wear-out, but rather to random causes.

Supposing customers report an average of 15 failures for every 200 control valves in service over the 5-year period, calculate the failure rate (λ) and the MTTF for these control valves.

Solution: The testing scenario is one where failures are repaired in a short amount of time, since these are working valves being maintained in a real process environment. Thus, we may calculate lambda as a simple fraction of failed components to total components.

Using the appropriate formula, relating number of failed components to the total number of components:

$$\lambda = \frac{N_f}{t} \frac{1}{N_o}$$

$$\lambda = \frac{15}{5 \text{ yr}} \frac{1}{200}$$

$$\lambda = \frac{3}{200 \text{ yr}}$$

$$\lambda = 0.015 \text{ per year (pa)} = 1.7123 \times 10^{-6} \text{ per hour}$$

With this value for lambda being so much larger than the microprocessor's transistors, it is not necessary to use a unit such as FIT to conveniently represent it.


$$\text{MTTF} = \frac{1}{\lambda} = 66.667 \text{ years} = 584000 \text{ hours}$$

Recall that Mean Time To Failure (MTTF) is the amount of time it would take²¹ to log a number of failures equal to the total number of valves in service, given the observed rate of failure due to random causes. Note that MTTF is largely synonymous with MTBF. The only technical difference between MTBF and MTTF is that MTTF more specifically relates to situations where components are repairable, which is the scenario we have here with well-maintained control valves.

²¹The typically large values we see for MTBF and MTTF can be misleading, as they represent a *theoretical* time based on the failure rate seen over relatively short testing times where all components are “young.” In reality, the wear-out time of a component will be less than its MTBF. In the case of these control valves, they would likely all “die” of old age and wear long before reaching an age of 66.667 years!

32.3.2 The “bathtub” curve

Failure rate tends to be constant during a component’s useful lifespan where the major cause of failure is random events (“Acts of God”). However, λ does not remain constant over the entire life of the component or system. A common graphical expression of failure rate is the so-called *bathtub curve* showing the typical failure rate profile over time from initial manufacture (brand-new) to wear-out:



This curve profiles the failure rate of a large sample of components (or a large sample of systems) as they age. Failure rate begins at a relatively high value starting at time zero due to defects in manufacture. Failure rate drops off rapidly during a period of time called the *burn-in period* where defective components experience an early death. After the burn-in period, failure rate remains relatively constant over the useful life of the components, and this is where we typically define and apply the failure rate (λ). Any failures occurring during this “useful life” period are due to random mishaps (“Acts of God”). Toward the end of the components’ working lives when the components enter the *wear-out period*, failure rate begins to rise until all components eventually fail. The *mean (average) life* of a component (t_m) is the time required for one-half of the components surviving up until the wear-out time (t_w) to fail, the other half failing after the mean life time.

Several important features are evident in this “bathtub” curve. First, component reliability is greatest between the times of burn-in and wear-out. For this reason, many manufacturers of high-reliability components and systems perform their own burn-in testing prior to sale, so that the customers are purchasing products that have already passed the burn-in phase of their lives. To express this using colloquial terms, we may think of “burnt-in” components as those having already passed through their “growing pains,” and are now “mature” enough to face demanding applications.

Another important measure of reliability is the *mean life*. This is an expression of a component’s (or system’s) operating lifespan. At first this may sound synonymous with MTBF, but it is not.

MTBF – and by extension lambda, since MTBF is the reciprocal of failure rate – is an expression of susceptibility to random (“chance”) failures. Both MTBF and λ_{useful} are quite independent of mean life²². In practice, values for MTBF often greatly exceed values for mean life.

To cite a practical example, the Rosemount model 3051C differential pressure transmitter has a suggested useful lifetime of 50 years (based on the expected service life of tantalum electrolytic capacitors used in its circuitry), while its demonstrated MTBF is 136 years. The larger value of 136 years is a projection based on the failure rate of large samples of these transmitters when they are all “young,” which is why one should never confuse MTBF for service life. In reality, components within the instrument will begin to suffer accelerated failure rates as they reach their end of useful lifetime, as the instrument approaches the right-hand end of the “bathtub” curve.

When determining the length of time any component should be allowed to function in a high-reliability system, the mean life (or even better, the *wear-out* time) should be used as a guide, not the MTBF. This is not to suggest the MTBF is a useless figure – far from it. MTBF simply serves a different purpose, and that is to predict the rate of random failures *during* the useful life span of a large number of components or systems, whereas mean life predicts the service life period where the component’s failure rate remains relatively constant.

²²One could even imagine some theoretical component immune to wear-out, but still having finite values for failure rate and MTBF. Remember, λ_{useful} and MTBF refer to *chance* failures, not the normal failures associated with age and extended use.

32.3.3 Reliability

Reliability (R) is the probability a component or system will perform as designed. Like all probability figures, reliability ranges in value from 0 to 1, inclusive. Given the tendency of manufactured devices to fail over time, reliability decreases with time. During the useful life of a component or system, reliability is related to failure rate by a simple exponential function:

$$R = e^{-\lambda t}$$

Where,

R = Reliability as a function of time (sometimes shown as $R(t)$)

e = Euler's constant (≈ 2.71828)

λ = Failure rate (assumed to be a constant during the useful life period)

t = Time

Knowing that failure rate is the mathematical reciprocal of mean time between failures (MTBF), we may re-write this equation in terms of MTBF as a "time constant" (τ) for random failures during the useful life period:

$$R = e^{\frac{-t}{MTBF}} \quad \text{or} \quad R = e^{\frac{-t}{\tau}}$$

This inverse-exponential function mathematically explains the scenario described earlier where we tested a large batch of components, counting the number of failed components and the number of surviving components over time. Like the dice experiment where we set aside each "failed" die and then rolled only the remaining "survivors" for the next trial in the test, we end up with a diminishing number of "survivors" as the test proceeds.

The same exponential function for calculating reliability applies to single components as well. Imagine a single component functioning within its useful life period, subject only to random failures. The longer this component is relied upon, the more time it has to succumb to random faults, and therefore the less likely it is to function perfectly over the duration of its test. To illustrate by example, a pressure transmitter installed and used for a period of 1 year has a greater chance of functioning perfectly over that service time than an identical pressure transmitter pressed into service for 5 years, simply because the one operating for 5 years has five times more opportunity to fail. In other words, the reliability of a component over a specified time is a function of time, and not just the failure rate (λ).

Using dice once again to illustrate, it is as if we rolled a single six-sided die over and over, waiting for it to “fail” (roll a “1”). The more times we roll this single die, the more likely it will eventually “fail” (eventually roll a “1”). With each roll, the probability of failure is $\frac{1}{6}$, and the probability of survival is $\frac{5}{6}$. Since survival over multiple rolls necessitates surviving the first roll *and* next roll *and* the next roll, all the way to the last surviving roll, the probability function we should apply here is the “AND” (multiplication) of survival probability. Therefore, the survival probability after a single roll is $\frac{5}{6}$, while the survival probability for two successive rolls is $(\frac{5}{6})^2$, the survival probability for three successive rolls is $(\frac{5}{6})^3$, and so on.

The following table shows the probabilities of “failure” and “survival” for this die with an increasing number of rolls:

Number of rolls	Probability of failure (1)	Probability of survival (2, 3, 4, 5, 6)
1	$1 / 6 = 0.16667$	$5 / 6 = 0.83333$
2	$11 / 36 = 0.30556$	$25 / 36 = 0.69444$
3	$91 / 216 = 0.42129$	$125 / 216 = 0.57870$
4	$671 / 1296 = 0.51775$	$625 / 1296 = 0.48225$
n	$1 - (\frac{5}{6})^n$	$(\frac{5}{6})^n$

A practical example of this equation in use would be the reliability calculation for a Rosemount model 1151 analog differential pressure transmitter (with a demonstrated MTBF value of 226 years as published by Rosemount) over a service life of 5 years following burn-in:

$$R = e^{\frac{-5}{226}}$$

$$R = 0.9781 = 97.81\%$$

Another way to interpret this reliability value is in terms of a large batch of transmitters. If three hundred Rosemount model 1151 transmitters were continuously used for five years following burn-in (assuming no replacement of failed units), we would expect approximately 293 of them to still be working (i.e. 6.564 random-cause failures) during that five-year period:

$$N_{surviving} = N_o e^{\frac{-t}{MTBF}}$$

$$\text{Number of surviving transmitters} = (300) \left(e^{\frac{-5}{226}} \right) = 293.436$$

$$N_{failed} = N_o \left(1 - e^{\frac{-t}{MTBF}} \right)$$

$$\text{Number of failed transmitters} = 300 \left(1 - e^{\frac{-5}{226}} \right) = 6.564$$

It should be noted that the calculation will be linear rather than inverse-exponential if we assume immediate replacement of failed transmitters (maintaining the total number of functioning units at 300). If this is the case, the number of random-cause failures is simply $\frac{1}{226}$ per year, or 0.02212 per transmitter over a 5-year period. For a collection of 300 (maintained) Rosemount model 1151 transmitters, this would equate to 6.637 failed units over the 5-year testing span:

$$\text{Number of failed transmitters} = (300) \left(\frac{5}{226} \right) = 6.637$$

32.3.4 Probability of failure on demand (PFD)

Reliability, as previously defined, is the probability a component or system will perform as designed. Like all probability values, reliability is expressed a number ranging between 0 and 1, inclusive. A reliability value of zero (0) means the component or system is totally unreliable (i.e. it is guaranteed to fail). Conversely, a reliability value of one (1) means the component or system is completely reliable (i.e. guaranteed to properly function). In the context of dependability (i.e. the probability that a safety component or system will function *when called upon to act*), the unreliability of that component or system is referred to as *PFD*, an acronym standing for *Probability of Failure on Demand*. Like dependability, this is also a probability value ranging from 0 to 1, inclusive. A PFD value of zero (0) means there is no probability of failure (i.e. it is 100% dependable – guaranteed to properly perform when needed), while a PFD value of one (1) means it is completely undependable (i.e. guaranteed to fail when activated). Thus:

$$\text{Dependability} + \text{PFD} = 1$$

$$\text{PFD} = 1 - \text{Dependability}$$

$$\text{Dependability} = 1 - \text{PFD}$$

Obviously, a system designed for high dependability should exhibit a small PFD value (very nearly 0). Just how low the PFD needs to be is a function of how critical the component or system is to the fulfillment of our human needs.

The degree to which a system must be dependable in order to fulfill our modern expectations is often surprisingly high. Suppose someone were to tell you the reliability of seatbelts in a particular automobile was 99.9 percent (0.999). This sounds rather good, doesn't it? However, when you actually consider the fact that this degree of probability would mean an average of one failed seatbelt for every 1000 collisions, the results are seen to be rather poor (at least to modern American standards of expectation). If the dependability of seatbelts is 0.999, then the PFD is 0.001:

$$\text{PFD} = 1 - \text{Dependability}$$

$$\text{PFD} = 1 - 0.999$$

$$\text{PFD} = 0.001$$

Let's suppose an automobile manufacturer sets a goal of only 1 failed seatbelt in any of its cars during a 1 million unit production run, assuming each and every one of these cars were to crash. Assuming four seatbelts per car, this equates to a $\frac{1}{4000000}$ PFD. The necessary dependability of this manufacturer's seatbelts must therefore be:

$$\text{Dependability} = 1 - \text{PFD} = 1 - \frac{1}{4000000} = 0.99999975$$

Thus, the dependability of these seatbelts must be 99.999975% in order to fulfill the goal of only 1 (potential) seatbelt failure out of 4 million seatbelts produced.

A common order-of-magnitude expression of desired reliability is the number of "9" digits in the reliability value. A reliability value of 99.9% would be expressed as "three nine's" and a reliability value of 99.99% as "four nine's." Expressed thusly, the seatbelt dependability must be "six nine's" in order to achieve the automobile manufacturer's goal.

32.4 High-reliability systems

As discussed at the beginning of this chapter, instrumentation safety may be broadly divided into two categories: the safety hazards posed by malfunctioning instruments, and special instrument systems designed to reduce safety hazards of industrial processes. This section regards the first category.

All methods of reliability improvement incur some extra cost on the operation, whether it be capital expense (initial purchase/installation cost) or continuing expense (labor or consumables). The choice to improve system reliability is therefore very much an economic one. One of the human challenges associated with reliability improvement is continually justifying this cost over time. Ironically, the more successful a reliability improvement program has been, the less important that program seems. The manager of an operation suffering from reliability problems does not need to be convinced of the economic benefit of reliability improvement as much as the manager of a trouble-free facility. Furthermore, the people most aware of the benefits of reliability improvement are usually those tasked with reliability-improving duties (such as preventive maintenance), while the people least aware of the same benefits are usually those managing budgets. If ever a disagreement erupts between the two camps, pleas for continued financial support of reliability improvement programs may be seen as nothing more than self-interest, further escalating tensions²³.

A variety of methods exist to improve the reliability of systems. The following subsections investigate several of them.

²³Preventive maintenance is not the only example of such a dynamic. Modern society is filled with monetarily expensive programs and institutions existing for the ultimate purpose of avoiding *greater* costs, monetary and otherwise. Public education, health care, and national militaries are just a few that come to my mind. Not only is it a challenge to continue justifying the expense of a well-functioning cost-avoidance program, but it is also a challenge to detect and remove unnecessary expenses (waste) within that program. To extend the preventive maintenance example, an appeal by maintenance personnel to continue (or further) the maintenance budget may happen to be legitimate, but a certain degree of self-interest will always be present in the argument. Just because preventive maintenance is actually necessary to avoid greater expense due to failure, does not mean *all* preventive maintenance demands are economically justified! Proper funding of any such program depends on the financiers being fair in their judgment *and* the executors being honest in their requests. So long as both parties are human, this territory will remain contentious.

32.4.1 Design and selection for reliability

Many workable designs may exist for electronic and mechanical systems alike, but not all are equal in terms of reliability. A major factor in machine reliability, for example, is *balance*. A well-balanced machine will operate with little vibration, whereas an ill-balanced machine will tend to shake itself (and other devices mechanically coupled to it) apart over time²⁴.

Electronic circuit reliability is strongly influenced by design as well as by component choice. An historical example of reliability-driven design is found in the Foxboro SPEC 200 analog control system. The reliability of the SPEC 200 control system is legendary, with a proven record of minimal failures over many years of industrial use. According to Foxboro technical literature, several design guidelines were developed following application experience with Foxboro electronic field instruments (most notably the “E” and “H” model lines), among them the following:

- All critical switches should spend most of their time in the *closed* state
- Avoid the use of carbon composition resistors – use wirewound or film-type resistors instead
- Avoid the use of plastic-cased semiconductors – use glass-cased or hermetically sealed instead
- Avoid the use of electrolytic capacitors wherever possible – use polycarbonate or tantalum instead

Each of these design guidelines is based on minimization of component failure. Having switches spend most of their lives in the closed state means their contact surfaces will be less exposed to air and therefore less susceptible to corrosion over time (leading to an “open” fault). Wirewound resistors are better able to tolerate vibration and physical abuse than brittle carbon-composition designs. Glass-cased and hermetically-sealed semiconductors are better at sealing out moisture than plastic-cased semiconductors. Electrolytic capacitors are famously unreliable compared to other capacitor types such as polycarbonate, and so their avoidance is wise.

In addition to high-quality component characteristics and excellent design practices, components used in these lines of Foxboro instruments were “burned in” prior to circuit board assembly, thus avoiding many “early failures” due to components “burning in” during actual service.

²⁴Sustained vibrations can do really strange things to equipment. It is not uncommon to see threaded fasteners undone slowly over time by vibrations, as well as cracks forming in what appear to be extremely strong supporting elements such as beams, pipes, etc. Vibration is almost never good for mechanical (or electrical!) equipment, so it should be eliminated wherever reliability is a concern.

32.4.2 Preventive maintenance

The term *preventive maintenance* refers to the maintenance (repair or replacement) of components prior to their inevitable failure in a system. In order to intelligently schedule the replacement of critical system components, some knowledge of those components' useful lifetimes is necessary. On the standard "bathtub curve," this corresponds with the *wear-out time* or $t_{wear-out}$.

In many industrial operations, preventive maintenance schedules (if they exist at all) are based on past history of component lifetimes, and the operational expenses incurred due to failure of those components. Preventive maintenance represents an up-front cost, paid in exchange for the avoidance of larger costs later in time.

A common example of preventive maintenance and its cost savings is the periodic replacement of lubricating oil and oil filters for automobile engines. Automobile manufacturers provide specifications for the replacement of oil and filters based on testing of their engines, and assumptions made regarding the driving habits of their customers. Some manufacturers even provide dual maintenance schedules, one for "normal" driving and another for "heavy" or "performance" driving to account for accelerated wear. As trivial as an oil change might seem to the average driver, regular maintenance to an automobile's lubrication system is absolutely critical not only to long service life, but also to optimum performance. Certainly, the consequences of not performing this preventive maintenance task on an automobile's engine will be costly²⁵.

Another example of preventive maintenance for increased system reliability is the regular replacement of light bulbs in traffic signal arrays. For rather obvious reasons, the proper function of traffic signal lights is critical for smooth traffic flow and public safety. It would not be a satisfactory state of affairs to replace traffic signal light bulbs only when they failed, as is common with the replacement of most light bulbs. In order to achieve high reliability, these bulbs must be replaced in advance of their expected wear-out times²⁶. The cost of performing this maintenance is undeniable, but then so is the (greater) cost of congested traffic and accidents caused by burned-out traffic light bulbs.

An example of preventive maintenance in industrial instrumentation is the installation and service of *dryer* mechanisms for compressed air, used to power pneumatic instruments and valve actuators. Compressed air is a very useful medium for transferring (and storing) mechanical energy, but problems will develop within pneumatic instruments if water is allowed to collect within air distribution systems. Corrosion, blockages, and hydraulic "locking" are all potential consequences of "wet" instrument air. Consequently, instrument compressed air systems are usually installed separate from utility compressed air systems (used for operating general-purpose pneumatic tools and equipment actuators), using different types of pipe (plastic, copper, or stainless steel rather than black iron or galvanized iron) to avoid corrosion and using *air dryer* mechanisms near the compressor to absorb and expel moisture. These air dryers typically use a beaded *desiccant* material to absorb

²⁵On an anecdotal note, a friend of mine once destroyed his car's engine, having never performed an oil or filter change on it since the day he purchased it. His poor car expired after only 70000 miles of driving – a mere fraction of its normal service life with regular maintenance. Given the type of car it was, he could have easily expected 200000 miles of service between engine rebuilds had he performed the recommended maintenance on it.

²⁶Another friend of mine used to work as a traffic signal technician in a major American city. Since the light bulbs they replaced still had some service life remaining, they decided to donate the bulbs to a charity organization where the used bulbs would be freely given to low-income citizens. Incidentally, this same friend also instructed me on the proper method of inserting a new bulb into a socket: twisting the bulb just enough to maintain some spring tension on the base, rather than twisting the bulb until it will not turn farther (as most people do). Maintaining some natural spring tension on the metal leaf within the socket helps extend the socket's useful life as well!

water vapor from the compressed air, and then this desiccant material is periodically purged of its retained water. After some time of operation, though, the desiccant must be physically removed and replaced with fresh desiccant.

32.4.3 Component de-rating

Some²⁷ control system components exhibit an inverse relationship between service load (how “hard” the component is used) and service life (how long it will last). In such cases, a way to increase service life is to *de-rate* that component: operate it at a load reduced from its design rating.

For example, a variable-frequency motor drive (VFD) takes AC power at a fixed frequency and voltage and converts it into AC power of varying frequency and voltage to drive an induction motor at different speeds and torques. These electronic devices dissipate some heat owing mostly to the imperfect (slightly resistive) “on” states of power transistors. Temperature is a wear factor for semiconductor devices, with greater temperatures leading to reduced service lives. A VFD operating at high temperature, therefore, will fail sooner than a VFD operating at low temperature, all other factors being equal. One way to reduce the operating temperature of a VFD is to over-size it for the application. If the motor to be driven requires 2 horsepower of electrical power at full load, and increased reliability is demanded of the drive, then perhaps a 5 horsepower VFD (programmed with reduced trip settings appropriate to the smaller motor) could be chosen to drive the motor.

In addition to extending service life, de-rating also has the ability to amplify the mean time between failure (MTBF) of load-sensitive components. Recall that MTBF is the reciprocal of failure rate during the low area of the “bathtub curve,” representing failures due to random causes. This is distinct from wear-out, which is an increase in failure rate due to irreversible wear and aging. The main reason a component will exhibit a greater MTBF value as a consequence of de-rating is that the component will be better able to absorb transient overloads, which is a typical cause of failure during the operational life of system components.

Consider the example of a pressure sensor in a process known to exhibit transient pressure surges. A sensor chosen such that the typical process operating pressure spans most of its range will have little overpressure capacity. Perhaps just a few over-pressure events will cause this sensor to fail well before its rated service life. A de-rated pressure sensor (with a pressure-sensing range covering much greater pressures than what are normally encountered in this process), by comparison, will have more pressure capacity to withstand random surges, and therefore exhibit less probability of random failure.

The costs associated with component de-rating include initial investment (usually greater, owing to the greater capacity and more robust construction compared to a “normally” rated component) and reduced sensitivity. The latter factor is an important one to consider if the component is expected to provide high accuracy as well as high reliability. In the example of the de-rated pressure sensor, accuracy will likely suffer because the full pressure range of the sensor is not being used for normal process pressure measurements. If the instrument is digital, resolution will certainly suffer as a result of de-rating the instrument’s measurement range. Alternative methods of reliability improvement (including more frequent preventive maintenance) may be a better solution than de-rating in such cases.

²⁷Many components do not exhibit any relationship between load and lifespan. An electronic PID controller, for example, will last just as long controlling an “easy” self-regulating process as it will controlling a “difficult” unstable (“runaway”) process. The same might not be said for the other components of those loops, however! If the control valve in the self-regulating process rarely changes position, but the control valve in the runaway process continually moves in an effort to stabilize it at setpoint, the less active control valve will most likely enjoy a longer service life.

32.4.4 Redundant components

The MTBF of any system dependent upon certain critical components may be extended by duplicating those components in parallel fashion, such that the failure of only one does not compromise the system as a whole. This is called *redundancy*. A common example of component redundancy in instrumentation and control systems is the redundancy offered by distributed control systems (DCSs), where processors, network cables, and even I/O (input/output) channels may be equipped with “hot standby” duplicates ready to assume functionality in the event the primary component fails.


Redundancy tends to extend the MTBF of a system without necessarily extending its service life. A DCS, for example, equipped with redundant microprocessor control modules in its rack, will exhibit a greater MTBF because a random microprocessor fault will be covered by the presence of the spare (“hot standby”) microprocessor module. However, given the fact that both microprocessors are continually powered, and therefore tend to “wear” at the same rate, their operating lives will not be additive. In other words, two microprocessors will not function twice as long before wear-out than one microprocessor.

The extension of MTBF resulting from redundancy holds true only if the random failures are truly independent events – that is, not associated by a common cause. To use the example of a DCS rack with redundant microprocessor control modules again, the susceptibility of that rack to a random microprocessor fault will be reduced by the presence of redundant microprocessors *only* if the faults in question are unrelated to each other, affecting the two microprocessors separately. There may exist common-cause fault mechanisms capable of disabling *both* microprocessor modules as easily as it could disable one, in which case the redundancy adds no value at all. Examples of such common-cause faults include power surges (because a surge strong enough to kill one module will likely kill the other at the same time) and a computer virus infection (because a virus able to attack one will be able to attack the other just as easily, and at the same time).

A simple example of component redundancy in an industrial instrumentation system is dual DC power supplies feeding through a diode module. The following photograph shows a typical example, in this case a pair of Allen-Bradley AC-to-DC power supplies for a DeviceNet digital network:



If either of the two AC-to-DC power supplies happens to fail with a low output voltage, the other power supply is able to carry the load by passing its power through the diode redundancy module²⁸:




²⁸This redundancy module has its own MTBF value, and so by including it in the system we are adding one more component that can fail. However, the MTBF rate of a simple diode network greatly exceeds that of an entire AC-to-DC power supply, and so we find ourselves at a greater level of reliability using this diode redundancy module than if we did not (and only had one power supply).

In order for redundant components to actually increase system MTBF, the potential for common-cause failures must be addressed. For example, consider the effects of powering redundant AC-to-DC power supplies from the exact same AC line. Redundant power supplies would increase system reliability in the face of a random power supply failure, but this redundancy would do *nothing at all* to improve system reliability in the event of the common AC power line failing! In order to enjoy the fullest benefit of redundancy in this example, we must source each AC-to-DC power supply from a different (unrelated) AC line.

Another example of redundancy in industrial instrumentation is the use of multiple transmitters to sense the same process variable, the notion being that the critical process variable will still be monitored even in the event of a transmitter failure. Thus, installing redundant transmitters should increase the MTBF of the system's sensing ability.

Here again, we must address common-cause failures in order to reap the full benefits of redundancy. If three liquid level transmitters are installed to measure the exact same liquid level, their combined signals represent an increase in measurement system MTBF *only* for independent faults. A failure mechanism common to all three transmitters will leave the system just as vulnerable to random failure as a single transmitter. In order to achieve optimum MTBF in redundant sensor arrays, the sensors must be immune to common faults.

In this example, three different types of level transmitter monitor the level of liquid inside a vessel, their signals processed by a *selector* function programmed inside a DCS:



Here, level transmitter 23a is a guided-wave radar (GWR), level transmitter 23b is a tape-and-

float, and level transmitter 23c is a differential pressure sensor. All three level transmitters sense liquid level using different technologies, each one with its own strengths and weaknesses. Better redundancy of measurement is obtained this way, since no single process condition or other random event is likely to fault more than one of the transmitters at any given time.

For instance, if the process liquid density happened to suddenly change, it would affect the measurement accuracy of the differential pressure transmitter (LT-23c), but not the radar transmitter nor the tape-and-float transmitter. If the process vapor density were to suddenly change, it might affect the radar transmitter (since vapor density generally affects dielectric constant, and dielectric constant affects the propagation velocity of electromagnetic waves, which in turn will affect the time taken for the radar pulse to strike the liquid surface and return), but this will not affect the float transmitter's accuracy nor will it affect the differential pressure transmitter's accuracy. Surface turbulence of the liquid inside the vessel may severely affect the float transmitter's ability to accurately sense liquid level, but it will have little effect on the differential pressure transmitter's reading nor the radar transmitter's measurement (assuming the radar transmitter is shrouded in a *stilling well*).

If the selector function takes either the median (middle) measurement or an average of the best 2-out-of-3 ("2oo3"), none of these random process occurrences will greatly affect the selected measurement of liquid level inside the vessel. True redundancy is achieved here, since the three level transmitters are not only less likely to (all) fail simultaneously than for any single transmitter to fail, but also because the level is being sensed in three completely different ways.

A crucial requirement for redundancy to be effective is that all redundant components must have precisely the same process function. In the case of redundant DCS components such as processors, I/O cards, and network cables, each of these redundant components must do nothing more than serve as "backup" spares for their primary counterparts. If a particular DCS node were equipped with two processors – one as the primary and another as a secondary (backup) – but yet the backup processor were tasked with some detail specific to it and not to the primary processor (or vice-versa), the two processors would *not* be truly redundant to each other. If one processor were to fail, the other would not perform *exactly* the same function, and so the system's operation would be affected (even if only in a small way) by the processor failure.

Likewise, redundant sensors must perform the exact same process measurement function in order to be truly redundant. A process equipped with triplicate measurement transmitters such as the previous example where a vessel's liquid level was being measured by a guided-wave radar, tape-and-float, and differential pressure based level transmitters, would enjoy the protection of redundancy if and only if all three transmitters sensed the exact same liquid level over the exact same calibrated range. This often represents a challenge, in finding suitable locations on the process vessel for three different instruments to sense the exact same process variable. Quite often, the pipe fittings penetrating the vessel (often called *nozzles*) are not conveniently located to accept multiple instruments at the points necessary to ensure consistency of measurement between them. This is often the case when an existing process vessel is retrofitted with redundant process transmitters. New construction is usually less of a problem, since the necessary nozzles and other accessories may be placed in their proper positions during the design stage²⁹.

If fluid flow conditions inside a process vessel are excessively turbulent, multiple sensors installed

²⁹Of course, this assumes good communication and proper planning between all parties involved. It is not uncommon for piping engineers and instrument engineers to mis-communicate during the crucial stages of process vessel design, so that the vessel turns out not to be configured as needed for redundant instruments.

to measure the same variable will sometimes report significant differences. Multiple temperature transmitters located in close proximity to each other on a distillation column, for example, may report significant differences of temperature if their respective sensing elements (thermocouples, RTDs) contact the process liquid or vapor at points where the flow patterns vary. Multiple liquid level sensors, even of the same technology, may report differences in liquid level if the liquid inside the vessel swirls or “funnels” as it enters and exits the vessel.

Not only will substantial measurement differences between redundant transmitters compromise their ability to function as “backup” devices in the event of a failure, such differences may actually “fool” a redundant system into thinking one or more of the transmitters has already failed, thereby causing the deviating measurement to be ignored. To use the triplicate level-sensing array as an example again, suppose the radar-based level transmitter happened to register two inches greater level than the other two transmitters due to the effects³⁰ of liquid swirl inside the vessel. If the selector function is programmed to ignore such deviating measurements, the system degrades to a duplicate-redundant instead of triplicate-redundant array. In the event of a dangerously low liquid level, for example, only the radar-based and float-based level transmitters will be ready to signal this dangerous process condition to the control system, because the pressure-based level transmitter is registering too high.

³⁰If a swirling fluid inside the vessel encounters a stationary baffle, it will tend to “pile up” on one side of that baffle, causing the liquid level to actually be greater in that region of the vessel than anywhere else inside the vessel. Any transmitter placed within this region will register a greater level, regardless of the measurement technology used.

32.4.5 Proof tests and self-diagnostics

A reliability enhancing technique related to preventive maintenance of critical instruments and functions, but generally not as expensive as component replacement, is periodic *testing* of component and system function. Regular “proof testing” of critical components enhances the MTBF of a system for two different reasons:

- Early detection of developing problems
- Regular “exercise” of components

First, proof testing may reveal weaknesses developing in components, indicating the need for replacement in the near future. An analogy to this is visiting a doctor to get a comprehensive exam – if this is done regularly, potentially fatal conditions may be detected early and crises averted.

The second way proof testing increases system reliability is by realizing the beneficial effects of regular function. The performance of many component and system types tends to degrade after prolonged periods of inactivity³¹. This tendency is most prevalent in mechanical systems, but holds true for some electrical components and systems as well. Solenoid valves, for instance, may become “stuck” in place if not cycled for long periods of time. Bearings may corrode and seize in place if left immobile. Both primary- and secondary-cell batteries are well known for their tendency to fail after prolonged periods of non-use. Regular cycling of such components actually *enhances* their reliability, decreasing the probability of a “stagnation” related failure well before the rated useful life has elapsed.

An important part of any proof-testing program is to ensure a ready stock of spare components is kept on hand in the event proof-testing reveals a failed component. Proof testing is of little value if the failed component cannot be immediately repaired or replaced, and so these warehoused components should be configured (or be easily configurable) with the exact parameters necessary for immediate installation. A common tendency in business is to focus attention on the engineering and installation of process and control systems, but neglect to invest in the support materials and infrastructure to keep those systems in excellent condition. High-reliability systems have special needs, and this is one of them.

³¹The father of a certain friend of mine has operated a used automobile business for many years. One of the tasks given to this friend when he was a young man, growing up helping his father in his business, was to regularly drive some of the cars on the lot which had not been driven for some time. If an automobile is left un-operated for many weeks, there is a marked tendency for batteries to fail and tires to lose their air pressure, among other things. The salespeople at this used car business jokingly referred to this as *lot rot*, and the only preventive measure was to routinely drive the cars so they would not “rot” in stagnation. Machines, like people, suffer if subjected to a lack of physical activity.

Methods of proof testing

The most direct method of testing a critical system is to stimulate it to its range limits and observe its reaction. For a process transmitter, this sort of test usually takes the form of a full-range calibration check. For a controller, proof testing would consist of driving all input signals through their respective ranges in all combinations to check for the appropriate output response(s). For a final control element (such as a control valve), this requires full stroking of the element, coupled with physical leakage tests (or other assessments) to ensure the element is having the intended effect on the process.

An obvious challenge to proof testing is how to perform such comprehensive tests without disrupting the process in which it functions. Proof-testing an out-of-service instrument is a simple matter, but proof-testing an instrument installed in a working system is something else entirely. How can transmitters, controllers, and final control elements be manipulated through their entire operating ranges without actually disturbing (best case) or halting (worst case) the process? Even if all tests may be performed at the required intervals during shut-down periods, the tests are not as realistic as they could be with the process operating at typical pressures and temperatures. Proof-testing components during actual “run” conditions is the most realistic way to assess their readiness.

One way to proof-test critical instruments with minimal impact to the continued operation of a process is to perform the tests on only some components, not all. For instance, it is a relatively simple matter to take a transmitter out of service in an operating process to check its response to stimuli: simply place the controller in manual mode and let a human operator control the process manually while an instrument technician tests the transmitter. While this strategy admittedly is not comprehensive, at least proof-testing some of the instruments is better than proof-testing none of them.

Another method of proof-testing is to “test to shutdown:” choose a time when operations personnel plan on shutting the process down anyway, then use that time as an opportunity to proof-test one or more critical component(s) necessary for the system to run. This method enjoys the greatest degree of realism, while avoiding the inconvenience and expense of an unnecessary process interruption.

Yet another method to perform proof tests on critical instrumentation is to accelerate the speed of the testing stimuli so that the final control elements will not react fully enough to actually disrupt the process, but yet will adequately assess the responsiveness of all (or most) of the components in question. The nuclear power industry sometimes uses this proof-test technique, by applying high-speed pulse signals to safety shutdown sensors in order to test the proper operation of shutdown logic, without actually shutting the reactor down. The test consists of injecting short-duration pulse signals at the sensor level, then monitoring the output of the shutdown logic to ensure consequent pulse signals are sent to the shutdown device(s). Various chemical and petroleum industries apply a similar proof-testing technique to safety valves called *partial stroke testing*, whereby the valve is stroked only part of its travel: enough to ensure the valve is capable of adequate motion without closing (or opening, depending on the valve function) enough to actually disrupt the process.

Redundant systems offer unique benefits and challenges to component proof-testing. The benefit of a redundant system in this regard is that any one redundant component may be removed from service for testing without any special action by operations personnel. Unlike a “simplex” system where removal of an instrument requires a human operator to manually take over control during the

duration of the test, the “backup” components of a redundant system should do this automatically, theoretically making the test much easier to conduct. However, the challenge of doing this is the fact that the portion of the system responsible for ensuring seamless transition in the event of a failure is in fact a component liable to failure itself. The only way to test this component is to actually disable one (or more, in highly redundant configurations) of the redundant components to see whether or not the remaining component(s) perform their redundant roles. So, proof-testing a redundant system harbors no danger if all components of the system are good, but risks process disruption if there happens to be an undetected fault.

Let us return to our triplicate level transmitter system once again to explore these concepts. Suppose we wished to perform a proof-test of the pressure-based level transmitter. Being one of three transmitters measuring liquid level in this vessel, we should be able to remove it from service with no preparation (other than notifying operations personnel of the test, and of the potential consequences) since the selector function should automatically de-select the disabled transmitter and continue measuring the process via the remaining two transmitters. If the proof-testing is successful, it proves not only that the transmitter works, but also that the selector function adequately performed its task in “backing up” the tested transmitter while it was removed. However, if the selector function happened to be failed when we disable the one level transmitter for proof-testing, the selected process level signal could register a faulty value instead of switching to the two remaining transmitters’ signals. This might disrupt the process, especially if the selected level signal went to a control loop or to an automatic shutdown system. We could, of course, proceed with the utmost caution by having operations personnel place the control system in “manual” mode while we remove that one transmitter from service, just in case the redundancy does not function as designed. Doing so, however, fails to fully test the system’s redundancy, since by placing the system in manual mode before the test we do not allow the redundant logic to fully function as it would be expected to in the event of an actual instrument failure.

Regular proof-testing is an essential activity to realize optimum reliability for any critical system. However, in all proof-testing we are faced with a choice: either test the components to their fullest degree, in their normal operating modes, and risk (or perhaps guarantee) a process disruption; or perform a test that is less than comprehensive, but with less (or no) risk of process disruption. In the vast majority of cases, the latter option is chosen simply due to the costs associated with process disruption. Our challenge as instrumentation professionals is to formulate proof tests that are as comprehensive as possible while being the least disruptive to the process we are trying to regulate.

Instrument self-diagnostics

One of the great advantages of digital electronic technology in industrial instrumentation is the inclusion of *self-diagnostic* ability in field instruments. A “smart” instrument containing its own microprocessor may be programmed to detect certain conditions known to indicate sensor failure or other problems, then signal the control system that something is wrong. Though self-diagnostics can never be perfectly effective in that there will inevitably be cases of undetected faults and even false positives (declarations of a fault where none exists), the current state of affairs is considerably better than the days of purely analog technology where instruments possessed little or no self-diagnostic capability.

Digital field instruments have the ability to communicate self-diagnostic error messages to their host systems over the same “fieldbus” networks they use to communicate regular process data. FOUNDATION Fieldbus instruments in particular have extensive error-reporting capability, including a “status” variable associated with every process signal that propagates down through all function blocks responsible for control of the process. Detected faults are efficiently communicated throughout the information chain in the system when instruments have full digital communication ability.

“Smart” instruments with self-diagnostic ability but limited to analog (e.g. 4-20 mA DC) signaling may also convey error information, just not as readily or as comprehensively as a fully digital instrument. The NAMUR recommendations for 4-20 mA signaling (NE-43) provide a means to do this:


Signal level	Fault condition
Output ≤ 3.6 mA	Sensing transducer failed low
$3.6 \text{ mA} < \text{Output} < 3.8 \text{ mA}$	Sensing transducer failed (detected) low
$3.8 \text{ mA} \leq \text{Output} < 4.0 \text{ mA}$	Measurement under-range
$21.0 > \text{Output} \geq 20.5 \text{ mA}$	Measurement over-range
Output ≥ 21.0 mA	Sensing transducer failed high

Proper interpretation of these special current ranges, of course, demands a receiver capable of accurate current measurement outside the standard 4-20 mA range. Many control systems with analog input capability are programmed to recognize the NAMUR error-indicating current levels.

A challenge for any self-diagnostic system is how to check for faults in the “brain” of the unit itself: the microprocessor. If a failure occurs within the microprocessor of a “smart” instrument – the very component responsible for performing logic functions related to self-diagnostic testing – how would it be able to detect a fault in logic? The question is somewhat philosophical, equivalent to determining whether or not a neurologist is able to diagnose his or her own neurological problems.

One simple method of detecting gross faults in a microprocessor system is known as a *watchdog timer*. The principle works like this: the microprocessor is programmed to output continuous a low-frequency pulse signal, with an external circuit “watching” that pulse signal for any interruptions or freezing. If the microprocessor fails in any significant way, the pulse signal will either skip pulses or “freeze” in either the high or low state, thus indicating a microprocessor failure to the “watchdog” circuit.

One may construct a watchdog timer circuit using a pair of solid-state timing relays connected to the pulse output channel of the microprocessor device:



Both the on-delay and off-delay timers receive the same pulse signal from the microprocessor, their inputs connected directly in parallel with the microprocessor's pulse output. The off-delay timer immediately actuates upon receiving a "high" signal, and begins to time when the pulse signal goes "low." The on-delay timer begins to time during a "high" signal, but immediately de-actuates whenever the pulse signal goes "low." So long as the time settings for the on-delay and off-delay timer relays are greater than the "high" and "low" durations of the watchdog pulse signal, respectively, neither relay contact will open as long as the pulse signal continues in its regular pattern.

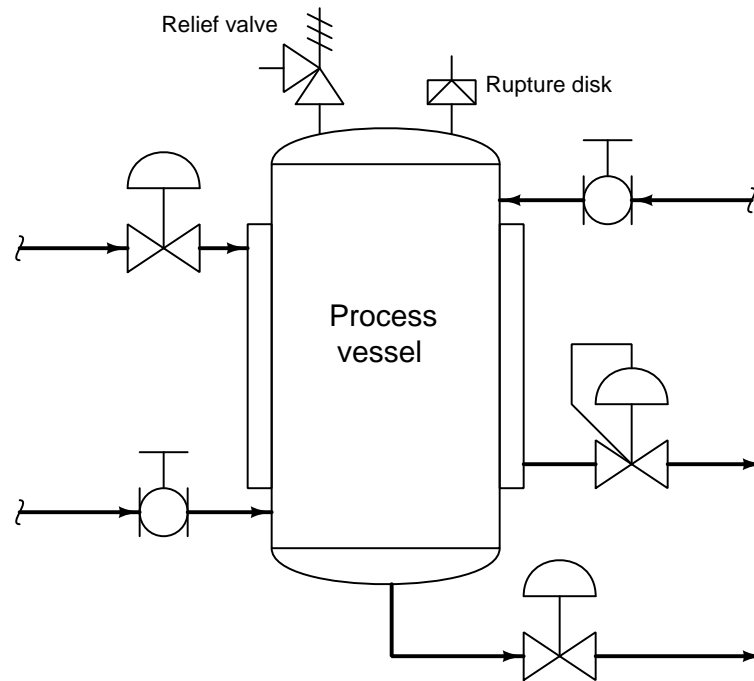
When the microprocessor is behaving normally, outputting a regular watchdog pulse signal, the off-delay timer's contact will hold in a closed state because it keeps getting energized with each "high" signal and never has enough time to drop out during each "low" signal. Likewise, the on-delay timer's contact will remain in its normally closed state because it never has enough time to pick up during each "high" signal before being de-actuated with each "low" signal. Both timing relay contacts will be in a closed state when all is well.

However, if the microprocessor's pulse output signal happens to freeze in the "low" state (or skip a "high" pulse), the off-delay timer will de-actuate, opening its contact and signaling a fault. Conversely, if the microprocessor's pulse signal happens to freeze in the "high" state (or skip a "low" pulse), the on-delay timer will actuate, opening its contact and signaling a fault. Either timing relay opening its contact signals an interruption or cessation of the watchdog pulse signal, indicating a serious microprocessor fault.

32.5 Overpressure protection devices

Fluid pressure exerts force on any surface area it contacts, as described by the formula $F = PA$. One practical consequence of this fact is that process vessels and pipelines may catastrophically burst if subjected to excessive fluid pressure. If subjected to excessive vacuum, some vessels and piping may implode (collapse in on themselves). Not only do these potential failures pose operational problems, but they may also pose severe safety and environmental hazards, especially if the process fluid in question is toxic, flammable, or both.

Special safety devices exist to help prevent such unfortunately events from occurring, among them being *rupture disks*, *relief valves*, and *safety valves*. The following subsections describe each of these protective devices and their intended operation. In a P&ID, rupture disks and relief valves are represented by the following symbols:



A rupture disk acts like an electrical fuse for overpressure protection: when the burst pressure is exceeded, the disk ruptures to let fluids escape through it. Safety and relief valves work like self-resetting circuit breakers: they open to relieve pressure, then re-close to seal the process system once more.

Two common causes of process overpressure are *piping blockages* and overheating caused by *fires*. Although it may sound ridiculous, a number of fatal industrial accidents have been caused by something as simple as shut block valves that should have been left open. When fluid cannot escape a process vessel, the pumping forces may exceed the burst rating of the vessel, causing catastrophic failure. Fires may also cause overpressure conditions, owing to the expansion of process fluids inside sealed vessels. Overpressure protection devices play a crucial role in such scenarios, venting process fluid so as to avoid bursting the vessel. It should be mentioned that these two causes of overpressure

may have vastly differing protection requirements: the required flow rate of exiting fluid to safely limit pressure may be far greater in a “fire case” than it is for a “blockage case,” which means overpressure protection devices sized for the latter may be insufficient to protect against the former.

Overpressure protection device selection is a task restricted to the domain of process safety engineers. Instrument technicians may be involved in the installation and maintenance of overpressure protection devices, but only a qualified and licensed engineer should decide which specific device(s) to use for a particular process system.

32.5.1 Rupture disks

One of the simplest forms of overpressure protection for process lines and vessels is a device known as a *rupture disk*. This is nothing more than a thin sheet of material (usually alloy steel) designed to rupture in the event of an overpressure condition. The amount of force applied to this thin metal sheet is given by the formula $F = PA$ (force equals pressure times area). The thin metal sheet is designed to rupture at a certain threshold of force equivalent to the burst pressure ($P = \frac{F}{A}$). Once the disk ruptures, the fluid vents through new-formed path, thus relieving pressure. Like an electrical fuse, a rupture disk is a one-time device which must be replaced after it has “blown.”

A photograph of a small rupture disk (prior to being placed in service) appears here:



The particular rupture disk shown here has a burst pressure of 30 PSI at a temperature of 130 °C. Temperature is an important factor in the rating of a rupture disk, as the physical strength of the thin metal rupture element changes with temperature. This metal disk is usually quite thin, usually in the order of 0.002 to 0.060 inches in thickness.

Some modern rupture disks use a *graphite* rupture element instead of metal. Not only does graphite exhibit better corrosion resistance to process fluids than metal, but it also does not fatigue in the same way that metal will over time. Burst pressure for a graphite rupture disk, therefore, may be more consistent than with a metal disk. A significant disadvantage of graphite rupture disks, however, is their tendency to shatter upon bursting. Metal rupture disks merely tear, but a graphite rupture disk tends to break into small pieces which are then carried away by the exiting

fluid. These graphite shards may exit as shrapnel, or even lodge inside the mechanism of a valve if one is installed downstream of the rupture disk.

32.5.2 Direct-actuated safety and relief valves

Pressure Relief Valves (PRVs) and *Pressure Safety Valves* (PSVs) are special types of valves designed to open up in order to relieve excess pressure from inside a process vessel or piping system. These valves are normally shut, opening only when sufficient fluid pressure develops across them to relieve that process fluid pressure and thereby protect the pipes and vessels upstream. Unlike regular control valves, PRVs and PSVs are actuated by the process fluid pressure itself rather than by some external pressure or force (e.g. pneumatic signal pressure, electrical motor or solenoid coil).

While the terms “Relief Valve” and “Safety Valve” are sometimes interchanged, there is a distinct difference in operation between them. A *relief valve* opens in direct proportion to the amount of overpressure it experiences in the process piping. That is, a PRV will open slightly for slight overpressures, and open more for greater overpressures. Pressure Relief Valves are commonly used in liquid services. By contrast, a *safety valve* opens fully with a “snap action” whenever it experiences a sufficient overpressure condition, not closing until the process fluid pressure falls significantly below that “lift” pressure value. In other words, a PSV’s action is *hysteretic*³². Pressure Safety Valves are commonly used in gas and vapor services, such as compressed air systems and steam systems.

Safety valves typically have two pressure ratings: the pressure value required to initially open (“lift”) the valve, and the pressure value required to reseal (close) the valve. The difference between these two pressure is called the *blowdown* pressure. A safety valve’s lift pressure will always exceed its reseal pressure, giving the valve a hysteretic behavior.

³²A simple “memory trick” I use to correctly distinguish between relief and safety valves is to remember that a safety valve has snap action (both words beginning with the letter “s”).

This photograph shows a Varec pressure relief valve on an industrial hot water system, designed to release pressure to atmosphere if necessary to prevent damage to process pipes and vessels in the system:




The vertical pipe is the atmospheric vent line, while the bottom flange of this PRV connects to the pressurized hot water line. A large spring inside the relief valve establishes the lift pressure.

A miniature pressure relief valve manufactured by Nupro, cut away to show its internal components, appears in this next photograph. The pipe fittings on this valve are 1/4 inch NPT, to give a sense of scale:



A close-up photograph shows the plug and seat inside this PRV, pointed to by the tip of a ball-point pen:



A simple tension-adjusting mechanism on a spring establishes this valve's lift pressure. The spring exerts a force on the stem to the right, pressing the plug against the face of the seat. A knob allows manual adjustment of spring tension, relating directly to lift pressure:



The operation of this relief valve mechanism is quite simple: process fluid pressure entering the right-hand side fitting exerts force against the plug, which normally blocks passage of the fluid through to the side fitting. The area of the plug serves as a piston for the fluid pressure to push against, the amount of force predicted by the familiar force-pressure-area formula $F = PA$. If the fluid pressure exerts enough force on the plug's end to lift it off the seat against the restraining force of the spring (on the left-hand side of the valve mechanism), the plug lifts and vents fluid pressure through the side port.

It is worthy to note that most relief valve mechanisms work on the exact same principle of actuation: *the valve's plug serves as its own actuator*. The pressure difference across this plug provides all the motive force necessary to actuate the valve. This simplicity translates to a high degree of reliability, a desirable quality in any safety-related system component.

Another style of overpressure valve appears in this next photograph. Manufactured by the Groth corporation, this is a combination pressure/vacuum safety valve assembly for an underground tank, designed to vent excess pressure to atmosphere *or* introduce air to the tank in the event of excess vacuum forming inside:



Even when buried, the threat of damage to the tank from overpressure is quite real. The extremely large surface area of the tank's interior walls represents an incredible amount of force potential even with low gas pressures³³. By limiting the amount of differential gas pressure which may exist between the inside and outside of the tank, the amount of stress applied to the tank walls by gas pressure or vacuum is correspondingly limited.


Large storage tanks – whether above-ground or subterranean – are typically thin-wall for reasons of economics, and cannot withstand significant pressures or vacuums. An improperly vented storage tank may burst with only slight pressure inside, or collapse inwardly with only a slight vacuum inside. Combination pressure/vacuum safety valves such as this Groth model 1208 unit reduce the chances of either failure from happening.

Of course, an alternative solution to this problem is to continuously vent the tank with an open vent pipe at the top. If the tank is always vented to atmosphere, it cannot build up either a pressure or a vacuum inside. However, continuous venting means vapors could escape from the tank if the liquid stored inside is volatile. Escaping vapors may constitute product loss and/or negative environmental impact, being a form of *fugitive emission*. In such cases it is prudent to vent the tank

³³To illustrate, consider a (vertical) cylindrical storage tank 15 feet tall and 20 feet in diameter, with an internal gas pressure of 8 inches water column. The total force exerted radially on the walls of this tank from this very modest internal pressure would be in excess of *39000 pounds!* The force exerted by the same pressure on the tank's circular lid would exceed *13000 pounds* (6.5 tons)!


with an automatic valve such as this only when needed to prevent pressure-induced stress on the tank walls.

An illustration shows the interior construction of this safety valve:



Like the miniature Nupro relief valve previously shown, the trim of this Groth safety valve acts as its own actuator: process gas pressure directly forces the vent plug off its seat, while process gas vacuum forces the vacuum plug off its seat. The lift pressure and vacuum ratings of the Groth valve are quite low, and so no spring is used to provide restraining force to the plugs. Rather, the *weight* of the plugs themselves holds them down on their seats against the force of the process gas.

This set of illustrations shows a pressure/vacuum safety valve in both modes of operation:



In each mode, the respective disk lifts up against the force of its own weight to allow gases to flow through the valve. If a greater lift pressure (or lift *vacuum*) rating is desired, precise weights may be fixed to the top of either disk. Greater weights equate to greater pressures, following the familiar equation $P = \frac{F}{A}$, where F is the force of gravity acting on the disk and weight(s) and A is the area of the disk.


For example, suppose the disk in one of these safety valves weighs 8 pounds and has a diameter of 9 inches. The surface area for a circular disk nine inches in diameter is 63.62 square inches ($A = \pi r^2$), making the lift pressure equal to 0.126 PSI ($P = \frac{F}{A}$). Such low pressures are often expressed in units other than PSI in order to make the numbers more manageable. The lift pressure of 0.126 PSI for this safety valve might alternatively be described as 3.48 inches water column or 0.867 kPa.

A close inspection of this valve design also provides clues as to why it is technically a *safety* valve rather than a *relief* valve. Recall that the distinction between these two types of overpressure-protection valves was that a relief valve opens proportionally to the experienced overpressure, while a safety valve behaves in a “snap” action manner³⁴, opening at the lift pressure and not closing again until a (lower) re-seating pressure is achieved.

The “secret” to achieving this snap-action behavior characteristic of safety valves is to design the valve’s plug in such a way that it presents a larger surface area for the escaping process fluid to act upon once open than it does when closed. This way, less pressure is needed to hold the valve open than to initially lift it from a closed condition.

³⁴Think: a safety valve has snap action!

Examining the pressure-relief mechanism of the Groth valve design closer, we see how the plug's diameter exceeds that of the seating area, with a "lip" extending down. This wide plug, combined with the lip forms an effective surface area when the plug is lifted that is larger than that exposed to the process pressure when the plug is seated. Thus, the process fluid finds it "easier" to hold the plug open than to initially lift it off the seat. This translates into a reseating pressure that is less than the lift pressure, and a corresponding "snap action" when the valve initially lifts off the seat.



The extra area on the plug's lower surface enclosed by the lip (i.e. the holding area minus the lifting area) is sometimes referred to as a *huddling chamber*. The size of this "huddling chamber" and the length of the lip establishes the degree of hysteresis (blowdown) in the safety valve's behavior.

A certain class of overpressure valve called a *safety relief valve* is designed with an adjustable "blowdown ring" to allow variations in the huddling chamber's geometry. Essentially, the blowdown ring acts as an inner lip on the valve seat to complement the outer lip on the plug. Adjusting this inner lip farther away from the plug allows more process fluid to escape laterally without touching the plug, thereby minimizing the effect of the huddling chamber and making the valve behave as a simple relief valve with no snap-action. Adjusting the blowdown ring closer to the plug forces the escaping fluid to travel toward the plug's face before reversing direction past the outer lip, making the huddling chamber more effective and therefore providing snap-action behavior. This adjustability allows the safety relief valve to act as a simple relief valve (i.e. opening proportional to overpressure) or as a safety valve (snap action) with varying amounts of blowdown ($P_{blowdown} = P_{lift} - P_{reseat}$) as determined by the user. This blowdown ring's position is typically locked into place with a seal to discourage tampering once the valve is installed in the process.


This next photograph shows a cutaway of a safety relief valve manufactured by Crosby, mounted on a cart for instructional use at Bellingham Technical College:



The adjusting bolt marked by the letter “A” at the top of the valve determines the lift pressure setting, by adjusting the amount of pre-load on the spring. Like the Nupro and Groth valves shown previously, the Crosby valve’s plug serves as its own actuator, the actuating force being a function of differential pressure across the valve and plug/seat area ($F = PA$).

The toothed gear-like component directly left of the letter “J” is called a *guide ring*, and it functions as a blowdown adjustment. This ring forms a “lip” around the valve seat’s edge much like the lip shown in the Groth valve diagrams. If the guide ring is turned to set it at a lower position (extending further past the seat), the volume of the huddling chamber increases, thereby increasing the blowdown value (i.e. keeping the valve open longer than it would be otherwise as the pressure falls).

An interesting combination of overpressure-protection technologies sometimes seen in industry are rupture disks combined with safety valves. Placing a rupture disk before a safety valve provides the benefits of ensuring zero leakage during normal operation as well as isolating the safety valve from potentially corrosive effects of the process fluid:




Potential problems with this strategy include the possibility of accumulating vapor pressure between the rupture disk and the safety valve (thereby increasing the effective burst pressure of the disk), and also the possibility of rupture disk shards becoming lodged in the safety valve mechanism, restricting flow and/or preventing re-closure.

32.5.3 Pilot-operated safety and relief valves

While many safety and relief valves actuate by the direct action of the process fluid forcing against the valve plug mechanism, others are more sophisticated in design, relying on a secondary pressure-sensing mechanism to trigger and direct fluid pressure to the main valve assembly to actuate it. This pressure-sensing mechanism is called a *pilot*, and usually features a widely-adjustable range to give the overall valve assembly a larger variety of applications.

In a pilot-operated overpressure-protection valve, the “lift” pressure value is established by a spring adjustment in the pilot mechanism rather than by an adjustment made to the main valve mechanism. A photograph³⁵ of a pilot-operated pressure relief valve used on a liquid petroleum pipeline appears here:



The relief valve mechanism itself is the white-painted flanged valve found in the center-right region of the photograph (RV-1919). This particular relief valve happens to be a Fisher model 760 with 8-inch, ANSI 300# flanges. The actuating pilot mechanism is the small unit connected to the relief valve body via stainless-steel tubing. When this pilot senses fluid pressure in the pipeline exceeding the lift pressure, it switches fluid pressure to the piston actuating mechanism of the main relief valve, opening it to relieve fluid pressure from the pipeline. Thus, the lift pressure value for the relief valve is set within the pilot rather than within the main valve mechanism. Altering this lift pressure setting is a matter of adjusting spring tension within the pilot mechanism, and/or replacing components within the pilot mechanism.

³⁵This photograph courtesy of the National Transportation Safety Board’s report of the 1999 petroleum pipeline rupture in Bellingham, Washington. Improper setting of this relief valve pilot played a role in the pipeline rupture, the result of which was nearly a quarter-million gallons of gasoline spilling into a creek and subsequently igniting. One of the lessons to take from this event is the importance of proper instrument maintenance and configuration, and how such technical details concerning industrial components may have consequences reaching far beyond the industrial facility where those components are located.

32.6 Safety Instrumented Functions and Systems

A *Safety Instrumented Function*, or *SIF*, is one or more components designed to execute a specific safety-related task in the event of a specific dangerous condition. The over-temperature shutdown switch inside a clothes dryer or an electric water heater is a simple, domestic example of an SIF, shutting off the source of energy to the appliance in the event of a detected over-temperature condition. Safety Instrumented Functions are alternatively referred to as *Instrument Protective Functions*, or *IPFs*.

A *Safety Instrumented System*, or *SIS*, is a collection of SIFs designed to bring an industrial process to a safe condition in the event of any dangerous detected conditions. Also known as *Emergency Shutdown* (ESD) or *Protective Instrument Systems* (PIS), these systems serve as an additional “layer” of protection against process equipment damage, adverse environmental impact, and/or human injury beyond the protection normally offered by a properly operating regulatory control system. Like all automatic control systems, an SIS consists of three basic sections: (1) Sensor(s) to detect a dangerous condition, (2) Controller to decide when to shut down the process, and (3) Final control element(s) to actually perform the shutdown action necessary to bring the process to a safe condition. Sensors may consist of process switches and/or transmitters separate from the regulatory control system. The controller for an SIS is usually called a *logic solver*, and is also separate from the regular control system. The final control elements for an SIS may be special on/off valves (often called “chopper” valves) or override solenoids used to force the normal control valve into a shutdown state.

Some industries, such as chemical processing and nuclear power, have extensively employed safety instrumented systems for many decades. Likewise, automatic shutdown controls have been standard on steam boilers and combustion furnaces for years. The increasing capability of modern instrumentation, coupled with the realization of enormous costs (both social and fiscal) resulting from industrial disasters has pushed safety instrumentation to new levels of sophistication and new breadths of application. It is the purpose of this section to explore some common safety instrumented system concepts as well as some specific industrial applications.

One of the challenges inherent to safety instrumented system design is to balance the goal of maximum safety against the goal of maximum economy. If an industrial manufacturing facility is equipped with enough sensors and layered safety shutdown systems to virtually ensure no unsafe condition will ever prevail, that same facility will be plagued by “false alarm” and “spurious trip” events³⁶ where the safety systems malfunction in a manner detrimental to the profitable operation of the facility. In other words, a process system designed with an emphasis on automatic shut-down will probably shut down more frequently than it actually needs to. While the avoidance of unsafe process conditions is obviously a noble goal, it cannot come at the expense of economically practical operation or else there will be no reason for the facility to exist at all³⁷. A safety system must fulfill


³⁶Many synonyms exist to describe the action of a safety system needlessly shutting down a process. The term “nuisance trip” is often (aptly) used to describe such events. Another (more charitable) label is “fail-to-safe,” meaning the failure brings the process to a safe condition, as opposed to a dangerous condition.

³⁷Of course, there do exist industrial facilities operating at a financial loss for the greater public benefit (e.g. certain waste processing operations), but these are the exception rather than the rule. It is obviously the point of a *business* to turn a profit, and so the vast majority of industries simply cannot sustain a philosophy of safety at *any* cost. One could argue that a “paranoid” safety system even at a waste processing plant is unsustainable, because too many “false trips” result in inefficient processing of the waste, posing a greater public health threat the longer it remains unprocessed.

its intended protective function, but not at the expense of compromising the intended purpose of the facility.

This tension is understood well within the electric power generation and distribution industries. Faults in high-voltage electrical lines can be very dangerous, as well as destructive to electrical equipment. For this reason, special protective devices are placed within power systems to monitor conditions and halt the flow of electricity if those conditions become threatening. However, the very presence of these devices means it is possible for power to accidentally shut off, causing unnecessary power outages for customers. In the electrical industry, the word “dependability” refers to the probability that the protective systems will cut power when required. By contrast, the word “security” is used in the electrical industry to refer to the avoidance of unnecessary outages. We will apply these terms to general process systems.

To illustrate the tension between dependability and security in a fluid process system, we may analyze a double-block shutoff valve³⁸ system for a petroleum pipeline:



The safety function of these block valves is, of course, to shut off flow from the petroleum source to the distribution pipeline in the event that the pipeline suffers a leak or rupture. Having two block valves in “series” adds an additional layer of safety, in that only one of the block valves need shut to fulfill the safety (dependability) function. Note the use of two different valve actuator technologies: one electric (motor) and the other a piston (either pneumatic or hydraulically actuated). This diversity of actuator technologies helps avoid common-cause failures, helping to ensure both valves will not simultaneously fail due to a single cause.

However, the typical operation of the pipeline demands both block valves be open in order for petroleum to flow through it. The presence of redundant (dual) block valves, while increasing safety, decreases security for the pipeline. If *either* of the two block valves happened to fail shut when there was no need to shut off the pipeline, flow through the pipeline would needlessly halt. Having two series-plumbed block valves instead of one block valve increases the probability of unnecessary pipeline shutdowns.

³⁸As drawn, these valves happen to be ball-design, the first actuated by an electric motor and the second actuated by a pneumatic piston. As is often the case with redundant instruments, an effort is made to diversify the technology applied to the redundant elements in order to minimize the probability of common-cause failures. If both block valves were electrically actuated, a failure of the electric power supply would disable both valves. If both block valves were pneumatically actuated, a failure of the compressed air supply would disable both valves. The use of one electric valve and one pneumatic valve grants greater independence of operation to the double-block valve system.


A precise notation useful for specifying dependability and security in redundant systems compares the number of redundant elements necessary to achieve the desired result compared to the total number of redundant elements. If the desired result for our double-block valve array is to shut down the pipeline in the event of a detected leak or rupture, we would say the system is *one out of two* (1oo2) redundant for dependability. In other words, only one out of the two redundant valves needs to function properly (shut off) in order to bring the pipeline to a safe condition. If the desired result is to allow flow through the pipeline when the pipeline is leak-free, we would say the system is *two out of two* (2oo2) redundant for security. This means *both* of the two block valves need to function properly (open up) in order to allow petroleum to flow through the pipeline.

This numerical notation showing the number of essential elements versus number of total elements is often referred to as *MooN* (“*M* out of *N*”) notation, or sometimes as *NooM* (“*N* out of *M*”) notation³⁹. When discussing safety instrumented systems, the ISA standard 84 defines redundancy in terms of the number of agreeing channels necessary to perform the safety (shutdown) function – in other words, the ISA’s usage of “MooN” notation implies dependability, rather than security.

A complementary method of quantifying dependability and security for redundant systems is to label in terms of how many element failures the system may sustain while still achieving the desired result. For this series set of double block valves, the safety (shutdown) function has a *fault tolerance* of one (1), since one of the valves may fail to shut when called upon but the other valve remains sufficient in itself to shut off the flow of petroleum to the pipeline. The normal operation of the system, however, has a fault tolerance of zero (0). Both block valves must open up when called upon in order to establish flow through the pipeline.

³⁹For what it’s worth, the ISA safety standard 84 defines this notation as “MooN,” but I have seen sufficient examples of the contrary (“NooM”) to question the authority of either label.


It should be clearly evident that a series set of block valves emphasizes dependability (the ability to shut off flow through the pipeline when needed) at the expense of security (the ability to allow normal flow through the pipeline when there is no leak). We may now analyze a parallel block valve scheme to compare its redundant characteristics:



In this system, the safety (dependability) redundancy function is 2oo2, since *both* block valves would have to shut off in order to bring the pipeline to a safe condition in the event of a detected pipeline leak. However, security would be 1oo2, since only one of the two valves would have to open up in order to establish flow through the pipeline. Thus, a parallel block valve array emphasizes production (the ability to allow flow through the pipeline) over safety (the ability to shut off flow through the pipeline).

Another way to express the redundant behavior of the parallel block valve array is to say that the safety function has a fault tolerance of zero (0), while the production function has a fault tolerance of one (1).

One way to avoid compromises between dependability and security is to increase the number of redundant components, forming arrays of greater complexity. Consider this quadruple block valve array, designed to serve the same function on a petroleum pipeline:



In order to fulfill its safety function of shutting off the flow of petroleum to the pipeline, both parallel pipe “branches” must be shut off. At first, this might seem to indicate a two-out-of-four (2oo4) dependability, because all we would need is for one valve in each branch (two valves total) out of the four valves to shut off in order to shut off flow to the pipeline. We must remember, however, that we do not have the luxury of assuming idealized faults. If only two of the four valves function properly in shutting off, they just might happen to be two valves *in the same branch*, in which case two valves properly functioning is not enough to guarantee a safe pipeline condition. Thus, this redundant system actually exhibits *three-out-of-four* (3oo4) dependability (i.e. it has a safety fault tolerance of one), because we need three out of the four block valves to properly shut off in order to *guarantee* a safe pipeline condition.

Analyzing this quadruple block valve array for security, we see that three out of the four valves need to function properly (open up) in order to guarantee flow to the pipeline. Once again, it may appear at first as though all we need are two of the four valves to open up in order to establish flow to the pipeline, but this will not be enough if those two valves happen to be in different parallel branches. So, this system exhibits three-out-of-four (3oo4) security (i.e. it has an production fault tolerance of one).

32.6.1 SIS sensors

Perhaps the simplest form of sensor providing process information for a safety instrumented function is a *process switch*. Examples of process switches include temperature switches, pressure switches, level switches, and flow switches⁴⁰. SIS sensors must be properly calibrated and configured to indicate the presence of a dangerous condition. They must be separate and distinct from the sensors used for regulatory control, in order to ensure a level of safety protection beyond that of the basic process control system.


Referring to the clothes dryer and domestic water heater over-temperature shutdown switches, these high-temperature shutdown sensors are distinctly separate from the regulatory (temperature-controlling) sensors used to maintain the appliance's temperature at setpoint. As such, they should only ever spring into action in the event of a high-temperature *failure* of the basic control system. That is, the over-temperature safety switch on a clothes dryer or a water heater should only ever reach its high-temperature limit if the normal temperature control system of the appliance fails to do its job of regulating temperature to normal levels.

Industrial Safety Instrumented Systems (SIS) always use dedicated transmitters and/or process switches to detect abnormal process conditions. As a rule, one should always use independent sensors for safety shutdown, and never rely on the regulatory control sensor(s) for safety functions. In the electric power industry we see this same segregation of functions: separate instrument transformers (PTs and CTs) are used to sense line voltage and line current for metering and control (regulatory) versus for protective relay (safety shutdown) equipment. It would be foolish to depend on one sensor for both functions. We see this general rule applied even in home appliances such as electric water heaters: the safety shutdown temperature switch is a separate component from the thermostat switch used to regulate water temperature. This way, a failure in the regulatory sensor does not compromise the integrity of the safety function.

A modern trend in safety instrumented systems is to use continuous process transmitters rather than discrete process switches to detect dangerous process conditions. Any process transmitter – analog or digital – may be used as a safety shutdown sensor if its signal is compared against a “trip” limit value by a comparator relay or function block. This comparator function provides an on-or-off (discrete) output based on the transmitter's signal value relative to the trip point.

⁴⁰For a general introduction to process switches, refer to chapter 9 beginning on page 653.

A simplified example of a continuous transmitter used as a discrete alarm and trip device is shown here, where analog comparators generate discrete “trip” and “alarm” signals based on the measured value of liquid in a vessel. Note the necessity of *two* level switches on the other side of the vessel to perform the same dual alarm and trip functions:



Benefits to using a continuous transmitter instead of discrete switches include the ability to easily change the alarm or trip value, and better diagnostic capability. The latter point is not as obvious as the former, and deserves more explanation. A transmitter continuously measuring liquid level will produce an output signal that varies over time with the measured process variable. A “healthy” transmitter should therefore exhibit a continuously changing output signal, proportional to the degree of change in the process. Discrete process switches, in contrast to transmitters, provide no indication of “healthy” operation. The only time a process switch should ever change states is when its trip limit is reached, which in the case of a safety shutdown sensor indicates a dangerous (rare) condition. A process switch showing a “normal” process variable may indeed be functional and indicating properly, but it might also be failed and incapable of registering a dangerous condition should one arise – there is no way to tell by monitoring its un-changing status. The continuously varying output of a process transmitter therefore serves as an indicator⁴¹ of proper function.

⁴¹Of course, the presence of some variation in a transmitter’s output over time is no guarantee of proper operation. Some failures may cause a transmitter to output a randomly “walking” signal when in fact it is not registering the process at all. However, being able to measure the continuous output of a process transmitter provides the instrument technician with far more data than is available with a discrete process switch. A safety transmitter’s output signal may be correlated against the output signal of another transmitter measuring the same process variable, perhaps even the transmitter used in the regulatory control loop. If two transmitters measuring the same process variable agree closely with one another over time, chances are extremely good are both functioning properly.


In applications where Safety Instrumented Function (SIF) reliability is paramount, *redundant* transmitters may be installed to yield additional reliability. The following photograph shows triple-redundant transmitters measuring liquid flow by sensing differential pressure dropped across an orifice plate:



A single orifice plate develops the pressure drop, with the three differential pressure transmitters “tubed” in parallel with each other, all the “high” side ports connected together through common⁴² impulse tubing and all the “low” side ports connected together through common impulse tubing. These particular transmitters happen to be FOUNDATION Fieldbus rather than 4-20 mA analog electronic. The yellow instrument tray cable (ITC) used to connect each transmitter to a segment coupling device may be clearly seen in this photograph.

⁴²It should be noted that the use of a single orifice plate and of common (parallel-connected) impulse lines represents a point of common-cause failure. A blockage at one or more of the orifice plate ports, or a closure of a manual block valve, would disable all three transmitters. As such, this might not be the best method of achieving high flow-measurement reliability.


The “trick” to using redundant transmitters is to have the system self-determine what the actual process value is in the event one or more of the redundant transmitters disagree with each other. *Voting* is the name given to this important function, and it often takes the form of signal selector functions:



Multiple selection criteria are typically offered by “voting” modules, including *high*, *low*, *average*, and *median*. A “high” select voter would be suitable for applications where the dangerous condition is a large measured value, the voting module selecting the highest-valued transmitter signal in an effort to err on the side of safety. This would represent a 1oo3 safety redundancy (since only one transmitter out of the three would have to register beyond the high trip level in order to initiate the shutdown). A “low” select voter would, of course, be suitable for any application where the dangerous condition is a small measured value (once again providing a 1oo3 safety redundancy).

The “average” selection function merely calculates and outputs the mathematical average of all transmitter signals – a strategy prone to problems if one of the redundant transmitters happens to fail in the “safe” direction (thus skewing the average value away from the “dangerous” direction and thereby possibly causing the system to respond to an actual dangerous condition later than it should).

The *median select* criterion is very useful in safety systems because it effectively ignores any measurements deviating substantially from the others. Median selector functions may be constructed of high- and low-select function blocks in either of the following⁴³ manners:



Three transmitters filtered through a median select function effectively provide a 2oo3 safety redundancy, since just a single transmitter registering a value beyond the safety trip point would be ignored by the voting function. *Two* or more transmitters would have to register values past the trip point in order to initiate a shutdown.

It should be stressed that redundant transmitter strategies are only effective if the transmitters all sense the exact same process variable, and if their failure modes are independent (i.e. no common-cause failure modes exist). If, for example, a set of redundant transmitters are attached to the

⁴³The best way to prove to yourself the median-selecting abilities of both function block networks is to perform a series of "thought experiments" where you declare three arbitrary transmitter signal values, then follow through the selection functions until you reach the output. For any three signal values you might choose, the result should always be the same: the *median* signal value is the one chosen by the voter.


process at different points such that they may legitimately sense different measurement values, the effectiveness of their redundancy will be compromised. Similarly, if a set of redundant transmitters are susceptible to failure from a shared condition (e.g. multiple liquid level transmitters that may be fooled by changes in process fluid density), then reliability will suffer.

32.6.2 SIS controllers (logic solvers)

Control hardware for safety instrumented functions should be separate from the control hardware used to regulate the process, if only for the simple reason that the SIF exists to bring the process to a safe state in the event of any unsafe condition arising, including dangerous failure of the basic regulatory controls. If a single piece of control hardware served the dual purposes of regulation *and* shutdown, a failure within that hardware resulting in loss of regulation (normal control) would not be protected because the safety function would be disabled by the same fault.


Safety controls are usually discrete with regard to their output signals. When a process needs to be shut down for safety reasons, the steps to implement the shutdown often take the form of opening and closing certain valves fully rather than partially. This sort of all-or-nothing control action is most easily implemented in the form of discrete signals triggering solenoid valves or electric motor actuators. A digital controller specially designed for and tasked with the execution of safety instrumented functions is usually called a *logic solver*, or sometimes a *safety PLC*, in recognition of this discrete-output nature.

A photograph of a “safety PLC” used as an SIS in an oil refinery processing unit is shown here, the controller being a Siemens “Quadlog” model:



Some logic solvers such as the Siemens Quadlog are adaptations of standard control systems (in the case of the Quadlog, its standard counterpart is called APACS). In the United States, where Rockwell's Allen-Bradley line of programmable logic controllers holds the dominant share of the PLC market, a version of the ControlLogix 5000 series called *GuardLogix* is manufactured specifically for safety system applications. Not only are there differences in hardware between standard and safety controllers (e.g. redundant processors), but some of the programming instructions are unique to these safety-oriented controllers as well.

An example of a safety-specific programming instruction is the GuardLogix DCSRT instruction, which compares two redundant input channels for agreement before activating a “start” bit which may be used to start some equipment function such as an electric motor:



In this case, the DCSRT instruction looks for two discrete inputs to be in the correct complementary states (Channel A = 1 and Channel B = 0) before allowing a motor to start. These states must not conflict for a time-span longer than 50 milliseconds, or else the DCSRT instruction will set a “Fault Present” (FP) bit. As you can see, the form-C pushbutton contacts are wired to two discrete inputs on the GuardLogix PLC, giving the PLC dual (complementary) indication of the switch status.

For specialized and highly critical applications, dedicated safety controllers exist which share no

legacy with standard control platforms. Triconex and ICS-Triplex are two such manufacturers, producing *triple-modular redundant* (TMR) control systems implementing 2oo3 voting at the hardware level, with redundant signal conditioning I/O circuits, redundant processors, and redundant communication channels between all components. The nuclear power industry boasts a wide array of application-specific digital control systems, with triple (or greater!) component redundancy for extreme reliability. An example of this is Toshiba's TOSMAP system for boiling-water nuclear power reactors, the digital controller and electro-hydraulic steam turbine valve actuator subsystem having a stated MTBF⁴⁴ of over 1000 years!


32.6.3 SIS final control elements

When a dangerous condition in a volatile process is sensed by process transmitters (or process switches), triggering a shutdown response from the logic solver, the final control elements must move with decisive and swift action. Such positive response may be obtained from a standard regulatory control valve (such as a globe-type throttling valve), but for more critical applications a rotary ball or plug valve may be more suitable. If the valve in question is used for safety shutdown purposes only and not regulation, it is often referred to as a *chopper* valve for its ability to “chop” (shut off quickly and securely) the process fluid flow. A more formal term for this is an *Emergency Isolation Valve*, or *EIV*.

Some process applications may tolerate the over-loading of both control and safety functions in a single valve, using the valve to regulate fluid flow during normal operation and fully stroke (either open or closed depending on the application) during a shutdown condition. A common method of achieving this dual functionality is to install a solenoid valve in-line with the actuating air pressure line, such that the valve's normal pneumatic signal may be interrupted at any moment, immediately driving the valve to a fail-safe position at the command of a discrete “trip” signal.

⁴⁴*MTBF* stands for *Mean Time Between Failure*, and represents the reliability of a large collection of components or systems. For any large batch of identical components or systems constantly subjected to ordinary stresses, MTBF is the theoretical length of time it will take for 63.2% of them to fail based on ordinary failure rates within the lifetime of those components or systems. Thus, MTBF may be thought of as the “time constant” (τ) for failure within a batch of identical components or systems.


Such a “trip” solenoid (sometimes referred to as a *dump* solenoid, because it “dumps” all air pressure stored in the actuating mechanism) is shown here, connected to a fail-closed (air-to-open) control valve:



Compressed air passes through the solenoid valve from the I/P transducer to the valve’s pneumatic diaphragm actuator when energized, the letter “E” and arrow showing this path in the diagram. When de-energized, the solenoid valve blocks air pressure coming from the I/P and vents all air pressure from the valve’s actuating diaphragm as shown by the letter “D” and arrow. Venting all actuating air pressure from a fail-closed valve will cause the valve to fail closed, obviously.

If we wished to have the valve fail open on demand, we could use the exact same solenoid and instrument air plumbing, but swap the fail-closed control valve for a fail-open control valve. When energized (regular operation), the solenoid would pass variable air pressure from the I/P transducer to the valve actuator so it could serve its regulating purpose. When de-energized, the solenoid would force the valve to the fully-open position by “dumping” all air pressure from the actuator.

For applications where it is safer to lock the control valve in its last position than to have it fail either fully closed or fully open, we might elect to use a solenoid valve in a different manner:



Here, de-energization of the solenoid valve causes the I/P transducer’s air pressure output to vent, while trapping and holding all air pressure inside the actuator at the trip time. Regardless of the valve’s “natural” fail-safe state, this system forces the valve to lock position⁴⁵ until the solenoid is re-energized.

⁴⁵This is assuming, of course, that there are no air leaks anywhere in the actuator, tubing, or solenoid which would cause the trapped pressure to decrease over time.

An example of a trip solenoid installed on a control valve appears in the following photograph. This valve also happens to have a *hand jack* wheel installed in the actuating mechanism, allowing a human operator to manually override the valve position by forcing it closed (or open) when the hand wheel is turned sufficiently:



Of all the components of a Safety Instrumented System (SIS), the final control elements (valves) are generally the least reliable, contributing most towards the system's probability of failure on demand (PFD). Sensors generally come in at second place in their contribution toward unreliability, and logic solvers a distant third place. Redundancy may be applied to control elements by creating valve networks where the failure of a single valve does not cause the system as a whole to fail. Unfortunately, this approach is extremely expensive, as valves have both high capital and high maintenance costs compared to SIS sensors and logic solvers.

A less expensive approach than redundancy to increasing safety valve reliability is to perform regular proof tests of their operation. This is commonly referred to in the industry as *partial stroke testing*. Rather than proof-test each safety valve to its full travel, which would interrupt normal process operations, the valve is commanded to move only part of its full travel. If the valve responds well to this "partial stroke" test, there is a high probability that it is able to move all the way, thus fulfilling the basic requirements of a proof test without actually shutting the process down⁴⁶.

⁴⁶Of course, if there is opportunity to fully stroke the safety valve to the point of process shutdown without undue interruption to production, this is the superior way of performing valve proof tests. Such "test-to-shutdown" proof testing may be scheduled at a time convenient to operations personnel, such as at the beginning of a planned process shutdown.

32.6.4 Safety Integrity Levels

A common way of ranking the dependability of a Safety Instrumented Function (SIF) is to use a simple numerical scale from one to four, with four being extremely dependable and one being only moderately dependable:

SIL number	Required Safety Availability (RSA)	Probability of Failure on Demand (PFD)
1	90% to 99%	0.1 to 0.01
2	99% to 99.9%	0.01 to 0.001
3	99.9% to 99.99%	0.001 to 0.0001
4	99.99% to 99.999%	0.0001 to 0.00001

The Required Safety Availability (RSA) value is synonymous with *dependability*: the probability⁴⁷ that a Safety Instrumented Function will perform its duty when faced with a dangerous process condition. Conversely, the Probability of Failure on Demand (PFD) is synonymous with *undependability*: the mathematical complement of RSA ($PFD = 1 - RSA$), expressing the probability that the SIF will fail to perform as needed, when needed.

Conveniently, the SIL number matches the minimum number of “nines” in the Required Safety Availability (RSA) value. For instance, a safety instrumented function with a Probability of Failure on Demand (PFD) of 0.00073, will have an RSA value of 99.927%, which equates to a SIL 3 rating.

It is important to understand what SIL is, and what SIL is not. The SIL rating refers to the reliability of a safety *function*, not to individual components of a system nor to the entire process itself. An overpressure protection system on a chemical reactor process with a SIL rating of 2, for example, has a Probability of Failure on Demand between 0.01 and 0.001 *for the specific shutdown function as a whole*. This PFD value incorporates failure probabilities of the sensor(s), logic solver, final control element(s), and the process piping including the reactor vessel itself plus any relief valves and other auxiliary equipment. If there arises a need to improve the PFD of this reactor’s overpressure protection, safety engineers have a variety of options at their disposal for doing so. The safety instruments themselves might be upgraded, a different redundancy strategy implemented, preventive maintenance schedules increased in frequency, or even process equipment changed to make an overpressure event less likely.

SIL ratings do not apply to an entire process. It is quite possible that the chemical reactor mentioned in the previous paragraph with an overpressure protection system SIL rating of 3 might have an *overtemperature* protection system SIL rating of only 2, due to differences in how the two different safety systems function.

Adding to this confusion is the fact that many instrument manufacturers rate their products as approved for use in certain SIL-rated applications. It is easy to misunderstand these claims, thinking that a safety instrumented function will be rated at some SIL value simply because instruments rated for that SIL value are used to implement it. In reality, the SIL value of any safety function is a much more complex determination. It is possible, for instance, to purchase and install a pressure transmitter rated for use in SIL 2 applications, and have the safety function as a whole be less than

⁴⁷ *Probability* is a quantitative measure of a particular outcome’s likelihood. A probability value of 1, or 100%, means the outcome in question is certain to happen. A probability value of 0 (0%) means the outcome is impossible. A probability value of 0.3 (30%) means it will happen an average of three times out of ten.

99% reliable (PFD greater than 0.01, or a SIL level no greater than 1) due to the effect of *Lusser's Law*⁴⁸.

As with so many other complex calculations in instrumentation engineering, there exist software packages with all the necessary formulae pre-programmed for engineers and technicians alike to use for calculating SIL ratings of safety instrumented functions. These software tools not only factor in the inherent reliability ratings of different system components, but also correct for preventive maintenance schedules and proof testing intervals so the user may determine the proper maintenance attention required to achieve a given SIL rating.

⁴⁸Lusser's Law of Reliability states that the total reliability of a system dependent on the function of several independent components is the mathematical product of those components' individual reliabilities. For example, a system with three essential components, each of those components having an individual reliability value of 70%, will exhibit a reliability of only 34.3% because $0.7 \times 0.7 \times 0.7 = 0.343$. This is why a safety function may utilize a pressure transmitter rated for use in SIL-3 applications, but exhibit a much lower total SIL rating due to the use of an ordinary final control element.

32.6.5 SIS example: burner management systems

One “classic” example of an industrial automatic shutdown system is a *Burner Management System* (or *BMS*) designed to monitor the operation of a combustion burner and shut off the fuel supply in the event of a dangerous condition. Sometimes referred to as *flame safety systems*, these systems watch for such potentially dangerous conditions as *low fuel pressure*, *high fuel pressure*, and *loss of flame*. Other dangerous conditions related to the process being heated (such as *low water level* for a steam boiler) may be included as additional trip conditions.

The safety shutdown action of a burner management system is to halt the flow of fuel to the burner in the event of any hazardous detected condition. The final control element is therefore one or more shutoff valves (and sometimes a vent valve in addition) to positively stop fuel flow to the burner.

A typical ultraviolet flame sensor appears in this photograph:



This flame sensor is sensitive to ultraviolet light only, not to visible or infrared light. The reason for this specific sensitivity is to ensure the sensor will not be “fooled” by the visible or infrared glow of hot surfaces inside the firebox if ever the flame goes out unexpectedly. Since ultraviolet light is emitted *only* by an active gas-fueled flame, the sensor acts as a true flame detector, and not a heat detector.

One of the more popular models of fuel gas safety shutoff valve used in the United States for burner management systems is shown here, manufactured by Maxon:



This particular model of shutoff valve has a viewing window on it where a metal tag linked to the valve mechanism marked “Open” (in red) or “Shut” (in black) positively indicates the valve’s mechanical status. Like most safety shutoff valves on burner systems, this valve is electrically actuated, and will automatically close by spring tension in the event of a power loss.


Another safety shutoff valve, this one manufactured by ITT, is shown here:



Close inspection of the nameplate on this ITT safety valve reveals several important details. Like the Maxon safety valve, it is electrically actuated, with a “holding” current indicated as 0.14 amps at 120 volts AC. Inside the valve is an “auxiliary” switch designed to actuate when the valve has mechanically reached the full “open” position. An additional switch, labeled *valve seal overtravel interlock*, indicates when the valve has securely reached the full “shut” position. This “valve seal” switch generates a *proof of closure* signal used in burner management systems to verify a safe shutdown condition of the fuel line. Both switches are rated to carry 15 amps of current at 120

VAC, which is important when designing the electrical details of the system to ensure the switch will not be tasked with too much current.


A simple P&ID for a gas-fired combustion burner system is shown here. The piping and valving shown is typical for a single burner. Multiple-burner systems are often equipped with individual shutoff valve manifolds and individual fuel pressure limit switches. Each burner, if multiple exist in the same furnace, *must* be equipped with its own flame sensor:



Note the use of double-block and bleed shutdown valves to positively isolate the fuel gas supply from the burner in the event of an emergency shutdown. The two block valves are specially designed for the purpose (such as the Maxon and ITT safety valves previously shown), while the bleed valve is often nothing more than an ordinary electric solenoid valve.

Most burner management systems are charged with a dual role: both to manage the safe shutdown of a burner in the event of a hazardous condition, *and* the safe start-up of a burner in normal conditions. Start-up of a large industrial burner system usually includes a lengthy *purge time* prior to ignition where the combustion air damper is left wide-open and the blower running for several minutes to positively purge the firebox of any residual fuel vapors. After the purge time, the burner management system will ignite the burner (or sometimes ignite a smaller burner called the *pilot*, which in turn will light the main burner). A burner management system executes all these pre-ignition and timing functions to ensure the burners will ignite safely and without incident.


While many industrial burners are managed by electromechanical relay or analog electronic control systems, the modern trend is toward microprocessor-based digital electronic controls. One popular system is the Honeywell 7800 series burner control system, an example of which is shown in this photograph:



Microprocessor controls provide numerous advantages over relay-based and analog electronic burner management systems. Timing of purge cycles is far more accurate with microprocessor control, and the requisite purge time is more difficult to override⁴⁹. Microprocessor-based burner controls usually have digital networking capability as well, allowing the connection of multiple controls to a single computer for remote monitoring.

⁴⁹Yes, maintenance and operations personnel alike are often tempted to bypass the purge time of a burner management system out of impatience and a desire to resume production. I have personally witnessed this in action, performed by an electrician with a screwdriver and a “jumper” wire, overriding the timing function of a flame safety system during a troubleshooting exercise simply to get the job done faster. The electrician’s rationale was that since the burner system was having problems lighting, and had been repeatedly purged in prior attempts, the purge cycle did not have to be full-length in subsequent attempts. I asked him if he would feel comfortable repeating those same words in court as part of the investigation of why the furnace exploded. He didn’t think this was funny.

The Honeywell 7800 series additionally offers local “annunciator” modules to visually indicate the status of permissive (interlock) contacts, showing maintenance personnel which switches are closed and what state the burner control system is in:



The entire “gas train” piping system for a dual-fuel boiler at a wastewater treatment facility appears in the following photograph. Note the use of double-block and bleed valves on both “trains” (one for utility-supplied natural gas and the other for “sludge gas” produced by the facility’s anaerobic digesters), the block valves for each train happening to be of different manufacture. A Honeywell 7800 flame safety control system is located in the blue enclosure:




32.6.6 SIS example: water treatment oxygen purge system

One of the processes of municipal wastewater treatment is the aerobic digestion of organic matter by bacteria. This process emulates one of many waste-decomposition processes in nature, performed on an accelerated time frame for the needs of large wastewater volumes in cities. The process consists of supplying naturally occurring bacteria within the wastewater with enough oxygen to metabolize the organic waste matter, which to the bacteria is food. In some treatment facilities, this aeration is performed with ambient air. In other facilities, it is performed with nearly pure oxygen.

Aerobic decomposition is usually part of a larger process called *activated sludge*, whereby the effluent from the decomposition process is separated into solids (sludge) and liquid (supernatant), with a large fraction of the sludge recycled back to the aerobic chamber to sustain a healthy culture of bacteria and also ensure adequate retention time for decomposition to occur. Separating liquids from solids and recycling the solids ensures a short retention time for the liquid (allowing high processing rates) and a long retention time for the solids (ensuring thorough digestion of organic matter by the bacteria).

A simplified P&ID of an activated sludge water treatment system is shown here, showing how both the oxygen flow into the aeration chamber and the sludge recycle flow back to the aeration chamber are controlled as a function of influent wastewater flow:




Aerobic decomposition performed with ambient air as the oxidizer is a very simple and safe process. Pure oxygen may be chosen instead of ambient air because it accelerates the metabolism of the bacteria, allowing more processing flow capacity in less physical space. For the same reason that pure oxygen accelerates bacterial metabolism, it also accelerates combustion of any flammable substances. This means if ever a flammable vapor or liquid were to enter the aeration chamber, there would be a risk of explosion.

Although flammable liquids are not a normal component of municipal wastewater, it is possible for flammable liquids to find their way to the wastewater treatment plant. One possibility is the event of a fuel carrier vehicle spilling its cargo, with gasoline or some other volatile fuel draining into a sewer system tunnel through holes in a grate. Such an occurrence is not normal, but certainly possible. Furthermore, it may occur without warning for the operations personnel to take preemptive action at the wastewater treatment plant.

To decrease this safety hazard, *Low Explosive Limit* (LEL) sensors installed on the aeration chamber detect and signal the presence of flammable gases or vapors inside the chamber. If any of the sensors register the presence of flammable substances, a safety shutdown system purges the chamber of pure oxygen by taking the following steps:

- Stop the flow of pure oxygen into the aeration chamber
- Open large vent valves to atmosphere
- Start air blowers to purge the chamber of residual pure oxygen



As with the P&ID, this diagram is a simplified representation of the real safety shutdown system. In a real system, multiple analytical high-alarm (LEL) sensors work to detect the presence of flammable gases or vapors, and the oxygen block valve arrangement would most likely be a double block and bleed rather than a single block valve.

The following photograph shows an LEL sensor mounted inside an insulated enclosure for protection from cold weather conditions at a wastewater treatment facility:



In this photograph, we see a purge air blower used to sweep the aeration chamber of pure oxygen (replacing it with ambient air) during an emergency shutdown condition:




Since this is a centrifugal blower, providing no seal against air flow through it when stopped, an automatic purge valve located downstream (not to be confused with the manually-actuated vent valve seen in this photograph) is installed to block off the blower from the oxygen-filled chamber. This purge valve remains shut during normal operation, and opens only after the blower has started to initiate a purge.

32.6.7 SIS example: nuclear reactor scram controls

Nuclear fission is a process by which the nuclei of specific types of atoms (most notably uranium-235 and plutonium-239) undergo spontaneous disintegration upon the absorption of an extra neutron, with the release of significant thermal energy and additional neutrons. A quantity of fissile material subjected to a source of neutron particle radiation will begin to fission, releasing massive quantities of heat which may then be used to boil water into steam and drive steam turbine engines to generate electricity. The “chain reaction” of neutrons splitting fissile atoms, which then eject more neutrons to split more fissile atoms, is inherently exponential in nature, but may be regulated by natural and artificial feedback loops.

A simplified diagram of a pressurized⁵⁰ water reactor (PWR) appears here:



⁵⁰Boiling-water reactors (BWR), the other major design type in the United States, output saturated steam at the top rather than heated water. Control rods enter a BWR from the bottom of the pressure vessel, rather than from the top as is standard for PWRs.


In the United States of America, nuclear reactors are designed to exhibit what is called a *negative temperature coefficient*, which means the chain reaction naturally slows as the temperature of the coolant increases. This physical tendency, engineered by the configuration of the reactor core and the design of the coolant system, adds a measure of self-stabilization to what would otherwise be an inherently unstable (“runaway”) process. This is an example of a “natural” negative-feedback loop in action: a process by which the very laws of physics conspire to regulate the activity of the fission reaction.

Additional regulation ability comes from the insertion of special *control rods* into the reactor core, designed to absorb neutrons and prevent them from “splitting” more atoms. With enough control rods inserted into a reactor core, a chain reaction cannot self-sustain. With enough control rods withdrawn from a freshly-fueled reactor core, the chain reaction will grow to an intensity strong enough to damage the reactor. Control rod position thus constitutes the primary method of power control for a fission reactor, and also the first⁵¹ means of emergency shutdown. These control rods are inserted and withdrawn in order to exert demand-control over the fission reaction. If the reaction rate is too low to meet demand, either a human operator or an automatic control system may withdraw the rods until the desired reactivity is reached. If the reaction rate becomes excessive, the rods may be inserted until the rate falls down to the desired level. Control rods are therefore the final control element (FCE) of an “artificial” negative-feedback loop designed to regulate reaction rate at a level matching power demand.

Due to the intense radiation flux near an operating power reactor, these control rods must be manipulated remotely rather than by direct human actuation. Nuclear reactor control rod actuators are typically special electric motors developed for this critical application.

⁵¹Other means of reactor shutdown exist, such as the purposeful injection of “neutron poisons” into the coolant system which act as neutron-absorbing control rods on a molecular level. The insertion of “scram” rods into the reactor, though, is by far the *fastest* method for quenching the chain-reaction.

A photograph⁵² showing the control rod array at the top of the ill-fated reactor at Three-Mile Island nuclear power plant appears here, with a mass of control cables connecting the rod actuators to the reactor control system:



Rapid insertion of control rods into a reactor core for emergency shutdown purposes is called a *scram*. Accounts vary as to the origin of this term, whether it has meaning as a technical acronym or as a colloquial expression to evacuate an area. Regardless of its etymology, a “scram” is an event to be avoided if possible. Like all industrial processes, a nuclear reactor fulfills its intended purpose only when operating. Shutdowns represent not only loss of revenue for the operating company, but also loss of power to local utilities and possible disruption of critical public services (heating, cooling, water pumping, fire protection, traffic control, etc.). An emergency shutdown system at a nuclear power plant must fulfill the opposing roles of dependability and security, with an extremely high degree of instrument reliability.

The electric motor actuators intended for normal operation of control rods are generally too slow to use for scram purposes. Hydraulic actuators capable of overriding the electric motor actuation may be used for scram insertion. Some early pressurized-water reactor scram system designs used a simple mechanical latch, disengaging the control rods from their motor actuators and letting gravity draw the rods fully into the reactor core.

⁵²This appears courtesy of the Nuclear Regulatory Commission’s special inquiry group report following the accident at Three Mile Island, on page 159.


A partial list of criteria sufficient to initiate a reactor scram is shown here:

- Detected earthquake
- Reactor pressure high
- Reactor pressure low
- Reactor water level low (BWR only)
- Reactor differential temperature high
- Main steam isolation valve shut
- Detected high radioactivity in coolant loop
- Detected high radioactivity in containment building
- Manual shutdown switch(es)
- Control system power loss
- Core neutron flux high
- Core neutron flux rate-of-change (period) high

The last two criteria bear further explanation. Since each fission event (the “splitting” of one fuel atom’s nucleus by an absorbed neutron) results in a definite amount of thermal energy release and also a definite number of additional neutrons released, the number of neutrons detected in the reactor core at any given moment is an approximate indication of the core’s thermal power as well as its reactivity. Neutron radiation flux measurement is therefore a fundamental process variable for fission reactor control, and also for safety shutdown. If sensors detect an excessive neutron flux, the reactor should be “scrammed” to avoid damage due to overheating. Likewise, if sensors detect a neutron flux level that is *rising* at an excessive *rate*, it indicates the possibility of a runaway chain-reaction which should also initiate a reactor “scram.”

In keeping with the high level of reliability and emphasis on safety for nuclear reactor shutdown controls, a common redundant strategy for sensors and logic is *two-out-of-four*, or *2oo4*. A contact logic diagram showing a 2oo4 configuration appears here:

2oo4 redundant logic for reactor scram systems



Any two contacts (A, B, C, or D) opening will interrupt power flow and "scram" the reactor

32.7 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Activation energy:** the amount of energy necessary to initiate a chemical reaction. Relevant to minimum ignition energy (MIE) for explosive mixtures of fuel and oxidizer, and also to intrinsic safety (preventing enough energy from reaching the process energy to possibly ignite a hazardous atmosphere).
- **Common-cause failures:** when multiple functions in a system depend on a single element, failure of that element will cause all dependent functions to fail. Relevant to design of safety shutdown systems and reliability calculations.
- **Defense-in-Depth:** a design philosophy relying on multiple layers of protection, the goal being to maintain some degree of protection in the event of one or more other layers failing.

References

Adamski, Robert S., *Design Critical Control or Emergency Shut Down Systems for Safety AND Reliability*, Revision 2, Premier Consulting Services, Irvine, CA.

Andrew, William G., *Applied Instrumentation in the Process Industries*, Volume I, Second Edition, Gulf Publishing Company, Houston, TX, 1979.

ANSI/ISA-84.00.01-2004 Part 1 (IEC 61151-1 Mod), "Functional Safety: Safety Instrumented Systems for the Process Industry Sector – Part 1: Framework, Definitions, System, Hardware and Software Requirements", ISA, Research Triangle Park, NC, 2004.

ANSI/ISA-84.00.01-2004 Part 2 (IEC 61151-2 Mod), "Functional Safety: Safety Instrumented Systems for the Process Industry Sector – Part 2: Guidelines for the Application of ANSI/ISA-84.00.01-2004 Part 1 (IEC 61151-1 Mod)", ISA, Research Triangle Park, NC, 2004.

Bazovsky, Igor, *Reliability Theory and Practice*, Prentice-Hall, Inc., Englewood Cliffs, NJ, 1961.

da Silva Cardoso, Gabriel; de Lima, Marcelo Lopes; dos Santos da Rocha, Maria Celia; Ferreira Lemos, Solange Soares, "Safety Instrumented Systems standardization for Fluid Catalytic Cracking Units at PETROBRAS", ISA, presented at ISA EXPO 2005, Chicago, IL, 2005.

"Engineer's Guide", Pepperl+Fuchs.

"Failure Mode / Mechanism Distributions" (FMD-97), Reliability Analysis Center, Rome, NY, 1997.

Grebe, John and Goble, William, *Failure Modes, Effects and Diagnostic Analysis; Project: 3051C Pressure Transmitter*, Report number Ros 03/10-11 R100, exida.com L.L.C., 2003.

"GuardLogix Safety Application Instruction Set", Publication 1756-RM095D-EN-P, Rockwell Automation, Inc., Milwaukee, WI, 2009.

Hattwig, Martin, and Steen, Henrikus, *Handbook of Explosion Prevention and Protection*, Wiley-VCH Verlag GmbH & Co. KGaA, Weinheim, Germany, 2004.

Hellemans, Marc *The Safety Relief Valve Handbook, Design and Use of Process Safety Valves to ASME and International Codes and Standards*, Elsevier Ltd, Oxford, UK, 2009.

Hicks, Tyler G., *Standard Handbook of Engineering Calculations*, McGraw-Hill Book Company, New York, NY, 1972.

"Identification and Description of Instrumentation, Control, Safety, and Information Systems and Components Implemented in Nuclear Power Plants", EPRI, Palo Alto, CA: 2001. 1001503.

"IEC 61508 Frequently Asked Questions", Rosemount website <http://mw4rosemount.usinternet.com/solution/faq61508.html>, updated December 1, 2003.

- IEEE PSRC, WG I 19, “Redundancy Considerations for Protective Relaying Schemes”, version 1.0, IEEE, 2007.
- Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Measurement and Analysis Volume I*, Fourth Edition, CRC Press, New York, NY, 2003.
- Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Control Volume II*, Third Edition, CRC Press, Boca Raton, FL, 1999.
- Lipták, Béla G. et al., *Instrument Engineers’ Handbook – Process Software and Digital Networks*, Third Edition, CRC Press, New York, NY, 2002.
- “Modern Instrumentation and Control for Nuclear Power Plants: A Guidebook”, Technical Reports Series No. 387, International Atomic Energy Agency (IAEA), Vienna, 2009.
- Newnham, Roger and Chau, Paul, *Safety Controls and Burner Management Systems (BMS) on Direct-Fired Multiple Burner Heaters*, Born Heaters Canada Ltd.
- “NFPA 70”, National Electrical Code, 2008 Edition, National Fire Protection Association.
- “NIOSH Pocket Guide to Chemical Hazards”, DHHS (NIOSH) publication # 2005-149, Department of Health and Human Services (DHHS), Centers for Disease Control and Prevention (CDC), National Institute for Occupational Safety and Health (NIOSH), Cincinnati, OH, September 2005.
- Perrow, Charles, *Normal Accidents: living with high-risk technologies*, Princeton University Press, Princeton, NJ, 1999.
- Rogovin, Mitchell and Frampton, George T. Jr., *Three Mile Island Volume I, A Report to the Commissioners and to the Public*, Nuclear Regulatory Commission Special Inquiry Group, Washington DC, 1980.
- Schultz, M. A., *Control of Nuclear Reactors and Power Plants*, McGraw-Hill Book Company, New York, NY, 1955.
- Showers, Glenn M., “Preventive Maintenance for Burner-Management Systems”, *HPAC – Heating/Piping/Air Conditioning Engineering*, February 2000.
- Svacina, Bob, and Larson, Brad, *Understanding Hazardous Area Sensing*, TURCK, Inc., Minneapolis, MN, 2001.
- “The SPEC 200 Concept”, Technical Information document TI 200-100, The Foxboro Company, Foxboro, MA, 1972.
- Ward, S.; Dahlin, T.; Higinbotham, W.; “Improving Reliability for Power System Protection”, paper presented before the 58th annual Protective Relay Conference in Atlanta, GA, April 28-30, 2004.
- Wehrs, Dave, “Detection of Plugged Impulse Lines Using Statistical Process Monitoring Technology”, Emerson Process Management, Rosemount Inc., Chanhassen, MN, December 2006.

Chapter 33

Instrumentation cyber-security

As digital technology finds greater application in industrial measurement and control systems, these systems become subject to digital vulnerabilities. Cyber-security, which used to be strictly limited to information technology (IT) systems such as those used in office and research environments (e.g. desktop computers, printers, internet routers), is now a pressing concern for industrial measurement and control systems.

There exist many points of commonality between digital IT and digital control systems, and it is at these points where mature protection concepts may be borrowed from the world of IT for use protecting industrial control systems. However, digital measurement and control systems have many unique features, and it is here we must develop protection strategies crafted specifically for industrial applications.

The chief difference between industrial controls and IT systems is, of course, the fact that industrial controls directly manage real physical processes. The purpose of an IT system, in contrast, is to manage *information*. While information can be dangerous in the wrong hands, physical processes such as chemical plants, nuclear power stations, water treatment facilities, hazardous waste treatment facilities, can be even more so.

This chapter will primarily focus on digital security as it applies to industrial measurement and control systems. The opening section is a case study on what has become a famous example of an industrial-scale cyber-attack: the so-called *Stuxnet* virus.

It should be noted that cyber-security is a very complex topic, and that this chapter of the book is quite unfinished at the time of this writing (2016). Later versions of the book will likely have much more information on this important topic.

33.1 Stuxnet

In November of 2007 a new computer virus was submitted to a virus scanning service. The purpose of this new virus was not understood at the time, but it was later determined to be an early version of the so-called *Stuxnet* virus which was designed to infiltrate and attack programmable logic controllers (PLCs) installed at the uranium enrichment facility in Iran, a critical part of that country's nuclear program located in the city of Natanz. Stuxnet stands as the world's first known computer virus ever designed to specifically attack an industrial control platform, in this case Siemens model S7 PLCs.

Later forensic analysis revealed the complexity and scope of Stuxnet for what it was: a digital weapon, directed against the Iranian nuclear program for the purpose of delaying that program's production of enriched uranium. Although the origins of Stuxnet are rather unique as viruses go, the lessons learned from Stuxnet help us as industrial control professionals to fortify our own control systems against similarly-styled digital attacks. The next such attack may not come from a nation-state like Stuxnet did, but you can be sure whoever attacks next will have gained from the lessons Stuxnet taught the world.

Since the Stuxnet attack was directed against a nuclear facility, it is worthwhile to know a little about what that facility did and how it functioned. The next subsection will delve into some of the details of modern uranium enrichment processes, while further subsections will outline how Stuxnet attacked those physical processes through the PLC control system.

The sections following this one on Stuxnet will broaden the scope of the conversation to vulnerabilities and fortification strategies common to many industrial control networks and systems.

33.1.1 A primer on uranium enrichment

Uranium is a naturally occurring metal with interesting properties lending themselves to applications of nuclear power and nuclear weaponry. Uranium is extremely dense, and also (mildly) radioactive. Of greater importance, though, is that some of the naturally occurring isotopes¹ of uranium are *fissile*, which means those atoms may be easily “split” by neutron particle bombardment, releasing huge amounts of energy as well as more neutrons which may then go on to split more uranium atoms in what is called a *chain reaction*. Such a chain-reaction, when controlled, constitutes the energy source of a fission reactor. Nuclear weapons employ violently uncontrolled chain reactions.

The most fissile isotope of uranium is uranium 235, that number being the total count of protons and neutrons within the nucleus of each atom. Unfortunately (or fortunately, depending on your view of nuclear fission), ^{235}U constitutes only 0.7% of all uranium found in the earth’s crust. The vast majority of naturally occurring uranium is the isotope ^{238}U which has all the same chemical properties of ^{235}U but is non-fissile (i.e. an atom of ^{238}U will not be “split” by neutron particle bombardment²).

Naturally-occurring uranium at a concentration of only 0.7% ^{235}U is too “dilute” for most³ nuclear reactors to use as fuel, and certainly is not concentrated enough to construct a nuclear weapon. Most power reactors require uranium fuel at a ^{235}U concentration of at least 3% for practical operation, and a concentration of at least 20% is considered the low threshold for use in constructing a uranium-based nuclear weapon. Mildly concentrated uranium useful for reactor fuel is commonly referred to “low-enriched uranium” or *LEU*, while uranium concentrated enough to build a nuclear weapon is referred to as “highly enriched uranium” or *HEU*. Modern uranium-based nuclear bombs rely on the uranium being concentrated to at least 90% ^{235}U , as do military power reactors such as the extremely compact designs used to power nuclear submarines. All of this means that an industrial-scale process for concentrating (enriching) ^{235}U is a necessary condition for building and sustaining a nuclear program of any kind, whether its purpose be civilian (power generation, research) or military (weapons, nuclear-powered vehicles).


Different technologies currently exist for uranium enrichment, and more are being developed. The technical details of uranium enrichment set the background for the Stuxnet story, the site of this cyber-attack being the Natanz uranium enrichment facility located in the middle-eastern nation of Iran.

¹The term *isotope* refers to differences in atomic mass for any chemical element. For example, the most common isotope of the element *carbon* (C) has six neutrons and six protons within each carbon atom nucleus, giving that isotope an atomic mass of twelve (^{12}C). A carbon atom having two more neutrons in its nucleus would be an example of the isotope ^{14}C , which just happens to be radioactive: the nucleus is unstable, and will over time *decay*, emitting energy and particles and in the process change into another element.

²It is noteworthy that ^{238}U can be converted into a different, fissile element called plutonium through the process of neutron bombardment. Likewise, naturally-occurring thorium 232 (^{232}Th) may be converted into ^{233}U which is fissile. However, converting non-fissile uranium into fissile plutonium, or converting non-fissile thorium into fissile uranium, requires intense neutron bombardment at a scale only seen within the core of a nuclear reactor running on some other fuel such as ^{235}U , which makes ^{235}U the critical ingredient for any independent nuclear program.

³Power reactors using “heavy” water as the moderator (such as the Canadian “CANDU” design) are in fact able to use uranium at natural ^{235}U concentration levels as fuel, but most of the power reactors in the world do not employ this design.

Like all 2-phase separation processes, uranium enrichment breaks a single input “feed” stream into two out-going streams of differing composition. Since in the case of uranium enrichment only one stream is of strategic interest, the stream containing concentrated ^{235}U is called the *product*. The other stream coming exiting the separation process, having been largely depleted of valuable ^{235}U , is called the *tails*:




During the United States’ Manhattan Project of World War Two, the main process chosen to enrich uranium for the first atomic weapons and industrial-scale reactors was *gaseous diffusion*. In this process, the uranium metal is first chemically converted into *uranium hexafluoride* (UF_6) gas so that it may be compressed, transported through pipes, processed in vessels, and controlled with valves. Then, the UF_6 gas is run through a long series of diffusion membranes (similar to fine-pore filters). At each membrane, those UF_6 molecules containing ^{235}U atoms will preferentially cross through the membranes because they are slightly less massive than the UF_6 molecules containing ^{238}U atoms. The mass difference between the two isotopes of uranium is so slight, though, that this membrane diffusion process must be repeated thousands of time in order to achieve any significant degree of enrichment. Gaseous diffusion is therefore an extremely inefficient process, but nevertheless one which may be scaled up to industrial size and used to enrich uranium at a pace sufficient for a military nuclear program. At the time of its construction, the world’s first gaseous diffusion enrichment plant (built in Oak Ridge, Tennessee) also happened to be the world’s largest industrial building.

An alternative uranium enrichment technology considered but later abandoned by the Manhattan Project scientists was *gas centrifuge* separation. A gas centrifuge is a machine with a hollow rotor spun at extremely high speed. Gas is introduced into the interior of the rotor, where centrifugal force causes the heavier molecules to migrate toward the walls of the rotor while keeping the lighter molecules toward the center. Centrifuges are commonly used for separating a variety of different liquids and solids dissolved in liquid (e.g. separating cells from plasma in blood, separating water from cream in milk), but gas centrifuges face a much more challenging task because the difference in density between various gas molecules is typically far less than the density differential in most liquid mixtures. This is especially true when the gas in question is uranium hexafluoride (UF_6), and the only difference in mass between the UF_6 molecules is that caused by the miniscule⁴ difference in mass between the uranium isotopes ^{235}U and ^{238}U .

⁴The formula weight for UF_6 containing fissile ^{235}U is 349 grams per mole, while the formula weight for UF_6 containing non-fissile ^{238}U is only slightly higher: 352 grams per mole. Thus, the difference in mass between the two molecules is less than one percent.

Gas centrifuge development was continued in Germany, and then later within the Soviet Union. The head of the Soviet gas centrifuge effort – a captured Austrian scientist named Gernot Zippe – was eventually brought to the United States where he shared the refined centrifuge design with American scientists and engineers. As complex as this technology is, it is far⁵ more energy-efficient than gas diffusion, making it the uranium enrichment technology of choice at the time of this writing (2016).

An illustration of Gernot Zippe's design is shown below. The unenriched UF_6 feed gas is introduced into the middle of the spinning rotor where it circulates in "counter-current" fashion both directions parallel to the rotor's axis. Lighter (^{235}U) gas tends to stay near the center of the rotor and is collected at the bottom by a stationary "scoop" tube where the inner gas current turns outward. Heavier (^{238}U) gas tends to stay near the rotor wall and is collected at the top by another stationary "scoop" where the outer current turns inward:

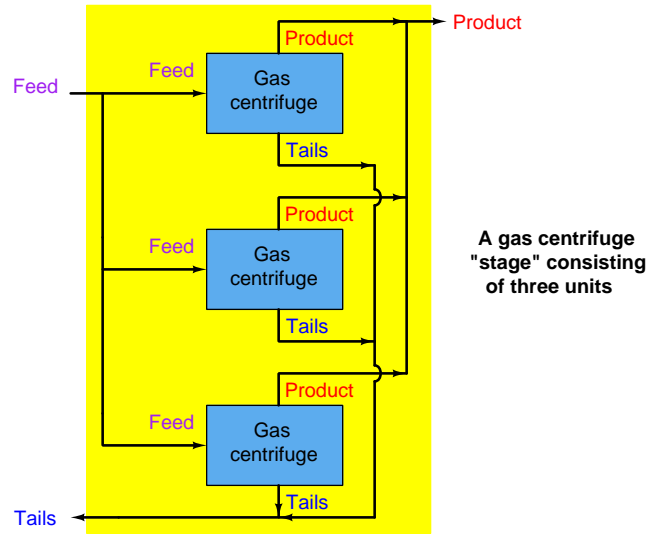


Like the separation membranes used in gaseous diffusion processes, each gas centrifuge is only able to enrich the UF_6 gas by a very slight amount. The modest enrichment factor of each centrifuge necessitates many be connected in series, with each successive centrifuge taking in the out-flow of the previous centrifuge in order to achieve any practical degree of enrichment. Furthermore, gas

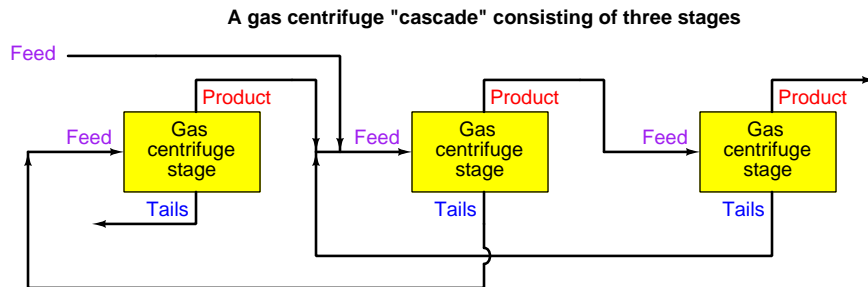
⁵By some estimates, gas centrifuge enrichment is 40 to 50 *times* more energy efficient than gaseous diffusion enrichment.

centrifuges are by their very nature rather limited in their flow capacity⁶. This low “throughput” necessitates parallel-connected gas centrifuges in order to achieve practical production rates for a national-scale nuclear program. A set of centrifuges connected in parallel for higher flow rates is called a *stage*, while a set of centrifuge stages connected in series for greater enrichment levels is called a *cascade*.

A gas centrifuge *stage* is very simple to understand, as each centrifuge’s feed, product, and tails lines are simply paralleled for additional throughput:




A gas centrifuge *cascade* is a bit more complex to grasp, as each centrifuge’s product gets sent to the feed inlet of the next stage for further enrichment, and the tails gets sent to the feed inlet of the previous stage for further depletion. The main feed line enters the cascade at one of the middle stages, with the main product line located at one far end and the main tails line located at the other far end:




⁶A typical gas centrifuge’s mass flow rating is on the order of milligrams per second. At their very low (vacuum) operating pressures, a typical centrifuge rotor will hold only a few grams of gas at any moment in time.

This US Department of Energy (DOE) photograph shows an array of 1980's-era American gas centrifuges located in Piketon, Ohio. Each of the tall cylinders is a single gas centrifuge machine, with the feed, product and tails tubing seen connecting to the spinning rotor at the top of the stationary casing:



The size of each stage in a gas centrifuge cascade is proportional to its feed flow rate. The stage processing the highest feed rate must be the largest (i.e. contain the most centrifuges), while the stages at the far ends of the cascade contain the least centrifuges. A cascade similar to the one at the Natanz enrichment facility in Iran – the target of the Stuxnet cyber-attack – is shown here without piping for simplicity, consisting of 164 individual gas centrifuges arranged in 15 stages. The main feed enters in the middle of the cascade at the largest stage, while enriched product exits at the right-hand end and depleted tails at the left-hand end:



The sheer number of gas centrifuges employed at a large-scale uranium enrichment facility is quite staggering. At the Natanz facility, where just one cascade contained 164 centrifuges, cascades were paralleled together in *sub-units* of six cascades each (984 centrifuges per sub-unit), and three of these sub-units made one cascade *unit* (2952 centrifuges total).

33.1.2 Gas centrifuge vulnerabilities

It would be an understatement to say that a gas centrifuge is a delicate machine. In order to perform their task efficiently⁷, gas centrifuge rotors must be long and made to rotate at extremely high rates of speed. Maintaining any rotating machine in a state of near-perfect balance is difficult, much more so when the rotating element is very long⁸. Furthermore, since the gas pressure inside each centrifuge rotor is sub-atmospheric, leak-free seals must be maintained between the spinning rotor and the stationary components (the casing and internal tubing). The extremely high rotational speeds of modern gas centrifuges (many tens of thousands of revolutions per minute!) necessitate advanced materials be used in rotor construction, optimizing light weight and high strength so that the rotors will not be torn to pieces by their own centrifugal force.

A peculiar problem faced by any high-speed rotating machine is a phenomenon called *critical speed*. Any object possessing both mass and resilience is capable of oscillating, which of course includes any and every rotating machine component. If the rotating component of a machine happens to spin at a rate equal to its own natural oscillating frequency, a condition of *mechanical resonance* occurs. *Any* amount of imbalance in the rotating component while spinning at this speed, however slight, will generate a force driving the assembly into continuous oscillation. The speed at which this resonance occurs is called the “critical speed” of the machine, and it should be avoided whenever possible.

Destructive resonance will be avoided so long as the machine is maintained at any speed significantly below or above its critical speed. Most modern gas centrifuges are classified as *supercritical* machines, because they are designed to operate at rotational speeds exceeding their critical speeds. The only time resonance becomes a problem in a supercritical machine is during start-up and shut-down, when the speed must momentarily pass through the critical value. So long as this moment is brief, however, oscillations will not have enough time to grow to destructive levels.

In addition to the problems faced by all high-speed rotating machines, a problem unique to gas centrifuges is gas pressure control. Since the rotor of a gas centrifuge spins inside of an evacuated⁹ stationary casing, the existence of any gas pressure inside the rotor creates additional stress acting in the same outward direction as the rotor’s own centrifugal force. This means rotor gas pressure must be maintained at a very low level in order to minimize rotor stress. Furthermore, if pressure and temperature conditions are not carefully controlled in a gas centrifuge, the gas may actually *sublimate* into a solid state which will deposit material on the inside wall of the rotor and surely throw it out of balance.

⁷Three major factors influence the efficiency of a gas centrifuge: rotor wall speed, rotor length, and gas temperature. Of these, rotor wall speed is the most influential. Higher speeds separate isotopes more effectively, because higher wall speeds result in greater amounts of radial acceleration, which increases the amount of centrifugal force experienced by the gas molecules. Longer rotors also separate isotopes more effectively because they provide more opportunity for the counter-flowing gas streams to separate lighter molecules toward the center and heavier molecules toward the wall. Higher temperatures reduce separation efficiency, because gas molecules at higher temperatures are more mobile and therefore diffuse (i.e. mix together) at higher rates. Therefore, the optimum gas centrifuge design will be long, spin as fast as possible, and operate as cool as possible.

⁸To give you an idea of just how long some gas centrifuge rotors are, the units built for the US Department of Energy facility in Ohio used rotors *40 feet in length!*

⁹This means the hollow casing exists in a state of vacuum, with no air or other gases present. This is done in order to help thermally insulate the rotor from ambient conditions, as well as avoid generating heat from air friction against the rotor’s outside surface. Remember, elevated temperatures cause the gas to diffuse at a faster rate, which in turn causes the gas to randomly mix and therefore not separate into light and heavy isotopes as intended.

One could argue that the temperamental nature of gas centrifuges is a good thing, because it makes the manufacture of enriched uranium difficult to achieve, which in turn complicates the development of nuclear weapons. This fragility also makes gas centrifuges an ideal target for anyone interested in halting or delaying nuclear weapons development, which was precisely the aim of the Stuxnet computer virus.

33.1.3 The Natanz uranium enrichment facility

Iran used an obsolete gas centrifuge design, perhaps the best they could obtain at the time, as the uranium enrichment platform of choice for their Natanz facility. By modern standards, this design was inefficient and troublesome, but the Iranians were able to coax serviceable performance from this centrifuge design by means of extensive instrumentation and controls.

Simply put, the Iranian strategy was to manufacture centrifuges faster than they would break and equip the centrifuge cascades with enough piping and supervisory instrumentation that they could detect and isolate failed centrifuges without stopping production, rather than wait until they had perfected the design of the centrifuges themselves. The extensive network of sensors, valves, piping, and PLCs (Programmable Logic Controllers) installed at the Natanz facility facilitated this fault-tolerant design.

The key to the Natanz system's fault tolerance was a set of isolation ("block") valves installed at each gas centrifuge. Each machine was also equipped with a sufficient array of sensors to detect malfunctions. If a centrifuge experienced trouble, such as excessive vibration, the PLC control system would automatically shut all the isolation valves for that failed centrifuge and turn off its drive motor. Since most stages in each cascade contained multiple centrifuges in parallel, the isolation of a single centrifuge within a stage would not shut down the entire cascade. Instead, maintenance personnel could repair the failed centrifuge while production continued, and return it to service when ready.

One undesired consequence of shutting isolation valves on operating centrifuges, though, was increased gas pressure in portions of the cascade. With fewer centrifuges left to handle a constant feed flow, the pressure drop across that stage increases. All upstream stages therefore experience more gas pressure, which as described earlier increases the stress imparted on the spinning centrifuge rotors. In answer to this problem was another innovation at the Natanz facility: using the "dump system" (a standard feature in any gas centrifuge cascade, for evacuating gas from the centrifuges in the event of an emergency shut-down event) as a pressure relief in the event of overpressure resulting from too many isolated centrifuges. Of course, engaging this "dump" system as a means of pressure control would reduce production rates, but it was a better outcome for the system operators than a complete shut-down of the cascade.

In summary, the instrumentation employed in the Natanz facility would automatically detect problems in each centrifuge, isolate any failed centrifuges from the running cascade, and open dump valves as necessary to reduce gas pressure on the remaining centrifuges. This so-called Cascade Protection System was implemented by Siemens model S7-417 PLCs, one per sub-unit (six cascades, each sub-unit containing 984 individual gas centrifuges). All-digital *Profibus* technology was used to communicate process data over network cables between the field instruments and the PLCs, as a means of reducing what would have otherwise been a huge amount of analog and discrete signal wiring.

Additional Siemens PLCs were used at the Natanz facility to control the gas centrifuges, notably the model S7-315 employed to issue commands to variable-frequency drive units sending power to the rotor drive motors. Like the larger S7-417 PLC units, one S7-315 PLC was used to control the motor drives of each cascade sub-unit (six cascades, 984 centrifuges). As subsequent portions of this chapter will detail, both of these Siemens PLC platforms were targets of the Stuxnet virus.

33.1.4 Stuxnet version 0.5

Multiple versions of the Stuxnet virus were aimed at the Natanz facility, at least two significantly different “major” versions which are publicly known at the time of this writing (2016). The first major Stuxnet version, developed as early as November of 2005 and labeled as version 0.5 by the Symantec Corporation, differed from later versions both in its means of delivery (the *dropper* portion of the virus code) and its means of attack (the *payload* portion of the virus code). Later versions of Stuxnet (compiled in 2009-2010 and dubbed versions 1.x by Symantec) employed a much more sophisticated “dropper” and a payload designed to affect a completely different portion of the Iranian centrifuge control system.

A summary of Stuxnet version 0.5 appears here:

- **Infection point:** The infection begins with files written to a removable drive (e.g. USB flash drive), automatically run by the Windows operating system upon connection to a personal computer.
- **Dropper vector:** Stuxnet searches for and infects any Siemens Step 7 PLC project archives found on the personal computer.
- **Payload target:** Siemens S7-417 programmable logic controllers (PLCs) implementing the Cascade Protection System for isolation and overpressure control of centrifuges.
- **Payload vector:** Install a DLL (Dynamically Linked Library) file in the Siemens Step 7 software library collection designed to alter any Step 7 programming code downloaded to a PLC, inserting attack code in the infected PLCs.
- **Payload task:** Shut off isolation valves and mis-calibrate the pressure sensors to cause mild over-pressuring of the centrifuges.
- **Goal:** Increase stress on operating centrifuges, leading to premature failure. Avoid catastrophic cascade failure, which would raise suspicion.
- **Stop date:** July 4, 2009.

The “dropper” portion of Stuxnet version 0.5 exploited a vulnerability in the Siemens “Step 7” PLC programming software which runs on Windows-based personal computers, but did not exploit any vulnerabilities within the Windows operating system itself. In fact, this early version of Stuxnet lacked the ability to self-propagate over the internet, and had to be installed on a personal computer running the Siemens Step 7 software. The most popular hypothesis to date is that the infection happened via a USB flash drive, or “memory stick” used to store digital data.

The “payload” portion of Stuxnet version 0.5 was incredibly sophisticated by comparison.

33.1.5 Stuxnet version 1.x

Subsequent versions of Stuxnet have been labeled as version 1.x and are treated here as one major release. A summary of Stuxnet versions 1.x appears here:

- **Infection point:** The infection begins with files written to a removable drive (e.g. USB flash drive), automatically run by the Windows operating system upon connection to a personal computer. The infection is then able to spread from one Windows PC to another over networks using multiple Windows vulnerabilities.
- **Dropper vector:** Exploit multiple “zero day¹⁰” vulnerabilities in Windows XP and Vista operating systems to aggressively propagate the virus over computer networks, then infect any Siemens Step 7 project files found on those computers.
- **Payload target:** Siemens S7-315 programmable logic controllers (PLCs) regulating centrifuge rotor speeds.
- **Payload vector:** Install a DLL (Dynamically Linked Library) file in the Siemens Step 7 software library collection designed to alter any Step 7 programming code downloaded to a PLC, inserting attack code in the infected PLCs.
- **Payload task:** Change rotor speeds over time so as to make them pass through their “critical speed” range.
- **Goal:** Increase stress on operating centrifuges, leading to premature failure. Again, avoid catastrophic cascade failure which would raise suspicion.
- **Stop date:** June 24, 2012.

33.2 Motives

There are multiple motives for compromising the security of an industrial control system, some of which overlap motives for attacking IT systems, and some of which are unique to the industrial world. This section details some of the reasons why people might wish to attack an industrial control system.

¹⁰The term *zero-day* in the digital security world refers to vulnerabilities that are unknown to the manufacturer of the software, as opposed to known vulnerabilities that have been on record with the manufacturer for some time. The fact that Stuxnet 1.x employed no less than four zero-day Windows exploits strongly suggests it was developed by an agency with highly sophisticated resources. In other words, Stuxnet 1.x wasn’t made by amateurs. This is literally world-class hacking in action!

33.2.1 Technical challenge

Computer experts tend to be a demographic of people motivated by technical challenges and problem-solving. To this type of person, the challenge of breaking in to a computer system designed to foil intruders may be too tempting to resist.

To the person interested in compromising a digital system just for the sake of seeing whether it can be done, the reward is in achieving access, not necessarily inflicting any damage. These people are generally not a direct threat, but may pose an indirect threat if they share their expertise with others harboring sinister motives.

Other individuals motivated by the technical challenge of accessing a digital system are interested in seeing just how much havoc they can wreak once they gain access. Such individuals are analogous to *digital arsonists*, interested in starting the biggest fire that they can simply for the sake of the fire's size.

33.2.2 Profit

The major motive driving IT cyber-attacks today is *profit*: the theft of credit card and other sensitive digital information which may be sold on the black market. Criminal organizations benefit from this style of digital attack, with many attackers becoming millionaires by way of their digital exploits.

Another form of profit-driven attack is commonly called *ransomware*, where an attacker inserts malicious software on the victim's computer(s) preventing access to the system or encrypting files such that they become unusable. This malware then presents a message to the victim asking for monetary payment in exchange for normal system access.

Neither of these attacks is novel to industrial systems, and in fact are commonplace in the IT world. What is novel in industrial systems is the severity of the repercussions. One might imagine the response from an oil drilling rig's management team to ransomware preventing startup-up of a new oil well, where downtime may be in the range of millions of US dollars per day of production. Not only is the imperative to get back online stronger than it would be for a private individual whose home computer was being held ransom, but the ability for an oil company to immediately pay the attacker is much greater than any private individual.

Another potential application of the profit motive in industrial system attacks is *commodities trading*. Traders who profit from the purchase and sale of commodities produced by industrial manufacturers might stand to gain by knowing the day-to-day operational status of those manufacturers. If such people were to access the production and inventory logs residing in a facility's digital control system, for example, they may be able to make more profitable trading decisions based on this privileged information. Eavesdropping on industrial control system data therefore poses another mode of *insider trading*.

33.2.3 Espionage

Aside from gathering data from industrial systems for the direct purpose of profit, less direct motives for attacking industrial control systems exist. One such motive is the theft of proprietary process data, for example recipes and formulae for producing chemical products such as craft foods and drinks, as well as pharmaceuticals.

Special control strategies and process designs critical to the manufacture of certain products are valuable to competing organizations as well. A chemical company eager to discover how to control a temperamental new chemical reaction process might wish to sample the controller algorithms and instrument configurations used by a successful competitor. Even if these design details were not stolen outright, the attacker may gather valuable test data and learn from the developmental mistakes of their competitor, thereby saving time and money pursuing their own design.

Militaries also stand to gain from espionage of industrial measurement and control systems, since the military capabilities of other nations are founded on industrial-scale operations. A country interested in tracking the development of an adversary's nuclear weapons potential, for example, would have a motive to perform digital espionage via the control systems of those foreign nuclear facilities.

33.2.4 Sabotage

Here, at least in my view, is where cyber-security as it relates to industrial control systems becomes really interesting. The major factor distinguishing digital control system security from IT system security is the former's supervision of a real physical process. This means a control system cyber-attack has far more *direct potential for harm* than any IT cyber-attack.

Corporations and nation-states both have an interest in industrial sabotage if it means they may diminish the economic productivity of a competitor. A country, for example, whose export market is dominated by a single product may be tempted to launch cyber-attacks against facilities producing that same product in other countries, as a means to either maintain or elevate their power in the world economy. Corporations have the exact same interest, just at a different level within the global economy.

Certain activists may also have an interest in sabotaging an industrial facility. Shutting down production of a facility they deem dangerous or unethical, or perhaps just causing the company financial loss through poor product quality and/or non-compliance, are potential motivators for activists to target specific industrial processes.

Military interest in industrial sabotage is practically a "given" assumption, as such a cyber-attack merely constitutes a new type of weapon to add to their existing arsenals. Unlike conventional weapons, cyber-weapons are relatively inexpensive.

Another category of sabotage relevant to cyber-attacks is that perpetrated by *malicious insiders*. This last category is especially troubling, as it involves personnel with in-depth knowledge of the digital systems in question. This simple fact makes defense against such attacks extremely challenging, because these are people normally authorized to access the system and therefore are able to bypass most (if not all) security measures. A few notable examples of internal sabotage are listed here:

- Secret agents of foreign nations
- Recently discharged (former) employees
- Disgruntled employees within a corporation

The destructive potential of a government operative with access to critical systems needs no further explanation. Employees, however, do. An employee who gets laid off or fired may still have access to their former employer's critical systems if their system account is not promptly closed. The same is true if the company maintains a lax password policy, such as multiple people sharing a common user account. Even current employees may be motivated to sabotage their employer's systems, especially where there might be an economic advantage¹¹ to doing so.

¹¹Consider what forms of sabotage *striking* employees might be willing to do in order to gain leverage at the bargaining table.

33.2.5 Terrorism

This last motive is especially troubling when one considers the proliferation of digital technology and the disconcerting rise of terror-related attacks around the world. The goal of terrorists is quite simply to instill terror as a means of manipulating and/or punishing perceived enemies. Driven by ideology, terrorists tend not to discriminate when selecting their targets. Like arsonists previously mentioned, success is measured by the magnitude of terror and carnage instilled by the event. Common concerns of ethics are trumped by the dictates of the ideology.

The attacks of September 11, 2001 taught the world how ordinary technologies and systems (in that case, fully-fueled jet passenger aircraft) may be exploited as weapons capable of killing and injuring thousands of people. Industrial process designers would do well to think in similar terms, examining their systems not just from the perspective of their intended purpose but also as potential weapons wielded by terrorists.

33.3 Lexicon of cyber-security terms

Cyber-security seems to have its own vocabulary, ranging from unwieldy technical acronyms to slang terms borrowed from amateur computer enthusiasts. What follows is a partial listing of some common terms and their definitions. This list is not only useful as a definitional reference when encountering such terms in cyber-security literature, but it also serves to outline a number of common attack strategies:

- **Active attack:** an attack involving data written to a network or to device. See *passive attack* for contrast.
- **Backdoor:** an easy-to-access pathway into a system, typically used by system developers for convenience in their work. There is nothing wrong with a backdoor during development, but backdoors are very dangerous when left in place on commissioned systems.
- **Blacklist:** a database of prohibited messages or users or software applications.
- **Broadcast:** a form of network where all transmissions are heard by all connected devices, even those devices the data is not intended for. Any communication network sharing a common physical channel is a broadcast network.
- **Brute-force attack:** attempting every combination of characters in an effort to forge a working password.
- **Comsec:** shorthand for “communications security”.
- **Crypto:** shorthand for “cryptography”, which is the purposeful scrambling of data to render it unintelligible to all but the intended recipient.
- **Data diode:** a device permitting only one-way (simplex) data communication. Data diodes eliminate the possibility of active attacks, because they make writing data to the protected system impossible.
- **Denial-of-Service:** a form of attack where the intended function of the system is either downgraded or entirely faulted.
- **Dictionary attack:** attempting common words and character combinations in an effort to forge a working password.
- **Firewall:** a software or hardware application intended to limit connectivity between networked devices.
- **Flooding:** an attack technique consisting of overloading a digital system with data or requests for data, generally the point of which being to achieve denial of service (DoS) when the target system becomes overloaded.
- **FTP:** an acronym standing for File Transfer Protocol, a protocol used for reading and writing files on one computer remotely from another computer. FTP is a predecessor to *SFTP* which includes public-private key encryption for much better security.

- **HTTP:** an acronym standing for Hyper Text Transfer Protocol, the method used for computers to exchange web page data (encoded in HTML files). HTTP is not encrypted.
- **HTTPS:** an acronym standing for Hyper Text Transfer Protocol Secure, the method used for computers to exchange web page data (encoded in HTML files) using encryption.
- **IP:** an acronym standing for Internet Protocol, the packaging of data into “packets” which may be routed independently of each other across a large network.
- **IT:** an acronym standing for Information Technology, used to broadly describe general-purpose digital data systems and communications.
- **Key:** a relatively small segment of digital data that serves to either encrypt or decrypt other digital data. The imagery here is that of a key used to engage or disengage a physical lock.
- **LAN:** an acronym standing for Local Area Network, a network connecting multiple devices over a limited distance, such as the span of an office building or campus. See *WAN* for contrast.
- **Malware:** software written with malicious intent.
- **Man-in-the-Middle:** an attack where the attacker is positioned directly in between sender and receiver, in such a way as to be able to modify messages sent over the network without either sender or receiver being aware.
- **Operating system:** software installed on a computer for the purpose of directly managing that computer’s hardware resources, functioning as an intermediate layer between the application and the hardware itself. The existence of operating system software vastly simplifies the design and development of application software. Popular consumer-grade operating systems at the time of this writing (2016) include Microsoft Windows, Apple OS X, Linux, and BSD.
- **Packet sniffing:** the act of passively monitoring data transmitted over a broadcast network.
- **Passive attack:** an attack only involving the reading of data from a network or device. See *active attack* for contrast.
- **Passphrase:** an easily-memorized sentence which may be used to generate complex passwords. For example, one could take the first letter of every word in the passphrase “What we have here is a failure to communicate” to generate the password **wwhiaftc**. Passphrases are useful because they make complex passwords easy to remember, and in fact may be used to generate multiple passwords from the same phrase (e.g. using the *last* letter of each word instead of the first, to create the password **teeesaeoe** from the same passphrase used previously).
- **Ping:** a simple network utility used on IP networks to test connectivity. The ping message is sent from one computer to another, with the receiving computer replying to declare successful receipt of the ping message.
- **Private key:** a cryptographic key useful for decrypting encrypted data. “Private” refers to the fact that this key must be held in confidence by authorized parties only, since it has the ability to unlock coded messages.

- **Public key:** a cryptographic key useful only for encrypting data. “Public” refers to the fact that this key may be shared openly, as it cannot be used to unlock a coded message, but instead is only useful for encoding messages sent to a party holding a *private key* which can decode the message.
- **Replay attack:** a form of attack where a message is intercepted, recorded, and later broadcast to the network in order to inflict damage. An interesting feature of replay attacks is that they may work on encrypted messages, and even when the purpose of the message is unknown to the attacker!
- **SFTP:** an acronym standing for Secure File Transfer Protocol, a protocol used for reading and writing files on one computer remotely from another computer. SFTP is a successor to *FTP* which lacked encryption.
- **Sniffing:** inspecting network communications for important data. So-called “packet sniffers” monitor data traffic on a broadcast network for certain information such as network addresses and data content.
- **Spoofing:** presenting a false identification to the receiver of digital data. This commonly takes the form of presenting fake network address information, to trick the receiver into thinking the source is from a legitimate location or device.
- **SSH:** an acronym standing for Secure Shell, a remote-access utility commonly used in Unix operating systems allowing users to log into a computer from another computer connected to the same network. SSH is a successor to *telnet*, which lacked encryption.
- **TCP:** an acronym standing for Terminal Control Protocol, the protocol used to ensure segments of data make it to their intended destinations after being routed by IP (see *Internet Protocol*).
- **Telnet:** a remote-access utility commonly used in Unix operating systems allowing users to log into a computer from another computer connected to the same network. Telnet is a predecessor to *SSH* which includes public-private key encryption for much greater security.
- **Trusted:** a component or section of a digital system that is assumed to be safe from intrusion.
- **VPN:** an acronym standing for Virtual Private Network, which encrypts every aspect of a transaction between two computers connected on a network. The effect is to form a “virtual network” between the machines, the privacy of which being ensured by the encryption used to scramble the data.
- **Vuln:** shorthand for “vulnerability” or weakness in a system.
- **Walled garden:** a term used to describe an area of a digital system assumed to be safe from intrusion. See *trusted*.
- **WAN:** an acronym standing for Wide Area Network, a network connecting multiple devices over a long range, such as the span of a city. See *LAN* for contrast.
- **War dialing:** the exploratory practice of dialing random phone numbers in search of telephone modem connections, which may connect to computer systems.

- **Whitelist:** a database of permitted messages or users or software applications.

33.4 Fortifying strategies

As control system professionals, it is in our interest to ensure our measurement and control systems are secure from unauthorized access. It is helpful to regard system *security* similarly to how we regard system *safety* or *reliability*, as these concerns share many common properties:

- Just as accidents and faults are inevitable, so are unauthorized access attempts
- Just as 100% perfect safety and 100% perfect reliability is unattainable, so is 100% security
- Digital security needs to be an important criterion in the selection and setup of industrial instrumentation equipment, just as safety and reliability are important criteria
- Maximizing security requires a security-savvy culture within the organization, just as maximizing safety requires a safety-savvy culture and maximizing reliability requires a reliability-centric design philosophy

Also similar to safety and reliability is the philosophy of *defense-in-depth*, which is simply the idea of having multiple layers of protection in case one or more fail. Applied to digital security, defense-in-depth means not relying on a single mode of protection (e.g. passwords only) to protect a system from attack.

33.4.1 Policy-based strategies

These strategies focus on human behavior rather than system design or component selection. In some ways these strategies are the simplest to implement, as they generally require little in the way of technical expertise. This is not to suggest, however, that policy-based security strategies are therefore the *easiest* to implement. On the contrary, changing human behavior is usually a very difficult feat. Policy-based strategies are not necessarily cheap, either: although little capital is generally required, operational costs will likely rise as a result of these policies. This may take the form of monetary costs, additional staffing costs, and/or simply costs associated with impeding normal work flow (e.g. pulling personnel away from their routine tasks to do training, requiring personnel to spend more time doing things like inventing and tracking new passwords, slowing the pace of work by limiting authorization).

Foster awareness

Ensure all personnel tasked with using and maintaining the system are fully aware of security threats, and of best practices to mitigate those threats. Given the ever-evolving nature of cyber-attacks, this process of educating personnel must be continuous.

A prime mechanism of cyber-vulnerability is the casual sharing of information between employees, and with people outside the organization. Information such as passwords and network design should be considered “privileged” and should only be shared on a need-to-know basis. Critical security information such as passwords should never be communicated to others or stored electronically in plain (“cleartext”) format. When necessary to communicate or store such information electronically, it should be encrypted so that only authorized personnel may access it.

In addition to the ongoing education of technical personnel, it is important to keep management personnel aware of cyber threat and threat potentials, so that the necessary resources will be granted toward cyber-security efforts.

Employ security personnel

For any organization managing important processes and services, “important” being defined here as *threatening* if compromised by the right type of cyber-attack, it is imperative to employ qualified and diligent personnel tasked with the ongoing maintenance of digital security. These personnel must be capable of securing the control systems themselves and not just general data systems.

One of the routine tasks for these personnel should be evaluations of risks and vulnerabilities. This may take the form of security audits or even simulated attacks whereby the security of the system is tested with available tools.

Utilize effective authentication

Simply put, it is imperative to correctly identify all users accessing a system. This is what “authentication” means: correctly identifying the person (or device) attempting to use the digital system. *Passwords* are perhaps the most common authentication technique.

The first and foremost precaution to take with regard to authentication is to never use default (manufacturer) passwords, since these are public information. This precautionary measure may seem so obvious as to not require any elaboration, but sadly it remains a fact that too many password-protected devices and systems are found operating in industry with default passwords.

Another important precaution to take with passwords is to not use the same password for all systems. The reasoning behind this precaution is rather obvious: once a malicious party gains knowledge of that one password, they have access to all systems protected by it. The scenario is analogous to using the exact same key to unlock every door in the facility: all it takes now is one copied key and suddenly intruders have access to every room.

Passwords must also be changed on a regular basis. This provides some measure of protection even after a password becomes compromised, because the old password(s) no longer function.

Passwords chosen by system users should be “strong,” meaning difficult for anyone else to guess. When attackers attempt to guess passwords, they do so in two different ways:

- Try using common words or phrases that are easy to memorize
- Try every possible combination of characters until one is found that works

The first style of password attack is called a *dictionary attack*, because it relies on a database of common words and phrases. The second style of password attack is called a *brute force attack* because it relies on a simple and tireless (“brute”) algorithm, practical only if executed by a computer.

A password resistant to dictionary-style attacks is one not based on a common word or phrase. Ideally, that password will appear to be nonsense, not resembling any discernible word or simple pattern. The only way to “crack” such a password, since a database of common words will be useless against it, will be to attempt every possible character combination (i.e. a brute-force attack).

A password resistant to brute-force-style attacks is one belonging to a huge set of possible passwords. In other words, there must be a very large number of possible passwords limited to the same alphabet and number of characters. Calculating the brute-force strength of a password is a matter of applying a simple exponential function:

$$S = C^n$$

Where,

S = Password strength (i.e. the number of unique password combinations possible)

C = Number of available characters (i.e. the size of the alphabet)

n = Number of characters in the password

For example, a password consisting of four characters, each character being a letter of the English alphabet where lower- and upper-case characters are treated identically, would give the following strength:

$$S = 26^4 = 456976 \text{ possible password combinations}$$

If we allowed case-sensitivity (i.e. lower- and upper-case letters treated differently), this would double the value of C and yield more possible passwords:

$$S = 52^4 = 7311616 \text{ possible password combinations}$$

Obviously, then, passwords using larger alphabets are stronger than passwords with smaller alphabets.

Cautiously grant authorization

While *authentication* is the process of correctly identifying the user, *authorization* is the process of assigning rights to each user. The two concepts are obviously related, but not identical. Under any robust security policy, users are given only as much access as they need to perform their jobs efficiently. Too much access not only increases the probability of an attacker being able to cause maximum harm, but also increases the probability that benevolent users may accidentally cause harm.

Perhaps the most basic implementation of this policy is for users to log in to their respective computers using the lowest-privilege account needed for the known task(s), rather than to log in at the highest level of privilege they *might* need. This is a good policy for all people to adopt when they use personal computers to do any sort of task, be it work- or leisure-related. Logging in with full (“administrator”) privileges is certainly convenient because it allows you to do anything on the system (e.g. install new software, reconfigure any service, etc.) but it also means any malware accidentally engaged¹² under that account now has the same unrestricted level of access to the system. Habitually logging in to a computer system with a low-privilege account helps mitigate this risk, for any accidental execution of malware will be similarly limited in its power to do harm.

Another implementation of this policy is called *application whitelisting*, where only trusted software applications are allowed to be executed on any computer system. This stands in contrast to “blacklisting” which is the philosophy behind anti-virus software: maintaining a list of software applications known to be harmful (malware) and prohibiting the execution of those pre-identified applications. Blacklisting (anti-virus) only protects against malware that has been identified and notified to that computer. Blacklisting cannot protect against “zero-day” malware known by no one except the attacker. In a whitelisting system, each computer is pre-loaded with a list of acceptable applications, and no other applications – benign or malicious – will be able to run on that machine.

¹²Consider the very realistic scenario of logging in as administrator (or “root” in Unix systems) and then opening an email message which happens to carry an attached file infected with malware. Any file executed by a user is by default run at that user’s level of privilege because the operating system assumes that is the user’s intent.

Maintain good documentation

While this is important for effective maintenance in general, thorough and accurate documentation is especially important for digital security because it helps identify vulnerabilities. Details to document include:

- Network diagrams
- Software version numbers
- Device addresses

Close unnecessary access pathways

All access points to the critical system must be limited to those necessary for system function. This means all other potential access points in the critical system must be closed so as to minimize the total number of access points available to attackers. Examples of access points which should be inventoried and minimized:

- Hardware communication ports (e.g. USB serial ports, Ethernet ports, wireless radio cards)
- Software TCP ports
- Shared network file storage (“network drives”)
- “Back-door” accounts used for system development

That last category deserves some further explanation. When engineers are working to develop a new system, otherwise ordinary and sensible authentication/authorizations measures become a major nuisance. The process of software development always requires repeated logins, shutdowns, and tests forcing the user to re-authenticate themselves and negotiate security controls. It is therefore understandable when engineers create simpler, easier access routes to the system under development, to expedite their work and minimize frustration.

Such “back-door” access points become a problem when those same engineers forget (or simply neglect) to remove them after the developed system is released for others to use. An interesting example of this very point was the so-called *basisk* vulnerability discovered in some Siemens S7 PLC products. A security researcher named Dillon Beresford working for NSS Labs discovered a telnet¹³ service running on certain models of Siemens S7 PLCs with a user account named “basisk” (the password for this account being the same as the user name). All one needed to do in order to gain privileged access to the PLC’s operating system was connect to the PLC using a telnet client and enter “basisk” for the user name and “basisk” for the password! Clearly, this was a back-door account used by Siemens engineers during development of that PLC product line, but it was not closed prior to releasing the PLC for general use.

¹³Telnet is a legacy software utility used to remotely access command-line computer operating systems. Inherently insecure, telnet exchanges login credentials (user name and password) unencrypted over the network connection. A modern replacement for telnet is SSH (Secure SHell).

Maintain operating system software

All operating system software manufacturers periodically release “patches” designed to improve the performance of their products. This includes patches for discovered security flaws. Therefore, it is essential for all computers belonging to a critical system to be regularly “patched” to ensure maximum resistance to attack.

This is a significant problem within industry because so much industrial control system software is built to run on consumer-grade operating systems such as Microsoft Windows. Popular operating systems are built with maximum convenience in mind, not maximum security or even maximum reliability. New features added to an operating system for the purpose of convenient access and/or new functionality often present new vulnerabilities¹⁴.

Another facet to the consumer-grade operating system problem is that these operating systems have relatively short lifespans. Driven by consumer demand for more features, software manufacturers develop new operating systems and abandon older products at a much faster rate than industrial users upgrade their control systems. Upgrading the operating systems on computers used for an industrial control system is no small feat, because it usually means disruption of that system’s function, not only in terms of the time required to install the new software but also (potentially) re-training required for employees. Upgrading may even be impossible in cases where the new operating system no longer supports features necessary for that control system¹⁵. This would not be a problem if operating system manufacturers provided the same long-term (multi-decade) support for their products as industrial hardware manufacturers typically do, but this is not the case for consumer-grade products such as Microsoft Windows¹⁶.

¹⁴I am reminded of an example from the world of “smart” mobile telephones, commonly equipped with *accelerometer* sensors for detecting physical orientation. Accelerometers detect the force of acceleration and of gravity, and are useful for a variety of convenient “apps” having nothing to do with telephony. Smart phone manufacturers include such sensors in their mobile devices and link those sensors to the phone’s operating system because doing so permits innovative applications, which in turn makes the product more desirable to application developers and ultimately consumers. It was discovered, though, that the signals generated by these accelerometers could be used to detect “keystrokes” made by the user, the sensors picking up vibrations made as the user taps their finger against the glass touch-screen of the smart phone. With the right signal processing, the accelerometers’ signals could be combined in such a way to identify which characters the user was tapping on the virtual keyboard, and thereby eavesdrop on their text-based communications!

¹⁵An example of this is where a piece of obsolete industrial software runs on the computer’s operating system, for example a data acquisition program or data-analysis program made by a company that no longer exists. If this specialized software was written to run on a particular operating system, and no others, future versions of that operating system might not permit proper function of that specialized software. I have seen such cases in industry, where industrial facilities continue to run obsolete (unsupported) operating systems in order to keep running some specialized industrial software (e.g. PLC programming editors), which is needed to operate or maintain some specialized piece of control hardware which itself is obsolete but still functions adequately for the task. In order to upgrade to a modern operating system on that computer (e.g. an obsolete version of Microsoft Windows), one must upgrade the specialized software (e.g. the PLC programming editor software), which in turn would mean upgrading the control hardware (e.g. the PLCs themselves). All of this requires time and money, much more than just what is required to upgrade the operating system software itself.

¹⁶As a case in point, there are still a great many industrial computers running Microsoft Windows XP at the time of this writing (2016), even though this operating system is no longer supported by Microsoft. This means no more Service Pack upgrades from Microsoft, security patches, or even research on vulnerabilities for this obsolete operating system. All users of Windows XP are “on their own” with regard to cyber-attacks.

Routinely archive critical data

The data input into and generated by digital control systems is a valuable commodity, and must be treated as such. Unlike material commodities, data is easily replicated, and this fact provides some measure of protection against loss from a cyber-attack. Routine “back-ups” of critical data, therefore, is an essential part of any cyber-security program. It should be noted that this includes not just operational data collected by the control system during operation, but also data such as:

- PID tuning parameters
- Control algorithms (e.g. function block programs, configuration data, etc.)
- Network configuration parameters
- Software installation files
- Software license (authorization) files
- Software drivers
- Firmware files
- User authentication files
- All system documentation (e.g. network cable diagrams, loop diagrams)

This archived data should be stored in a medium immune to cyber-attacks, such as read-only optical disks. It would be foolish, for example, to store this sort of critical data only as files on the operating drives of computers susceptible to attack along with the rest of the control system.

Create response plans

Just as no industrial facility would be safe without incident response plans to mitigate physical crises, no industrial facility using digital control systems is secure without response plans for cyber-attacks. This includes such details as:

- A chain of command for leading the response
- Instructions on how to restore critical data and system functions
- Work-arounds for minimal operation while critical systems are still unavailable

Limit mobile device access

Mobile digital devices such as cell phones and even portable storage media (e.g. USB “flash” drives) pose digital security risks because they may be exploited as an attack vector bypassing air gaps and firewalls. It should be noted that version 0.5 of Stuxnet was likely inserted into the Iranian control system in this manner, through an infected USB flash drive.

A robust digital security policy will limit or entirely prohibit personal electronic devices into areas where they might connect to the facility’s networks or equipment. Where mobile devices are essential for job functions, those devices should be owned by the organization and registered in such a way as to authenticate their use. Computers should be configured to automatically reject non-registered devices such as removable flash-memory storage drives. Portable computers not owned and controlled by the organization should be completely off-limits¹⁷ from the process control system.

Above all, one should never underestimate the potential harm allowing uncontrolled devices to connect to critical, trusted portions of an industrial control system. The degree to which any portion of a digital system may be considered “trusted” is a function of *every* component of that system. Allowing connection to untrusted devices violates the confidence of that system.

Secure all toolkits

A special security consideration for industrial control systems is the existence of software designed to create and edit controller algorithms and configurations. The type of software used to write and edit Ladder Diagram (LD) code inside of programmable logic controllers (PLCs) is a good example of this, such as the Step7 software used to program Siemens PLCs in Iran’s Natanz uranium enrichment facility. Instrumentation professionals use such software on a regular basis to do their work, and as such it is an essential tool of the trade. However, this very same software is a weapon in the hands of an attacker, or when hijacked by malicious code.

A common practice in industry is to leave computers equipped with such “toolkit” software connected to the control network for convenience. This is a poor policy, and one that is easily remedied by simply disconnecting the programming computer from the control network immediately after downloading the edited control code. An even more secure policy is to never connect such “toolkit” computers to a network at all, but only to controllers directly, so that the toolkit software cannot be hijacked.

Another layer of defense is to utilize robust password protection on the programmable control devices when available, rather than leaving password fields blank which then permits any user of the toolkit software full access to the controller’s programming.

¹⁷This raises a potential problem from the perspective of outside technical support, since such support often entails contracted or manufacturer-employed personnel entering the site and using their work computers to perform system configuration tasks. For any organization implementing a strong security access policy, this point will need to be negotiated into every service contract to ensure all the necessary pieces of hardware and software exist “in-house” for the service personnel to use while on the job.

Close abandoned accounts

Given the fact that disgruntled technical employees constitute a significant security threat to organizations, it stands to reason that the user accounts of terminated employees should be closed as quickly as possible. Not only do terminated employees possess authentication knowledge in the form of user names and passwords, but they may also possess extensive knowledge of system design and vulnerabilities.

33.4.2 Design-based strategies

A *design-based* security strategy is one rooted in technical details of system architecture and functionality. Some of these strategies are quite simple (e.g. air gaps) while others are quite complex (e.g. encryption). In either case, they are strategies ideally addressed at the inception of a new system, and at every point of system alteration or expansion.

Advanced authentication

The authentication security provided by passwords, which is the most basic and popular form of authentication at the time of this writing, may be greatly enhanced if the system is designed to not just reject incorrect passwords, but to actively inconvenience the user for entering wrong passwords.

Password timeout systems introduce a mandatory waiting period for the user if they enter an incorrect password, typically after a couple of attempts so as to allow for innocent entry errors. *Password lockout* systems completely lock a user out of their digital account if they enter multiple incorrect passwords. The user's account must then be reset by another user on that system possessing high-level privileges.

The concept behind both password timeouts and password lockouts is to greatly increase the amount of time required for any dictionary-style or brute-force password attack to be successful, and therefore deter these attacks. Unfortunately timeouts and lockouts also present another form of system vulnerability to a *denial of service* attack. Someone wishing to deny access to a particular system user need only attempt to sign in as that user, using incorrect passwords. The timeout or lockout system will then delay (or outright deny) access to the legitimate user.

Authentication based on the user's knowledge (e.g. passwords) is but one form of identification, though. Other forms of authentication exist which are based on the possession of physical items called *tokens*, as well as identification based on unique features of the user's body (e.g. retinal patterns, fingerprints) called *biometric* authentication.

Token-based authentication requires all users to carry tokens on their person. This form of authentication so long as the token does not become stolen or copied by a malicious party.

Biometric authentication enjoys the advantage of being extremely difficult to replicate and nearly¹⁸ impossible to lose. The hardware required to scan fingerprints is relatively simple and inexpensive. Retinal scanners are more complex, but not beyond the reach of organizations possessing expensive digital assets. Presumably, there will even be DNA-based authentication technology available in the future.

¹⁸Before you laugh at the idea of losing one's own body, consider something as plausible as a fingerprint scanner programmed to accept the image of just one finger, and then that user suffering an injury to that same finger either obscuring the fingerprint or destroying the finger entirely.

Air gaps

An *air gap* is precisely what the name implies: a physical separation between the critical system and any network, preventing communication. Although it seems so simple that it ought to be obvious, an important design question to ask is whether or not the system in question really needs to have connectivity at all. Certainly, the more networked the system is, the easier it will be to access useful information and perform useful operational functions. However, connectivity is also a liability: that same convenience makes it easier for attackers to gain access.

While it may seem as though air gaps are the ultimate solution to digital security, this is not entirely true. A control system that *never* connects to a network is still vulnerable to cyber-attack, and that is through detachable programming and data-storage devices. For example, a PLC without a permanent network connection may become compromised by way of an infected portable computer used to edit the PLC's code. A DCS completely isolated from the facility's local area network (LAN) may become compromised if someone plugs in an infected data storage device such as a USB "flash" memory module.

In order for air gaps to be completely effective, they must be permanent and include portable devices as well as network connections. This is where effective security policy comes into play, ensuring portable devices are not allowed into areas where they might connect (intentionally or otherwise) to critical systems. Effective air-gapping of critical networks also necessitates physical security of the network media: ensuring attackers cannot gain access to the network cables themselves, so as to "tap" into those cables and thereby gain access. This requires careful planning of cable routes and possibly extra infrastructure (e.g. separate cable trays, conduits, and access-controlled equipment rooms) to implement.

Wireless (radio) data networks pose a special problem for the "air gap" strategy, because the very purpose of radio communication is to bridge physical air gaps. A partial measure applicable to some wireless systems is to use *directional antennas* to link separated points together, as opposed to *omnidirectional* antennas which transmit and receive radio energy equally in all directions. This complicates the task of "breaking in" to the data communication channel, although it is not 100 percent effective since no directional antenna has a perfectly focused radiation pattern, nor do directional antennas preclude the possibility of an attacker intercepting communications directly between the two antennae. Like all security measures, the purpose of using directional antennas is to make an attack *less probable*.

Network segregation

This simply means to divide digital networks into separate entities (or at least into layers) in order to reduce the exposure of digital control systems to any sources of harm. Air gaps constitute an elementary form of network segregation, but are not practical when at least some data must be communicated between networks.

At the opposite end of the network segregation spectrum is a scenario where all digital devices, control systems and office computers alike, connect to the facility's common Local Area Network (LAN). This is almost always a bad policy, as it invites a host of problems not limited to cyber-attacks but extending well beyond that to innocent mistakes and routine faults which may compromise system integrity. At the very least, control systems deserve their own dedicated network(s) on which to communicate, free of traffic from general information technology (IT) office systems.

In facilities where control system data absolutely must be shared on the general LAN, a *firewall* should be used to connect those two networks. Firewalls are either software or hardware entities designed to filter data passed through based on pre-set rules. In essence, each network on either side of a firewall is a "zone" of communication, while the firewall is a "conduit" between zones allowing only certain types of messages through. A rudimentary firewall might be configured to "blacklist" any data packets carrying hyper-text transfer protocol (HTTP) messages, as a way to prevent web-based access to the system. Alternatively, a firewall might be configured to "whitelist" only data packets carrying Modbus messages for a control system and block everything else.

Some specialized firewalls are manufactured specifically for industrial control systems. One such firewall at the time of this writing (2016) is manufactured by the Tofino, and has the capability to screen data packets based on rules specific to industrial control system platforms such as popular PLC models. Industrial firewalls differ from general-purpose data firewalls in their ability to recognize control-specific data.

In systems where different communication zones must have different levels of access to the outside world, multiple firewalls may be set up in such a way as to create a so-called *demilitarized zone* (DMZ). A DMZ is a network existing between a pair of firewalls, one firewall filtering data to and from a protected network, and the other firewall filtering data to and from an unprotected (or less-protected) network. Any devices connected to the DMZ will have access to either network, through different firewall rule sets. Any data exchanged between the protected and unprotected networks, though, must pass through both firewalls.

Encryption

Encryption refers to the intentional scrambling of data by means of a designated code called a *key*, a similar (or in some cases identical) key being used to un-scramble (decrypt) that data on the receiving end. The purpose of encryption, of course, is to foil passive attacks by making the data unintelligible to anyone but the intended recipient, and also to foil active attacks by making it impossible for an attacker's transmitted message to be successfully received.

A popular form of encrypted communication used for general networking is a *Virtual Private Network* or VPN. This is where two computer systems use VPN software (or two VPN hardware devices) to encrypt messages sent to each other over an unsecure medium. Since every single packet of data exchanged between the two computers is encrypted, the communication will be unintelligible to anyone else who might be "listening" on that unsecure network. In essence, VPNs create a secure "tunnel" for data to travel between points on an otherwise unprotected network.

Encryption may also be applied to dedicated, non-broadcast networks such as telephone channels and serial data communication lines. Special cryptographic modems and serial data translators are manufactured specifically for this purpose, and may be applied to legacy SCADA and control networks based on telephony or serial communication cables.

Important files stored on computer drives may also be encrypted, such that only users possessing the proper key(s) may decrypt and use the files.

It should be noted that encryption does not necessarily protect against so-called *replay* attacks, where the attacker records a communicated message and later re-transmits that same message to the network. For example, if a control system uses an encrypted message to command a remotely-located valve to shut, an attacker might simply re-play that same message at any time in the future to force the valve to shut without having to decrypt the message.

An interesting form of encryption applicable to certain wireless (radio) data networks is *spread-spectrum* communication. This is where all the data is not communicated over the same radio carrier frequency, but rather is divided or "spread" among a range of frequencies. Various techniques exist for spreading digital data across a spectrum of radio frequencies, but they all comprise a form of data encryption because the spreading of that data is orchestrated by means of a cryptographic key. Perhaps the simplest spread-spectrum method to understand is *frequency-hopping* or *channel-hopping*, where the transmitters and receivers both switch frequencies on a keyed schedule. Any receiver uninformed by the same key will not "know" which channels will be used, or in what order, and therefore will be unable to intercept anything but isolated pieces of the communicated data. Spread-spectrum technology was invented during the second World War as a means for Allied forces to encrypt their radio transmissions such that Axis forces could not interpret them.

Spread-spectrum capability is built into several wireless data communication standards, including Bluetooth and *WirelessHART*.

Read-only system access

One way to thwart so-called “active” attacks (where the attacker inserts or modifies data in a digital system to achieve malicious ends) is to engineer the system in such a way that all communicated data is *read-only* and therefore cannot be written or edited by anyone. This, of course, by itself will do nothing to guard against “passive” (read-only) attacks such as eavesdropping, but passive attacks are definitely the lesser of the two evils with regard to industrial control systems.

In systems where the digital data is communicated serially using protocols such as EIA/TIA-232, read-only access may be ensured by simply disconnecting one of the wires in the EIA/TIA-232 cable. By disconnecting the wire leading to the “receive data” pin of the critical system’s EIA/TIA-232 serial port, that system cannot receive external data but may only transmit data. The same is true for EIA/TIA-485 serial communications where “transmit” and “receive” connection pairs are separate.

Certain serial communication schemes are inherently simplex (i.e. one-way communication) such as EIA/TIA-422. If this is an option supported by the digital system in question, the use of that option will be an easy way to ensure remote read-only access.

For communication standards such as Ethernet which are inherently duplex (bi-directional), devices called *data diodes* may be installed to ensure read-only access. The term “data diode” invokes the functionality of a semiconductor rectifying diode, which allows the passage of electric current in one direction only. Instead of blocking reverse current flow, however, a “data diode” blocks reverse *information* flow.

The principle of read-only protection applies to computing systems as well as communication networks. Some digital systems do not strictly require on-board data collection or modification of operating parameters, and in such cases it is possible to replace read/write magnetic data drives with read-only (e.g. optical disk) drives in order to create a system that cannot be compromised. Admittedly, applications of this strategy are limited, as there are few control systems which never store operational data nor require any editing of parameters. However, this strategy should be considered where it applies¹⁹.

Many digital devices offer *write-protection* features in the form of a physical switch or key-lock preventing data editing. Just as some types of removable data drives have a “write-protect” tab or switch located on them, some “smart” field instruments also have write-protect switches inside their enclosures which may be toggled only by personnel with direct physical access to the device. Programmable Logic Controllers (PLCs) often have a front-panel write-protect switch allowing protection of the running program.

¹⁹An example of this strategy in action is an internet-connected personal computer system I once commissioned, running the Linux operating system from a DVD-ROM optical disk rather than a magnetic hard drive. The system would access the optical disk upon start-up to load the operating system kernel into its RAM memory, and then access the disk as needed for application executable files, shared library files, and other data. The principal use of this system was web browsing, and my intent was to make the computer as “hacker-proof” as I possibly could. Since the operating system files were stored on a read-only optical disk, it was impossible for an attacker to modify that data without having physical access to the machine. In order to thwart attacks on the data stored in the machine’s RAM memory, I configured the system to automatically shut down and re-start every day at an hour when no one would be using it. Every time the computer re-booted, its memory would be a *tabula rasa* (“clean slate”). Of course, this meant no one could permanently store downloaded files or other data on this machine from the internet, but from a security perspective that was the very point.

Not only do write-protect switches guard against malicious attacks, but they also help prevent innocent mistakes from causing major problems in control systems. Consider the example of a PLC network where each PLC connected to a common data network has its own hardware write-protect switch. If a technician or engineer desires to edit the program in one of these PLCs from their remotely-located personal computer, that person must first go to the location of that PLC and disable its write protection. While this may be seen as an inconvenience, it ensures that the PLC programmer will not mistakenly access the wrong PLC from their office-located personal computer, which is especially easy to do if the PLCs are similarly labeled on the network.

Making regular use of such features is a policy measure, but ensuring the exclusive use of equipment with this feature is a system design measure.

Control platform diversity

In control and safety systems utilizing redundant controller platforms, an additional measure of security is to use different models of controller in the redundant array. For example, a redundant control or safety system using two-out-of-three voting (2oo3) between three controllers might use controllers manufactured by three different vendors, each of those controllers running different operating systems and programmed using different editing software. This mitigates against device-specific attacks, since no two controllers in the array should have the exact same vulnerabilities.

A less-robust approach to process control security through diverse platforms is simply the use of effective Safety Instrumented Systems (SIS) applied to critical processes, which always employ controls different from the base-layer control system. An SIS system is designed to bring the process to a safe (shut down) condition in the event that the regular control system is unable to maintain normal operating conditions. In order to avoid common-cause failures, the SIS must be implemented on a control platform independent from the regular control system. The SIS might even employ analog control technology (and/or discrete relay-based control technology) in order to give it complete immunity from digital attacks.

In either case, improving security through the use of multiple, diverse control systems is another example of the *defense in depth* philosophy in action: building the system in such a way that no essential function depends on a single layer or single element, but rather multiple layers exist to ensure that essential function.

33.5 Review of fundamental principles

Shown here is a partial listing of principles applied in the subject matter of this chapter, given for the purpose of expanding the reader's view of this chapter's concepts and of their general inter-relationships with concepts elsewhere in the book. Your abilities as a problem-solver and as a life-long learner will be greatly enhanced by mastering the applications of these principles to a wide variety of topics, the more varied the better.

- **Blacklisting:** the concept of flagging certain users, software applications, etc. as “forbidden” from accessing a system.
- **Chemical isotopes:** variants of chemical elements differing fundamentally in atomic mass. Relevant to the subject of uranium enrichment for nuclear reactors and nuclear weapons, where one particular isotope must be separated from (“enriched”) another isotope in order to be useful.
- **Defense-in-Depth:** a design philosophy relying on multiple layers of protection, the goal being to maintain some degree of protection in the event of one or more other layers failing.
- **Reliability:** a statistical measure of the probability that a system will perform its design function. Relevant here with regard to control systems, in that proper control system design can significantly enhance the reliability of a large system if the controls are able to isolate faulted redundant elements within that system. This is the strategy used by designers of the Iranian uranium enrichment facility, using PLC controls to monitor the health of many gas centrifuges used to enrich uranium, and taking failed centrifuges off-line while maintaining continuous production.
- **Whitelisting:** the concept of only permitting certain users, software applications, etc. to access a system.

References

- “21 Steps to Improve Cyber Security of SCADA Networks”, Department of Energy, USA, May 2011.
- Bartman, Tom and Carson, Kevin, “Securing Communications for SCADA and Critical Industrial Systems”, Technical Paper 6678-01, Schweitzer Engineering Laboratories, Inc., Pullman, WA, January 22, 2015.
- Beresford, Dillon, “Siemens Simatic S7 PLC Exploitation”, technical presentation at Black Hat USA conference, 2011.
- Byres, Eric, “Building Intrinsically Secure Control and Safety Systems – Using ANSI/ISA-99 Security Standards for Improved Security and Reliability”, Byres Security Inc., May 2009.
- “Common Cybersecurity Vulnerabilities in Industrial Control Systems”, Department of Homeland Security, Control Systems Security Program, National Cyber Security Division, USA, May 2011.
- Falliere, Nicolas; Murchu, Liam O.; Chien, Eric; “W32.Stuxnet Dossier”, version 1.4, Symantec Corporation, Mountain View, CA, February 11, 2011.
- Fischer, Ted, “Private and Public Key Cryptography and Ransomware”, Center for Internet Security, Inc., Pullman, WA, December 2014.
- Grennan, Mark, “Firewall and Proxy Server HOWTO”, version 0.8, February 26, 2000.
- Horak, Ray, *Webster’s New World Telecom Dictionary*, Wiley Publishing, Inc., Indianapolis, IN, 2008.
- Kemp, R. Scott, “Gas Centrifuge Theory and Development: A Review of US Programs”, Program on Science and Global Security, Princeton University, Princeton, NJ, Taylor & Francis Group, LLC, 2009.
- Langner, Ralph, “To Kill A Centrifuge – A Technical Analysis of What Stuxnet’s Creators Tried to Achieve”, The Langner Group, Arlington, MA, November 2013.
- Lee, Jin-Shyan; Su, Yu-Wei; Shen, Chung-Chou, “A Comparative Study of Wireless Protocols: Bluetooth, UWB, ZigBee, and Wi-Fi”, Industrial Technology Research Institute, Hsinchu, Taiwan, November 2007.
- Leidigh, Christopher, “Fundamental Principles of Network Security”, White Paper #101, American Power Conversion (APC), 2005.
- Leischner, Garrett and Whitehead, David, “A View Through the Hacker’s Looking Glass”, Technical Paper 6237-01, Schweitzer Engineering Laboratories, Inc., Pullman, WA, April 2006.
- Makhijani, Arjun Ph.D.; Chalmers, Lois; Smith, Brice Ph.D.; “Uranium Enrichment – Just Plain Facts to Fuel an Informed Debate on Nuclear Proliferation and Nuclear Power”, Institute for Energy and Environmental Research, October 15, 2004.

McDonald, Geoff; Murchu, Liam O.; Doherty, Stephen; Chien, Eric; “Stuxnet 0.5: The Missing Link”, version 1.0, Symantec Corporation, Mountain View, CA, February 26, 2013.

Oman, Paul W.; Risley, Allen D.; Roberts, Jeff; Schweitzer, Edmund O. III, “Attack and Defend Tools for Remotely Accessible Control and Protection Equipment in Electric Power Systems”, Schweitzer Engineering Laboratories, Inc., Pullman, WA, March 12, 2002.

Postel, John, *Internet Protocol – DARPA Internet Program Protocol Specification*, RFC 791, Information Sciences Institute, University of Southern California, Marina Del Ray, CA, September 1981.

Rescorla, E. and Korver, B.; “Guidelines for Writing RFC Text on Security Considerations” (RFC 3552), The Internet Society, July 2003.

Risley, Allen; Marlow, Chad; Oman, Paul; Dolezilek, Dave, “Securing SEL Ethernet Products With VPN Technology”, Application Guide 2002-05, Schweitzer Engineering Laboratories, Inc., Pullman, WA, July 11, 2002.

“Seven Strategies to Effectively Defend Industrial Control Systems”, National Cybersecurity and Communications Integration Center (NCCIC), Department of Homeland Security (DHS), USA.

“W32.DuQu – The Precursor to the next Stuxnet”, version 1.4, Symantec Corporation, Mountain View, CA, November 23, 2011.

Whitehead, David and Smith, Rhett, “Cryptography: A Tutorial for Power Engineers”, Technical Paper 6345-01, Schweitzer Engineering Laboratories, Inc., Pullman, WA, October 20, 2008.

Zippe, Gernot, “A Progress Report: Development of Short Bowl Centrifuges”, Department of Physics, University of Virginia, July 1, 1959.

Chapter 34

Problem-solving and diagnostic strategies

The ability to solve complex problems is the most valuable technical skill an instrumentation professional can cultivate. A great many tasks associated with instrumentation work may be broken down into simple step-by-step instructions that any marginally qualified person may perform, but effective problem-solving is different. Problem-solving requires creativity, attention to detail, and the ability to approach a problem from multiple mental perspectives.

“Problem-solving” often refers to the solution of abstract problems, such as “word” problems in a mathematics class. However, in the field of industrial instrumentation it most often finds application in the form of “troubleshooting:” the diagnosis and correction of problems in instrumented systems. Troubleshooting is really just a form of problem-solving, applied to real physical systems rather than abstract scenarios. As such, many of the techniques developed to solve abstract problems work well in diagnosing real system problems. As we will see in this chapter, problem-solving in general and troubleshooting in particular are closely related to *scientific method*, where hypotheses are proposed, tested, and modified in the quest to discern cause and effect.

Like all skills, problem-solving may be improved with practice and persistence. The goal of this chapter is to outline several problem-solving tools and techniques.

34.1 Learn principles, not procedures

“As to methods there may be a million and then some, but principles are few. The man who grasps principles can successfully select his own methods. The man who tries methods, ignoring principles, is sure to have trouble.” – **Ralph Waldo Emerson**

Effective problem-solvers always reason from fundamental principles, rather than follow memorized procedures. Following this logic, the wise strategy for any student learning how to solve any type of problem is to internalize as many general principles as possible, and to connect common threads of principle to different applications and scenarios.

A good way for any student to self-check that their focus is indeed on principles rather than procedures is to continually ask themselves if they are able to explain *why* what they have learned is true. It is one thing to be able to explain what you have learned, but it is quite another to explain *why* your new knowledge is valid. “Why” may be the most important question you ever ask, and you need to ask it often.

34.2 Active reading

Learning from reading printed text is a kind of problem-solving activity in its own right. The problem is how to acquire new information from the pages of a book, and the solution requires active engagement of your mind as you read. This is often more difficult than it may seem at first, especially when the subject matter is complex and/or the source text is poorly written. Given the fact that much more technical information is available in text form than in any other format, and also the fact that continuous learning is absolutely essential in the field of instrumentation, active reading is an indispensable tool for student success as well as for continued professional growth in the field of instrumentation.

34.2.1 Don't limit yourself to one text!

A very common mistake made by struggling students is to limit themselves to the reading assigned by their instructor(s), when better books might be available. If homework assignments are given from a particular assigned text, it is understandable why a student might think this is the *only* appropriate text to use. However, textbook selection is an imperfect process, often influenced by factors other than optimal learning (e.g. instructor bias, publisher influence, etc.). Sadly, textbooks are sometimes chosen not on their merits as a learning resource but rather by whether or not the textbook comes complete with pre-made exercises and examinations to be used in class (thus relieving the instructors of much work creating these on their own). As a student you must take responsibility for your own learning and seek the best books available for learning, even if this means only using the assigned text as a source of homework problems!

Thanks to the internet, searching for high-quality books is very easy. Not only may you peruse titles and reviews, but many booksellers also allow limited access to these texts online so you can see firsthand how the books are written and decide whether or not each book would suit your needs.

If you are a student and you approach your instructor asking about other texts to read on the subject, you may find your instructor doesn't like the assigned text any better than you, and has alternatives ready to suggest. This is often the case in courses where the text has been selected by committee.

If you examine the personal library of any highly competent technician or engineer, you will likely find multiple books covering the same topics. Building your own collection of useful texts for learning is a sound strategy. Not only will you find yourself referencing these books throughout your career, but you will find their multiple explanations of common concepts an easier way to learn than by limiting yourself to just one point of view.

34.2.2 Marking versus outlining a text

A practical and common method to increase engagement while reading is to “mark” the paper pages of a book with notations, the idea being to note points of interest and thereby stimulate thinking as you read. For a brief primer on this subject I recommend Mortimer Adler’s essay *How to Mark a Book*.

We will begin our exploration of active reading with an example taken from page 101 of the classic text *The Measurement of High Temperatures* written by George Burgess and Henry Louis Le Châtelier in 1912. One reason for choosing such an old text is that the style of writing adds another challenge to the task of reading. For those already familiar with the subject of temperature measurement, the archaic writing style will help give you the perspective of a new student, encountering something unfamiliar for the first time:

CHAPTER IV.

THERMOELECTRIC PYROMETER

Principle. -- The junction of two metals heated to a given temperature is the seat of an electromotive force which is a function of the temperature only, at least under certain conditions which we shall define further on. In a circuit including several different junctions at different temperatures, the total electromotive force is equal to their algebraic sum. In a closed circuit there is produced a current equal to the quotient of this resultant electromotive force and the total resistance.

Experiments of Becquerel, Pouillet, and Regnault. -- It was Becquerel who first had the idea to profit from the discovery of Seebeck to measure high temperature (1830). He used a platinum-palladium couple, and estimated the temperature of the flame of an alcohol lamp, finding it equal to 135°. In reality the temperature of a wire heated in a flame is not that of the gases in combustion; it is inferior to this.

The method was studied and used for the first time in a systematic manner by Pouillet; he employed an iron-platinum couple which he compared with the air thermometer previously described (page 61). In order to protect the platinum from the action of the furnace gases, he inclosed it in an iron gun barrel which constituted the second metal of the junction. Pouillet does not seem to have made applications of this method, which must have given him very discordant results.

Edm. Becquerel resumed the study of his father's couple (platinum-palladium). He was the first to remark the great importance of using in these measurements a galvanometer of high resistance. It is the electromotive force which is a function

101

This is the original text as it appears on page 101 of the book. Next, we will explore different ways of “marking” the text, as if we were a new reader to this subject.

The most rudimentary method of marking a book consists of underlining and/or highlighting with a felt-tipped pen:

CHAPTER IV.

THERMOELECTRIC PYROMETER

Principle. -- The junction of two metals heated to a given temperature is the seat of an electromotive force which is a function of the temperature only, at least under certain conditions which we shall define further on. In a circuit including several different junctions at different temperatures, the total electromotive force is equal to their algebraic sum. In a closed circuit there is produced a current equal to the quotient of this resultant electromotive force and the total resistance.

Experiments of Becquerel, Pouillet, and Regnault. -- It was Becquerel who first had the idea to profit from the discovery of Seebeck to measure high temperature (1830). He used a platinum-palladium couple, and estimated the temperature of the flame of an alcohol lamp, finding it equal to 135°. In reality the temperature of a wire heated in a flame is not that of the gases in combustion; it is inferior to this.

The method was studied and used for the first time in a systematic manner by Pouillet; he employed an iron-platinum couple which he compared with the air thermometer previously described (page 61). In order to protect the platinum from the action of the furnace gases, he inclosed it in an iron gun barrel which constituted the second metal of the junction. Pouillet does not seem to have made applications of this method, which must have given him very discordant results.

Edm. Becquerel resumed the study of his father's couple (platinum-palladium). He was the first to remark the great importance of using in these measurements a galvanometer of high resistance. It is the electromotive force which is a function

101

Here you see how the reader has highlighted the paragraph introductory words in yellow ink, underlined names and dates with green lines, and underlined concepts with red lines. While such exercises may have helped the reader remain awake as well as generate cues for later cramming, it is doubtful they assisted the reader in understanding the concepts. A large degree of blame for this rather shallow and unproductive practice may be laid on instructional curricula emphasizing memorization and execution of procedures over conceptual understanding. Simply put, *this is what you get when students expect to be quizzed on isolated bits of data:* you get poor study habits such as this. Unfortunately, not only will this approach fail to yield deep understanding of the concepts, but it also reduces the act of reading to drudgery. When such practices are so common, it's no wonder a great many students loathe academic reading.

The choice of highlighting versus underlining, or of one color over another, is relatively unimportant. Any process of “marking” a book merely by drawing attention to certain words within it suffers the same weakness: all you are doing is emphasizing specific words and phrases, not incorporating your own thoughts into the text. In order to be *actively engaged* in your reading, you must expose your own thoughts and reflections on what you read, not just emphasize specific statements made by the author. As Adler points out:

“Understanding is a two-way operation; learning doesn’t consist in being an empty receptacle. The learner has to question himself and question the teacher. He even has to argue with the teacher, once he understands what the teacher is saying. And marking a book is literally an expression of differences, or agreements of opinion, with the author.”

The thrust of Adler’s essay is that the reader gains the greatest understanding of a text by expressing their own thoughts about what they are reading: articulating questions, drawing conclusions, linking ideas, and otherwise being an active participant in the reading process rather than a passive observer.

With this in mind, let's re-approach the text. Instead of simply emphasizing words and phrases, we will now show how a reader may articulate some of their own thoughts on the same page:

CHAPTER IV.

THERMOELECTRIC PYROMETER

Voltage only depends on temperature

Principle. -- The junction of two metals heated to a given temperature is the seat of an electromotive force which is a function of the temperature only, at least under certain conditions which we shall define further on. *What conditions?*

In a circuit including several different junctions at different temperatures, the total electromotive force is equal to their algebraic sum.

In a closed circuit there is produced a current equal to the quotient of this resultant electromotive force and the total resistance. *Ohm's Law? Kirchhoff's Voltage Law*

Experiments of Becquerel, Pouillet, and Regnault. -- It was Becquerel who first had the idea to profit from the discovery of Seebeck to measure high temperature (1830). He used a platinum-palladium couple, and estimated the temperature of the flame of an alcohol lamp, finding it equal to 135°. In reality the temperature of a wire heated in a flame is not that of the gases in combustion; it is inferior to this. *"inferior" = "less" (?) Why is this???*

The method was studied and used for the first time in a systematic manner by Pouillet; he employed an iron-platinum couple which he compared with the air thermometer previously described (page 61). In order to protect the platinum from the action of the furnace gases, he inclosed it in an iron gun barrel which constituted the second metal of the junction. Pouillet does not seem to have made applications of this method, which must have given him very discordant results.

Edm. Becquerel resumed the study of his father's couple (platinum-palladium). He was the first to remark the great importance of using in these measurements a galvanometer of high resistance. *Why is this important?*

101

Here you can see the focus has completely shifted away from facts and figures, and toward concepts. Note the questions raised by the reader, either doubting their own understanding of the book or the author's assertions. These questions need not be fully resolved at the conclusion of reading. Indeed, these may be excellent topics to raise in class once the student returns to school and has the instructor's attention. Questions should be taken as a positive sign for active reading and not as an indication of trouble: if you read a large body of prose and have no questions of your own, you probably weren't thinking deeply enough!

If this manner of marking the text seems messy and cluttered, one may opt to make all the notations on a separate piece of paper, or even typed into a computer for later printing and retrieval.

Even deeper engagement may be achieved if one takes the time to write an *outline* of ideas in the text. This is an exercise best done on a separate sheet of paper, or using a computer. An example is shown here, side-by-side with the original text. Note how an outline may include graphical sketches as well as words, and how it references the page number of the source text to make it easier to go back to the source and re-read if necessary:

<p style="text-align: center;">CHAPTER IV. THERMOELECTRIC PYROMETER</p> <p>Principle. -- The junction of two metals heated to a given temperature is the seat of an electromotive force which is a function of the temperature only, at least under certain conditions which we shall define further on. In a circuit including several different junctions at different temperatures, the total electromotive force is equal to their algebraic sum. In a closed circuit there is produced a current equal to the quotient of this resultant electromotive force and the total resistance.</p> <p>Experiments of Becquerel, Pouillet, and Regnault. -- It was Becquerel who first had the idea to profit from the discovery of Seebeck to measure high temperature (1830). He used a platinum-palladium couple, and estimated the temperature of the flame of an alcohol lamp, finding it equal to 135°. In reality the temperature of a wire heated in a flame is not that of the gases in combustion; it is inferior to this.</p> <p>The method was studied and used for the first time in a systematic manner by Pouillet; he employed an iron-platinum couple which he compared with the air thermometer previously described (page 61). In order to protect the platinum from the action of the furnace gases, he inclosed it in an iron gun barrel which constituted the second metal of the junction. Pouillet does not seem to have made applications of this method, which must have given him very discordant results.</p> <p>Edm. Becquerel resumed the study of his father's couple (platinum-palladium). He was the first to remark the great importance of using in these measurements a galvanometer of high resistance. It is the electromotive force which is a function</p> <p style="text-align: center;">101</p>	<p style="text-align: center;">(From page 101)</p> <p><i>Joining two different metals together and heating that junction produces a voltage dependent on temperature. What "certain conditions" affect voltage/temp relationship?</i></p> <p><i>The voltages of multiple thermoelectric junctions adds in series, with current = V/R.</i></p> <p><i>Becquerel used a platinum-palladium junction to try to measure a candle flame, but he got a result that was too low. Why is the wire's temp cooler than the flame?</i></p> <p><i>Pouillet used a platinum-iron junction using a hollow iron tube as the second conductor. This tube protected the from -- maybe corrosion?</i></p> <div style="text-align: center;"> </div> <p><i>Becquerel's son discovered you need to use a meter with a high resistance. Why is this??</i></p>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

What you see here are notes for just one page of Burgess' and Le Châtelier's text. Although there is no absolute rule for how detailed one's notes should be, they should be in *one's own words* rather than copying the author's words, and should at minimum express *at least one statement and/or question from the reader per paragraph of source text*, since each paragraph typically contains a distinct idea on the topic at hand. A complete outline of thermoelectric pyrometers would of course cover multiple pages of the source text, and not just the one page shown. Note how the outline includes cues for future learning – hints that there is more to this subject than what is immediately presented on this page of the text.

Writing your own outline of a text is especially useful when the text in question is densely packed with ideas and/or difficult to understand, because the act of outlining serves as a self-check for your own comprehension. With highlighting and underlining it is all too easy to lazily read the words and make these marks without really understanding what the text is saying. While outlining you know quite well when you haven't understood the text because you simply won't be able to express it in your own words. Instead of continuing to highlight in a state of blissful unawareness, you will find yourself stuck and unable to continue outlining. This prompts you to go back and scrutinize the text, going over it more carefully than before, until you find yourself able to continue outlining

again.

Perhaps the most common objection to outlining text as you read is that the process is slow. This raises a very important point, namely that *active reading should be slow*. Facts and figures may be skimmed, but complex ideas take time to penetrate into your mind. Not only will outlining force you to slow down when you need to, but it will also serve as a gauge for later study when you review your own notes to see if you still agree with them. In the course of studying some topic, you will often find that your understanding of that topic changes from your first impression. Seeing this change for yourself allows you to better understand how you learn, and thereby gives you practical insight into the workings of your own mind when it comes time to learn something new.

34.3 General problem-solving techniques

A variety of problem-solving techniques have been presented for students over the years which are all helpful in tackling problems both in the classroom and in the real world. Several of these techniques are presented here in this section.

34.3.1 Identifying and classifying all “known” conditions

An important step in solving certain types of problems, especially quantitative problems where calculations are necessary to obtain precise answers, it is often useful to list all the known quantities available to us relevant to the problem. Similarly, taking the time to list all relevant (and possibly relevant) mathematical formulae we might apply to the solution is a helpful step.

One way to save time applying the latter suggestion in a classroom setting is to keep a concise reference card or file filled with formulae you’ve been learning within that course. This reference may be referred to as often as necessary, without having to re-write the equations for each and every problem, thus eliminating unnecessary effort.

34.3.2 Re-cast the problem in a different format

Many people find it easier to grasp the nature of a problem – and by extension, that problem’s solution – if they can look at an illustration of the problem. Therefore, a helpful step in solving problems described to you in words is to translate those words into a picture to look at.


If you are one of those people for whom drawing is a challenge, take heart in the fact that this is a skill you can build. Practice is the key to honing this skill. With this in mind, make it a habit to sketch some kind of illustration for every problem you are asked to solve. If you are working in teams to solve a problem, a collaborative sketch goes a long way toward coordinating problem-solving efforts and ensuring everyone on the team has the same view of the problem.

For some people, describing a problem verbally is helpful in solving it. If your brain tends to work like this – understanding concepts and situations better when they are cast into clear prose – then you may find it helpful to first draft an explanatory paragraph of the problem in your own words. This is also an exercise lending itself well to team-based problem solving, as the entire team can help each other describe the nature of the problem.

34.3.3 Working backwards from a known solution


Sometimes we may gain insight into the solution of a problem by assuming we already know the answer to a similar problem, then working “backward” to find the problem from that assumed solution.

An application of this problem-solving strategy is found learning how to decode binary bits that have been encoded using the Manchester standard. With Manchester encoding, binary bits are represented by the rising and falling *edges* of square-shaped waveforms rather than high and low states themselves. For example:



Seeing this example, we note how each binary “0” bit is represented by a *falling* edge, while each “1” bit is represented by a *rising* edge.

Where most students encounter trouble is in situations where they have been given a Manchester encoded waveform and must decode it into its representative bit stream. Take this for example:




Most students’ first inclination is to ask their instructor or their classmates for an algorithm to decode the waveform. “What steps should I take to figure out where the data bits are?” they will ask. This sort of “give me the answer” mind-set should always be discouraged, because it is the polar opposite of true problem-solving technique, where the student methodically searches for patterns and develops algorithms on their own.

A better approach is to encourage the strategy of working the problem backwards: begin with a known series of binary bits, and then develop a Manchester waveform from that. The act of *encoding* a binary string provides insight that will be useful in *decoding* the next Manchester waveform they encounter.


For example, let's begin with the binary string 100011101:

1 0 0 0 1 1 1 0 1

We may begin the process of encoding this into Manchester format by sketching the rising- and falling-edges we know we will need for each bit:




Next, we can try connecting the tops and bottoms of these pulse edges to form a complete waveform. Soon, however, we will find that this is only possible where opposite bit states are adjacent to each other. Where identical bits follow in sequence, we are faced with sequential rising edges or sequential falling edges, which we cannot simply bridge at the tops or bottoms to make a full pulse:



This observation leads to the realization of why we need *reversals* in a Manchester waveform. The only way to connect repeating bits' edges together is if the waveform goes through another rising or falling edge in order to be properly set up for the next edge we need to represent a bit:


Reversals needed to "set up" for next real bit transitions




Here we see the power and utility of working a problem “backwards”: it reveals to us the *reason why* things are the way they are. Without this understanding, problem-solving is nothing more than rote recall of algorithms, and limited in application. Any problem becomes simpler to solve once we fully understand its rationale.

Once we realize the purpose for reversals in a Manchester waveform, it becomes obvious to see that these reversals always fall *between* the bit transitions, and thus are always *out of step* with the frequency of the bits. Those edge transitions representing real data bits must always fall along a regular timed interval, with reversals being “half-steps” in between those intervals. We need only to look for the widest-spaced intervals in a Manchester waveform to distinguish those pulse edges representing real data bits, and then we know to ignore any pulse edges out of step with them.


Returning to our sample problem, where we were given a Manchester waveform and asked to decode it:



First, we identify the real data bit edges by widest spacing:



Now that we know which pulse edges represent bits, we may ignore those that do not (the reversals), and decode the waveform:



34.3.4 Using thought experiments

One of the most powerful problem-solving techniques available for general use is something called a *thought experiment*. Scientists use experiments to confirm or refute hypotheses, testing their explanations by seeing whether or not they can successfully predict the outcome of a certain situation by comparing their predictions against real outcomes. While this technique is extremely useful, it might not always be practical or expedient. A useful alternative to real experiments is to mentally “model” the system and then imagine changing certain elements or variables within that model to deduce the effects.

Albert Einstein famously applied “thought experiments” to the formulation of his Theory of Relativity, for the very simple reason that he lacked the resources and technology to actually test his ideas in real life. Working as a patent clerk, he would imagine what might happen if an observer were to travel at or near the speed of light. One particular example of this is the anecdote of Einstein observing a clock tower as he rode a trolley traveling away from the tower. “What would an observer see,” he wondered, “as he viewed the clock’s face while traveling away from it at the speed of light?” Concluding that the clock’s face would appear to be frozen in time was one of the surprising “experimental” results leading Einstein to a more rigorous examination of physical effects at extremely high velocities.

“Thought experiments” are useful in solving a wide variety of problems, because they allow us to test our understanding of a system’s behavior. By imagining certain conditions or variables changing in a system and then asking ourselves what the effects will be, we probe our own understanding of that system, often times with the result being that we are able to predict its behavior under conditions that baffle us at first.

You will find “thought experiments” scattered throughout this book, used both as illustrations of problem-solving strategies and also as a tool to explain how certain technologies function. An example of this is the section explaining non-dispersive analyzers, which are instruments employing the absorption of light by certain species of chemicals in order to detect the presence and measure the quantities of those chemicals. Beginning in section 23.6 on page 1815, a series of “thought experiments” are used to explore the principles used to identify chemicals by light absorption. This series of virtual experiments becomes most valuable when this section explores the analyzer’s ability to *selectively* measure the presence of one light-absorbing chemical to the exclusion of other light-absorbing chemicals within the same mixture.

34.3.5 Explicitly annotating your thoughts

Suppose you were asked to solve this multiplication problem, without the use of a calculating machine of any kind, but with access to paper and a writing tool:

$$\begin{array}{r} 3418 \\ \times 572 \\ \hline \end{array}$$

Your primary school education should have prepared you to solve elementary arithmetic problems of this kind, by a process of digit-by-digit multiplication and addition, to arrive at an answer of 1,955,096. The procedure, while tedious, is rather simple: manually multiply the top numeral three times over by successive digits of the bottom numeral, noting any “carried” quantities as you do so, then sum those three subtotals together (padded with zeros to represent the place of the bottom numeral’s digit) to arrive at the final product.

Now suppose you were asked to solve the exact same multiplication problem, but this time doing the same digit-by-digit arithmetic all in your mind, without the use of a writing tool to annotate your work. Suddenly this elementary task becomes nearly impossible for anyone who isn’t a mathematical savant. What made the difference between this problem as an elementary exercise and this same problem as a nearly impossible feat? The answer to this question is *short-term memory*: most people do not possess a good enough short-term memory to mentally manage all the intermediate calculations necessary to complete the calculation. This is why people learn to annotate their work when performing manual multiplication, so they don’t have to rely on their limited short-term memories. The freedom to write your steps on paper converts what would otherwise be a Herculean feat of arithmetic into a rather trivial exercise.

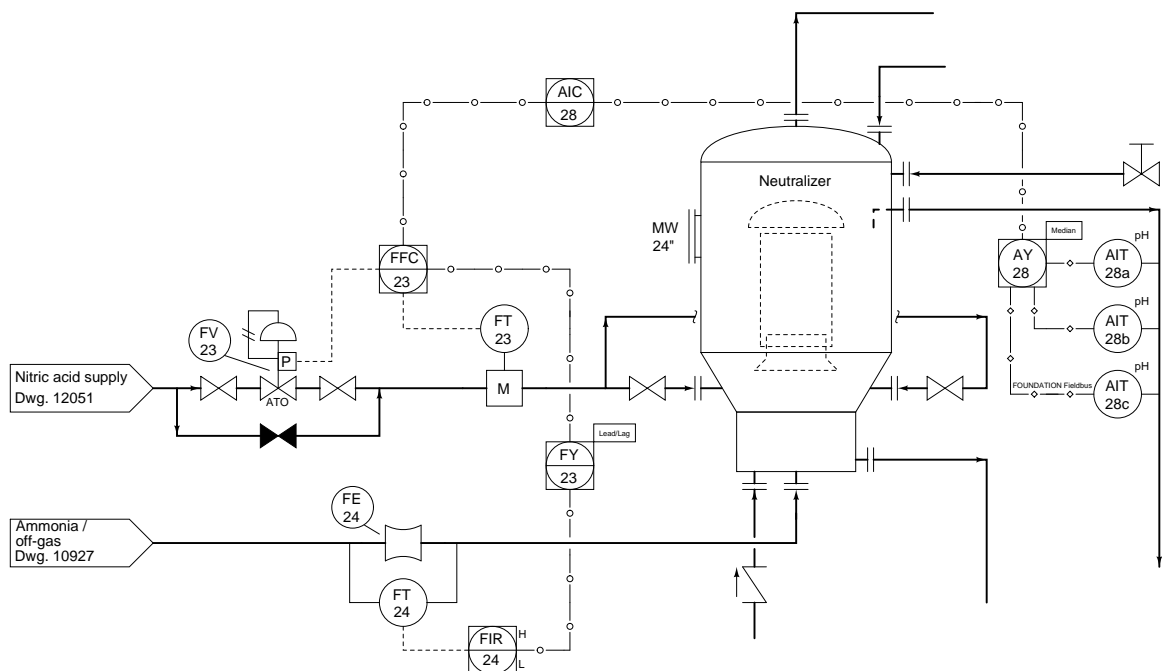
Annotating your intermediate steps as you solve a problem is actually an excellent general problem-solving strategy, applicable to far more than just arithmetic. Some examples of annotating intermediate steps are listed here:

- **Reading a complex document:** *annotating your thoughts, questions, and epiphanies as you read the text allows you to derive a better understanding of the text as a whole.*
- **Learning a new computer application:** *noting how features are accessed and identifying the necessary conditions for each feature to work helps you navigate the software more efficiently.*
- **Following a route on a map:** *marking where you started, where your destination is, and where you have traveled thus far helps you see how far you still need to go, and which alternative routes are open to you.*
- **Analyzing an electric circuit:** *annotating all calculated voltages, currents, and impedances on the diagram helps you keep track of what you know about the circuit and where to go next in your analysis.*
- **Troubleshooting a system fault:** *noting all your diagnostic steps and conclusions along the way helps you confirm or disprove hypotheses.*

Sadly, many students attempt to solve new types of problems analogously to performing multiplication without paper and pencil: they attempt to mentally manage all their intermediate steps, not writing anything down that would help them later. As a result, students tend to get “lost” when trying to solve new problems simply because they cannot readily reference of all their thoughts along the way. Most people simply give up when they begin to feel “lost” in solving a problem, thinking that if they cannot mentally picture the solution in its entirety then they have no hope of attaining it. Let’s face it: how soon would you give up on multiplying 3418×572 without a calculator if you believed the only alternative was to manage all the arithmetic in your head?

One reason why students default to the “mental-only” approach when approaching new problems is that their educational experience has only presented annotation for specific types of problems. Thus, marking all the carry digits and subtotals is something they “only do” when performing multiplication by hand; marking calculated voltages and currents on a schematic diagram is something they “only do” when solving DC resistor circuits; taking notes when reading is something they “only do” when completing a book report. In other words, *students see annotation only in very specific contexts, and so they may fail to see just how widely applicable annotation is as a problem-solving strategy.* What teachers should do is model and encourage annotation as a problem-solving technique for *all* types of problems, not just for *some* types of problems.


To illustrate how this might be done in the context of control system analysis, let us suppose we were asked to determine the effect of flow transmitter FT-24 failing with a low (no-flow) signal in this ratio control system, part of a process for manufacturing ammonium nitrate fertilizer:



Before it is possible to analyze the effects of a transmitter failure, we must first determine what the system ought to do in normal operation. Natural questions to ask might include the following:

- Where do the instrument signals come from and where do they go to?
- What does each instrument signal represent?
- What is the direction of action for each controller in the system?


With just a basic understanding of ratio control systems, we may answer all of these questions by close examination of the P&ID segment, and also annotate those thoughts and conclusions on the diagram in order to help us analyze the system's response. Starting with the first two questions of where signals originate and terminate and what each signal represents, we may annotate this with arrows and text (shown in red):



We know that all transmitters *output* data, and so all signal arrows should point *away* from all transmitters and *toward* all controllers. We know that all valves *receive* data, which means arrows must point toward the control valve. The first tag letter of each transmitter (AIT, FT) tells us its measurement function: chemical pH and flow, respectively. The fact that FT-23 is mounted on the same pipe as FV-23 tells us FT-23 must send controller FFC-23 its process variable (captive flow), making the other flow signal (from FT-24) the “wild” flow in this ratio control scheme. AIC-28’s task is to control pH exiting the neutralizer, so we know its output signal must call for a neutralizing reagent, in this case nitric acid. This tells us the signal between AIC-28 and FFC-23 must be a cascade output-setpoint link, with AIC-28 as the master controller and FFC-23 as the slave controller.

Now we turn to the question of controller action, since we know the direction of each controller's action (e.g. direct or reverse) is significant to how each controller will react to any given change in signal. Here, "thought experiments" are useful as we imagine the process variable changing due to some load condition, and then determine how the controller must respond to bring that process variable back to setpoint.

When we annotate the action of each controller, it is best to use symbols more descriptive than the words "direct" and "reverse," especially due to the confusion this often causes when distinguishing the effects of a changing PV signal versus a changing SP signal. In this case, we will write a short formula next to each controller denoting its action according to how the error is calculated ($e = PV - SP$ for direct action and $e = SP - PV$ for reverse action). We may also write "+" and "-" symbols next to each input on each controller to further reinforce the direction of each signal's influence:




FFC-23 is the best controller to start with, since it is the slave controller (in the "inner-most" control loop of this cascade/ratio system). Here, we see that FFC-23 must be reverse-acting, for if FT-23 reports a higher flow we will want FV-23 to close down. This means the remote SP input must have a non-inverting effect on the output: a greater signal from AIC-28 will increase nitric acid flow into the neutralizer. Following this reasoning, we see that AIC-28 should be direct-acting, calling for more nitric acid flow into the neutralizer as product pH becomes more alkaline (pH increases).

The purpose of the ratio control strategy is to balance the "wild" flow of ammonia into the neutralizer with a proportional flow of nitric acid. This is in keeping with principles of chemical reactions (stoichiometry) and mass balance. Therefore, we would expect an increase in ammonia flow to call for a proportionate increase in nitric acid flow, giving the wild flow signal a non-inverting

effect on FFC-23.

Only at this point in time are we fully ready to analyze the effects of FT-24 failing with a low-flow signal. Once again, we may annotate the failure on the diagram as well, arbitrarily electing to use blue “up” and “down” arrows and bold text to indicate the directions of change for each signal immediately following the failure of FT-24:



As FT-24’s signal fails low, the “wild” flow signal to FFC-23 goes low as well. Since we have already determined that input has a non-inverting effect on the ratio controller, we may conclude control valve FV-23 will close as a result, decreasing the flow of nitric acid into the neutralizer. This analysis becomes trivial after we’ve done the work of annotating the diagram with our own notes showing how the instruments are supposed to function. Without this set-up, the task of analyzing the effects of FT-24 failing low would be much more difficult.


34.4 Mathematical problem-solving techniques

Some problem-solving techniques are unique to quantitative problems, involving mathematical calculations. In this section we will explore some useful tips to help you solve such problems.


34.4.1 Manipulating algebraic equations

One of the most useful problem-solving techniques in all of algebra is the art of manipulating, or re-writing, equations to solve for a particular variable. The key to this technique is the fact that we may subject an equation to any mathematical operation we desire, so long as we apply that same operation to both sides of the equation.

Let us begin with an illustration showing why this is true. We will represent the equation $x + 9 = y$ as a balance-beam scale, showing the quantity $x + 9$ in one pan of the scale and y in the other pan:



The balance-beam analogy merely represents the fact that the two expressions $x + 9$ and y are equal; that is to say, they possess the same mathematical value as described by the equation $x + 9 = y$. It should be intuitively obvious that this equality will remain unaltered if we were to add some equal quantity to both sides of the equation, such as the number 3:




If the left-hand quantity is now larger by 3 units, and the right-hand quantity is also larger by 3 units, and those two quantities began as equal to each other, then those two quantities *must still remain equal to one another*. We have not altered the “truth” of the equation by adding 3 to both sides, any more than we would have upset the balance of a real scale by adding 3 units of mass to both pans.

We could similarly multiply each of these quantities by the same factor (say, 4) and still remain equal. The validity of the equation $x + 9 = y$ is not harmed by multiplying both sides by 4 to get $4(x + 9) = 4y$. Similarly, we could take the square-root of both sides to get $\sqrt{x + 9} = \sqrt{y}$ and still have an equality. We could reciprocate both sides to get $\frac{1}{x+9} = \frac{1}{y}$ and still have an equality. We could take the logarithm of both sides to get $\log(x + 9) = \log y$ and still have an equality. We could raise both sides to the power z to get $(x + 9)^z = y^z$ and still have an equality. The lesson here should be perfectly clear: if we start with an equation (two equal expressions), and apply the same mathematical operation to both sides of that equation, we'll still have an equation (two equal expressions).

This fact – that we may apply any mathematical operation equally to both sides of an equation without harming the validity of that equation – might seem at first to be pointless. However, it turns out to be a powerful tool for isolating variables within an equation. If we are creative in the operation(s) we choose to apply to an equation, we may do so in such a way that isolates one variable by itself, placing all other portions of the equation on the other side of the “equals” sign.

Returning to our expression $x + 9 = y$, suppose we were tasked with solving for x . In other words, we need to manipulate this equation to get x by itself on one side of the equals sign, with everything else on the other side of the equals sign. Clearly, the problem here is how to get rid of the “9” that’s being added to x on the left-hand side of the equation. One way to do this is to subtract 9 from both sides of the equation, thus canceling the “9” on the left-hand side and shifting it over to the right-hand side. This results in x being left all by itself on the left-hand side of the equation, just like we want:



Although we have the mathematical freedom to do *anything* we wish to both sides of the equation so long as we do it equally, we can see here that only one operation will work to isolate x by itself. This, therefore, is the challenge of algebraic manipulation: *how to determine which operation(s) we should apply to both sides in order to end up with the equation re-written the way we want?*

A general principle to help us answer this challenge is that of *inverse functions*: that is, pairs of mathematical operations known to “un-do” each other. The inverse function to addition is subtraction, which is why we *subtracted* 9 from each side of the equation $x + 9 = y$ in order to “un-do” the *addition* of 9 to x on the left-hand side. In other words, we determine the mathematical operation being applied to the variable we wish to isolate, and then we intentionally apply the opposite mathematical operation to both sides of the equation to “strip away” that interfering operation.

Here is a table showing a few inverse mathematical functions:

Function $f(x)$	Inverse function $f^{-1}(x)$
Addition (+)	Subtraction (−)
Multiplication (×)	Division (÷)
Power (x^n)	Root ($\sqrt[n]{x}$)
Exponent (n^x)	Logarithm ($\log_n x$)

It should be noted that sometimes we can use the same function to “un-do” itself if we apply the quantity creatively. For example, instead of subtracting 9 from each side of the equation $x + 9 = y$ to solve for x , we could have just as easily *added* -9 to both sides, which is mathematically equivalent. Likewise, instead of using division to “un-do” multiplication, we may simply multiply by the reciprocal, as in the following example where we wish to solve for x in the equation $\frac{5}{3}x = y$. Since we can see that the fraction $\frac{5}{3}$ is being *multiplied* by x , we may get rid of the fraction by multiplying both sides by its reciprocal ($\frac{3}{5}$), leaving x multiplied by a factor of 1, which is the same as x all by itself:

$$\begin{aligned}\frac{5}{3}x &= y \\ \left(\frac{3}{5}\right)\frac{5}{3}x &= \left(\frac{3}{5}\right)y \\ \frac{1}{1}x &= \left(\frac{3}{5}\right)y \\ x &= \frac{3}{5}y\end{aligned}$$

Had we chosen to *divide* both sides of this equation by the fraction $\frac{5}{3}$, we still would have arrived at a mathematically correct result, but it would have taken the form of a *compound fraction*, which is not as “elegant” a presentation:

$$\begin{aligned}\frac{5}{3}x &= y \\ \frac{\frac{5}{3}x}{\frac{5}{3}} &= \frac{y}{\frac{5}{3}} \\ x &= \frac{y}{\frac{5}{3}}\end{aligned}$$

The task of solving for a variable becomes more complicated when there are multiple operations needing to be stripped away to isolate a particular variable. Take for example the equation $5x - 4 = y$, where we wish to solve for x . Clearly, we must find a way to “un-do” the subtraction of 4, as well as the multiplication by 5, but which inverse operation should we apply first? The key here is to properly recognize the *order of operations* expressed in the original equation.

If we were to evaluate the equation $5x - 4 = y$ to arrive at a value for y given a value for x , our proper order of operations would be to first multiply x by 5, and then subtract 4. Multiplication/division precedes addition/subtraction, if there are no other influences such as parentheses. If our goal is to “un-do” each of these operations in order to arrive at x by itself, we must do so in *reverse order of operations*. This means first un-doing the subtraction of 4, and then un-doing the multiplication by 5. The following steps show how this is done:

Step	Equation	Explanation
1	$5x - 4 = y$	Original equation
2	$5x - 4 + 4 = y + 4$	Adding 4 to both sides
3	$5x = y + 4$	Simplifying
4	$\frac{5x}{5} = \frac{y+4}{5}$	Dividing both sides by 5
5	$x = \frac{y+4}{5}$	Simplifying

Note how these steps would have been different if the original equation were written with a different order of operations, such as $5(x - 4) = y$. With the parentheses forcing the order of operations such that the subtraction occurs before the multiplication, our steps for isolating x must reverse as well – we must first divide by 5, then add 4:


Step	Equation	Explanation
1	$5(x - 4) = y$	Original equation
2	$\frac{5(x-4)}{5} = \frac{y}{5}$	Dividing both sides by 5
3	$x - 4 = \frac{y}{5}$	Simplifying
4	$x - 4 + 4 = \frac{y}{5} + 4$	Adding 4 to both sides
5	$x = \frac{y}{5} + 4$	Simplifying

34.4.2 Linking formulae to solve mathematical problems


Most practical problems with mathematical solutions do not come to us with the proper formula pre-packaged for our use. Instead, we must identify relevant formulae relating the given variables together, and then use those formulae in combination to solve for the quantity we seek. The task of linking formulae to solve mathematical problems is one many students find quite challenging, and so it is worth our time to explore this in some detail.

Example: sensing valve position

For our first example we will consider a scenario where we wish to have a computer sense the stem position of a control valve. A linear potentiometer coupled to the valve stem will provide us with a suitable sensor to convert the valve stem's position into an electrical signal the computer can measure with appropriate analog I/O hardware:



This potentiometer may be used as one-half of a voltage divider network, with a constant-voltage power supply for the source, and a fixed resistor across which the computer's analog I/O can read a voltage drop. As valve stem position changes, the potentiometer's resistance in this circuit will change, thereby changing the voltage seen by the analog input module connected to the computer:



The problem we are now faced with is this: how do we program the computer to display the valve stem's position in *percent* (on a 0% to 100% range), when all it directly senses from the circuit is a varying DC voltage? It certainly would not be helpful to have the computer display the raw signal in units of volts. What we need is a mathematical formula to translate that sensed voltage drop across the 2200 ohm resistor into a percentage value representing valve stem position.

Our first step needs to be identifying all the relevant mathematical formulae in this problem. Since the computer is sensing the voltage drop across one of two resistors in a series circuit, it would appear the voltage divider formula is relevant here (relating source voltage and circuit resistances to voltage drops):

$$V = V_{source} \left(\frac{R}{R_{total}} \right)$$

We also need to somehow relate the valve's stem position (we'll use x to represent the per unit value of valve opening) to the resistance of the potentiometer. If 0% stem position (i.e. $x = 0$, valve fully closed) drives the potentiometer wiper fully down and 100% stem position (i.e. $x = 1$, valve fully open) drives the wiper fully up, the potentiometer's resistance in this circuit should be a simple proportion of its full 10000 ohm value:

$$R_{pot} = 10000x$$

Now that we have a formula relating valve stem position to electrical resistance, and another formula relating electrical resistance to voltage drop, we are ready to link these two formulae together and derive a function expressing voltage drop as a function of valve stem position.

A useful strategy for identifying how multiple formulae link together is to "mark up" those formulae on paper to show how their variables relate to the given information as well as to each other. A standard I find easy to remember and apply is to draw a *circle* around the variable I'm need to solve the value of, and to draw a *square* or a *rectangle* around any variables whose values are given. To begin this process, I will write both the voltage divider and potentiometer proportion formulae near each other, then circle x (valve stem position, per unit) as the variable I need to solve for:

$$R_{pot} = 10000(x)$$

Of course, a basic rule of algebra is that for any *one* formula it is only possible to solve for the value of *one* variable. This means *all* other quantities in a formula must be known in order to solve for that one variable. In the R_{pot} formula we have x which we're trying to solve the value of, the full-scale resistance value of 10000 ohms which is given to us in the problem, and the potentiometer's manifested resistance value in the circuit (R_{pot}). Since 10000 is a constant, the only other piece of information we need in this formula to solve for x is the resistance value R_{pot} . The fact that the variable R_{pot} remains unmarked makes this point clear: R_{pot} is the missing piece of the puzzle to solve for x .

At this juncture we look to the other formula on hand to see if any of its variables will provide us this missing information. Here, it should be clear we need to find out where R_{pot} might fit in the voltage divider formula. The R variable in the numerator of the fraction in this formula refers to the resistance across which we are measuring the voltage drop V . In this case, R must be the 2200 ohm fixed resistor, so we will enclose R in a square to represent the fact we know its value. Since the power supply's voltage is constant, we may also enclose V_{source} in a rectangle to show that value will be available to us:

$$R_{pot} = 10000(x)$$

$$V = \boxed{V_{source}} \left(\frac{\boxed{R}}{R_{total}} \right)$$

Power supply voltage
2200

What we are trying to find
 What we already know

The only other resistance variable in the voltage divider formula is R_{total} , which refers to the total resistance of the series-connected resistors in a voltage divider circuit. This particular circuit has two resistors: the 2200 ohm fixed resistor and the potentiometer. We know that total resistance in a two-resistor series circuit is the sum of those two resistors' individual values ($R_{total} = R_1 + R_2$), so we will write this as a third formula in our collection, with R_1 being the 2200 ohm fixed resistor and R_2 being the potentiometer's resistance:

$$R_{pot} = 10000(x)$$


$$V = \boxed{V_{source}} \left(\frac{\boxed{R}}{R_{total}} \right)$$

Power supply voltage
2200

What we are trying to find
 What we already know


$$R_{total} = 2200 + R_{pot}$$

Now we are ready to link these formulae together. Recall how we needed to find the value of R_{pot} in the potentiometer's formula before we could calculate the value of x (valve stem position). Note how the R_{total} formula relates that potentiometer resistance value to the other resistances in the circuit. This means the total resistance formula can provide us the value of R_{pot} we need to solve for x . We will draw a circle around R_{pot} in the total resistance formula reminding us we need that value, then show a link between this and the potentiometer formula by drawing an arrow extending from one to the other:




So far, this markup tells us R_{pot} is the missing puzzle piece to calculate x , and that R_{total} is the missing puzzle piece to calculate R_{pot} .

Looking at our three formulae, we see that the voltage divider formula is able to provide us with the value of R_{total} which we need to calculate R_{pot} which we need to calculate x . We will circle R_{total} in the voltage divider formula and draw another arrow showing the link:



The only variable unmarked and unlinked now is V , which is the voltage sensed by the computer's analog input module. This is now the one independent variable which will tell us the position of the valve stem (x).

Each arrow linking formulae together shows us where one formula will be *substituted* for a variable in another formula. The only thing we must do now prior to this substitution is algebraically manipulate each formula to solve for the one circled variable within it. I will skip these algebraic steps for brevity, and simply re-write the three formulae in manipulated form:



The logic chain of dependency linking these three formulae together is now crystal-clear: we begin with a measured voltage drop value of V to give us the value of R_{total} , which then plugs into the total resistance formula to give us the potentiometer's value R_{pot} , which then plugs into the potentiometer formula to give us valve stem position (x) as a per unit value. The algebraic substitutions are shown here:

Substituting $R_{total} - 2200$ for R_{pot} in the $x = \frac{R_{pot}}{10000}$ formula:

$$x = \frac{R_{total} - 2200}{10000}$$

Substituting $\frac{V_{source} 2200}{V}$ for R_{total} in the $x = \frac{R_{total} - 2200}{10000}$ formula:

$$x = \frac{\frac{V_{source} 2200}{V} - 2200}{10000}$$


This final formula can now be programmed into the computer, telling the computer how to calculate x as a function of V .

To summarize this problem-solving strategy:

1. Begin by writing every formula you can think of relevant to the problem at hand.
2. Identify the final value you're trying to solve for, and circle it.
3. Identify all given values, and show them by drawing squares or rectangles around those variables.
4. Identify any and all "missing puzzle pieces" in the formula with the circled variable.
5. If there is another formula containing a "missing puzzle piece," circle that variable and draw an arrow linking it to the previous formula, then identify any and all "missing puzzle pieces" needed to solve for that variable. There should only be one circled variable per formula.
6. Repeat the previous step as often as needed until there are no missing pieces left.
7. Algebraically manipulate all formulae to solve for their circled variables.
8. Algebraically substitute all variables as shown by the arrows.

Example: gas pressure inside a cylinder

Let's apply this technique to a practical problem related to engine and compressor mechanisms – how to determine the amount of gas pressure generated inside a cylinder when a piston is moved to compress the gas, given the distance of the piston's motion (x):



Compression is nothing more than a forced reduction in volume, and so a good first step for us to take is to recall any mathematical formulae relating *pressure* and *volume* for a confined gas. In this case, the Ideal Gas Law is our clear choice, describing the relationship between gas pressure, volume, temperature, and the molar gas quantity inside any enclosed volume:

$$PV = nRT$$

Where,

P = Absolute pressure (atmospheres)

V = Volume (liters)

n = Gas quantity (moles)

R = Universal gas constant (0.0821 L · atm / mol · K)

T = Absolute temperature (K)

However, the Ideal Gas Law formula contains no variable x to represent piston position. Somehow, we must find a way to relate x to the Ideal Gas Law formula in order to progress to a solution for this problem.

The first step in applying this formula is to identify which variables are given to us in the problem, and which we need to solve for. In order to solve for the value of any one variable in an equation, there cannot be any other unknown quantities – this is a basic law of algebra.

In this case, the variable we’re ultimately trying to solve for is *pressure* (P). The only other quantity found in the Ideal Gas Law formula that we know the value of at this point is R , which is a natural constant. V , n , and T are all unknown to us at this point in time. Once we do know the values of these three variables, we may algebraically manipulate the Ideal Gas Law formula to solve for P .

Once again we will begin by “marking up” the formula by drawing a circle around the variable we wish to solve for, and drawing squares around any variables or constants for which we already know the values. Any variables left unmarked (i.e. not enclosed in a circle or a square) are unknown quantities which we *must* determine before we can solve for the circled variable:

$$\textcircled{P} V = n \boxed{R} T$$

○ *What we are trying to find*
□ *What we already know*

This first step of identifying known and unknown quantities is critically important to the process of problem-solving, because it directs us to what we need to do next. An obstacle so many students and professionals alike experience is paralysis in problem-solving: they get part-way into solving a problem and encounter a point where they have no idea what to do or where to go next. This is why having a strategy to identify the next step is so vitally important.

After identifying these unknown variables of V , n , and T , we now know we need to either find other formulae to solve for those values, or investigate the real mechanism to see if we may directly measure their values.

It should be clear that temperature (T) is one of those variables we should be able to directly measure in the mechanism. We know in order to solve for pressure (P) given piston position (x), one of the real-world parameters we must first measure is temperature.

It should also be clear that volume (V) is a variable related to piston position (x), and as such we should be able to identify a formula relating the two, which we may then combine with the Ideal Gas Law formula to arrive at one step closer to our solution. Since the interior volume of the machine’s cylinder is, well, *cylindrical*, we know that the formula for calculating the volume of a cylinder is what we need to incorporate next:

$$V = \pi r^2 l$$

Where,


V = Volume of cylinder

$\pi \approx 3.1415927$

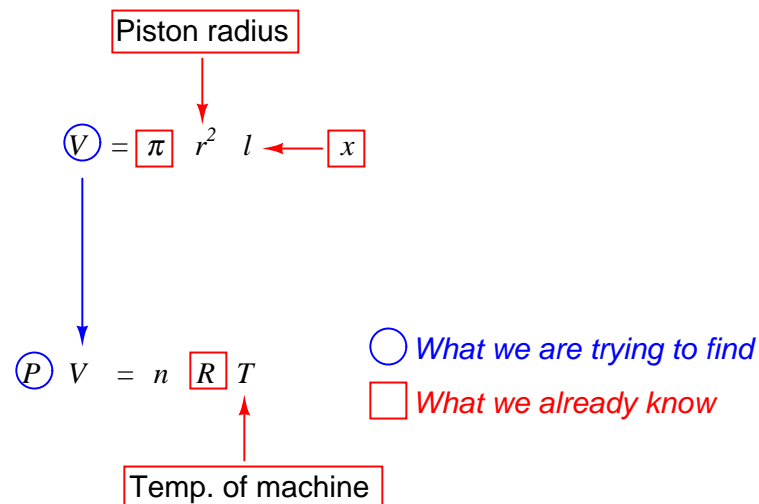
r = Radius of cylinder

l = Length of cylinder

As we did with the Ideal Gas Law formula, our next step is to identify in the cylinder volume formula those quantities we currently know versus those we don't (and therefore those we need to find). Marking up the formula with circles and squares is helpful, as is drawing an arrow between the variable we're trying to solve for with the cylinder formula and where it goes in the Ideal Gas Law formula. This arrow visually links the two formulae together, showing us where we will eventually need to perform algebraic substitution:



Pi (π) of course is a natural constant, just like R , so we enclose it in a square to show that we know its value. r is the radius of the cylinder, which in the mechanism's case is the same as the radius of the piston. As with temperature, this is a quantity we will need to determine in order to solve this problem. Length (l) is exactly the same as the piston position given to us in the problem: x , meaning we may re-write the cylinder volume formula as $V = \pi r^2 x$. Representing all of these relationships graphically:



The only unknown still existing is n , the number of moles of gas held inside the cylinder. This is not something one would be able to directly measure, like temperature, piston radius, or piston position. Either this is a quantity that would have to be provided to us as a given condition in the problem, or we must know something else about the problem in order to find the value of n . We

will set aside this issue for the time being and concentrate for now on how to combine the cylinder and Ideal Gas Law formulae.

As our graphical markup of the two formulae show us, volume (V) calculated by the cylinder formula is what gets “plugged in” to the Ideal Gas Law formula (V as well). In algebra, this operation is called *substitution*: replacing a single variable in an equation with another equation.

Substituting $V = \pi r^2 x$ for V in $PV = nRT$. . .

$$PV = nRT$$

$$P(\pi r^2 x) = nRT$$

$$P\pi r^2 x = nRT$$

Now, solving for P

$$P = \frac{nRT}{\pi r^2 x}$$

We now have a formula for P written as a function of x . Provided we know the temperature (T), the piston’s dimensions, and the molecular quantity of gas inside the cylinder (n), we may plug any arbitrary value for piston position (x) and very easily calculate pressure (P).

Returning to the question of knowing the value of molecular gas quantity (n), we might ask ourselves, “What if no one provided us with a value for n ? How could we possibly solve for pressure then?” Recall once more the basic law of algebra telling us we can only solve for one unknown quantity in a single equation at one time. What would we need to know in our custom formula in order to solve for n ? A graphic markup of the formula is helpful once again:

$$P = \frac{\textcircled{n} \text{R} \text{T}}{\pi \text{r}^2 \text{x}}$$

Temp. of machine

↓

n
 R
 T

π
 r^2
 x

↑ ↑


Piston radius Piston position

○ *What we are trying to find*
□ *What we already know*

From this we can see it would be possible to calculate n , if only we knew P . Clearly, this is a Catch-22: our goal is to calculate the value of P from n , yet the only algebraic means we have of determining n is to first know the value of P .

However, there is a solution to this conundrum. Most engine and compressor mechanisms work by moving the piston back and forth in coordination with valves to vent and direct the flow of gas into and out of the cylinder. This usually means there will be *other* piston positions where vent valve(s) are open to ensure atmospheric pressure inside the cylinder. In other words, there will be other values of x for which the pressure P will be known. In such a case where both x and P are known, we may solve for n , and then be confident that value of n will be the same for other piston positions because the vent valves close off before compression begins, trapping all gas molecules inside the cylinder.

Representing two such conditions, using subscripts to distinguish pressures and piston positions in each of the two conditions. Here, we will assume that the gas temperature is the same in both conditions, to simplify the problem:



Now that we have a means to solve for n , we may use substitution again to replace n in the “compressed” formula (condition 2) with the n from the “vented” formula (condition 1). A key assumption in making this substitution is that we know n will actually be the same value for those two conditions, which is a safe assumption because compression machines seal off the cylinder (preventing gas entry or escape) prior to the compression stroke of the piston.

$$P_1 = \frac{\overset{\circlearrowleft}{n} \boxed{R} \boxed{T}}{\boxed{\pi} \boxed{r^2} \boxed{x_1}}$$

Vented condition

$$\overset{\circlearrowleft}{P_2} = \frac{n \boxed{R} \boxed{T}}{\boxed{\pi} \boxed{r^2} \boxed{x_2}}$$

Compressed condition

(A blue arrow points from the circled n in the vented condition to the n in the compressed condition.)

To do this substitution, we must first manipulate the first equation to solve for n , then replace n in the second equation with our manipulated first equation:

$$P_1 = \frac{nRT}{\pi r^2 x_1}$$

$$n = \frac{\pi r^2 x_1 P_1}{RT}$$

Substituting . . . n

$$P_2 = \frac{\left(\frac{\pi r^2 x_1 P_1}{RT}\right) RT}{\pi r^2 x_2}$$

Here we see something very interesting: following the substitution, many of the variables and constants cancel each other out, leaving us with a much simpler formula solving for P_2 in terms of x_1 and x_2 . First, we cancel RT out of the numerator:

$$P_2 = \frac{\left(\frac{\pi r^2 x_1 P_1}{RT}\right) RT}{\pi r^2 x_2}$$

$$P_2 = \frac{\pi r^2 x_1 P_1}{\pi r^2 x_2}$$

Next, we cancel out πr^2 from numerator and denominator:

$$P_2 = \frac{x_1 P_1}{x_2}$$

$$P_2 = \frac{x_1}{x_2} P_1$$

Thus, the machine's gas pressure in the compressed condition is a simple ratio of its two piston positions multiplied by the absolute pressure in the vented condition.


34.4.3 Double-checking calculations

When performing calculations to arrive at an answer to some problem, it is important to check your work. Even if all the algebraic work you've done is perfectly correct, it is still possible to commit simple "keystroke" errors while entering numbers or executing operations on your calculator, and/or to make simple mental-math calculation errors. For this reason, teachers from time immemorial have encouraged their students to *double-check* their work.

What many students tend to do, unfortunately, is check their work by following all the same steps one more time and seeing if they get the same answer as before. While this might make sense to do at first, it actually invites the exact same errors, because our short-term memory for calculator keystrokes and mental math operations tends to be quite good. If we made a keystroke error the first time, we are very likely to make the exact same keystroke error the second time, simply following muscle memory.

A good way to avoid this error is to check your mathematical work *backwards*, beginning with the answer you previously calculated, working backwards through the equations to see if you can arrive at one of the values given to you at the start of the problem. This technique forces you to approach the problem differently, using different keystrokes in different orders.

For example, suppose you were tasked with calculating the pressure generated by a vertical column of liquid inside a process vessel, in order to properly set the lower- and upper-range values (LRV and URV) for the hydrostatic level transmitter:



The pressure seen by the transmitter at a 0% level condition (empty vessel) should be only that produced by the 2.5 foot height of the impulse tubing filled with process fluid:

$$P = \gamma h$$

$$P = (58.3 \text{ lb/ft}^3)(2.5 \text{ ft})$$

$$P = 145.75 \text{ lb/ft}^2 = 1.012 \text{ PSI}$$

Likewise, the pressure seen by the transmitter at a 100% level condition should be that produced by the impulse tubing height *and* the vessel's internal fill of 8 feet:

$$P = \gamma h$$

$$P = (58.3 \text{ lb/ft}^3)(2.5 \text{ ft} + 8 \text{ ft})$$

$$P = (58.3 \text{ lb/ft}^3)(10.5 \text{ ft})$$

$$P = 612.15 \text{ lb/ft}^2 = 4.251 \text{ PSI}$$

If we wished to check our final answer of 4.25 PSI, we could work backwards from this result to try to calculate the total fluid height of 10.5 feet, or work backwards to calculate the fluid density of 58.3 lb/ft³. Let's do the latter and see what we get:

$$\gamma = \frac{P}{h}$$

$$\gamma = \frac{612.15 \text{ lb/ft}^2}{10.5 \text{ ft}}$$

$$\gamma = 58.3 \text{ lb/ft}^3$$

Although this technique may seem obvious, it nevertheless avoids the pitfalls of repeating keystroke errors, which is an error plaguing the work of many students!

A variation on this theme is to calculate a quantity not given or reached in any of the initial problems. Here, for example, we were tasked with determining the transmitter's lower- and upper-range values (pressures at 0% fill and 100% fill). One way to check our work is to see if a different fill condition such as 50% gives us a pressure value lying exactly half-way between the LRV of 1.012 PSI and the URV of 4.251 PSI.

If we imagine the vessel half-full, our total liquid height seen by the transmitter will be 4 feet plus the suppression value of 2.5 feet, or 6.5 feet total. Calculating hydrostatic pressure at this liquid height:

$$P = \gamma h$$

$$P = (58.3 \text{ lb/ft}^3)(2.5 \text{ ft} + 4 \text{ ft})$$

$$P = (58.3 \text{ lb/ft}^3)(6.5 \text{ ft})$$

$$P = 378.95 \text{ lb/ft}^2 = 2.632 \text{ PSI}$$

A simple way to check that 2.632 lies half-way in between 1.012 and 4.251 is to calculate the average value of 1.012 and 4.251:

$$\frac{1.012 + 4.251}{2} = 2.632$$

Sure enough, it checks out to be correct. This is good validation of our initial work in calculating the transmitter's LRV of 1.012 PSI and URV of 4.251 PSI.

Yet another variation on this same theme of checking your work different ways is to approach the problem differently altogether. We know we may calculate the LRV and URV pressure values for this process vessel and fluid if we convert the fluid's density into a specific gravity value ($58.3 \text{ lb/ft}^3 = 0.9339$) and then calculate hydrostatic pressure as though we were dealing with inches water column, corrected for the specific gravity of this process fluid.

A height of 2.5 feet (0% level) is 30 inches, which when multiplied by the specific gravity value of 0.9339 yields 28.016 inches WC, which is equivalent to our first calculated pressure value of 1.012 PSI.

A height of 10.5 feet (100% level) is 126 inches, which when multiplied by the specific gravity value of 0.9339 yields 117.67 inches WC, which is equivalent to our first calculated pressure value of 4.251 PSI.

As you can see, solving for LRV and URV pressures using an entirely different technique yields the same answers as before, which is good validation of our original work.

It is important to note that double-checking your work in this manner merely *helps confirm* your original work, but it does not *conclusively prove* your original answers are correct. For example, one of the numbers we keep using in the URV calculations is the total liquid height of 10.5 feet. Suppose, though, that this height value was something we incorrectly calculated mentally by adding the suppression and the range heights. For example, suppose our actual given heights in the problem were a 2.5 feet suppression height and a 7 foot measurement range, and we mistakenly added 2.5 feet and 7 feet in our heads to get 10.5 feet. If that were the case, both the original URV pressure calculation and the subsequent double-checks might still agree with one another, but could still be wrong because they all depend on that same 10.5 foot height value.


The most realistic attitude to maintain in all problem-solving is *scientific skepticism*: there is no such thing as *absolute proof* so long as the potential for error exists. Given that you are a human being – liable to all manner of fallacies and mistakes – absolute proof will forever lie outside your grasp. The best we can do is incrementally minimize the potential for errors and mistakes, and accept the inevitable uncertainty.

34.4.4 Maintaining relevance of intermediate calculations

A very common attitude I've noticed with students is something I like to call “any procedure, so long as it works.” When learning to solve particular types of quantitative problems, the natural tendency is to identify procedural techniques for calculating the correct answer(s), and then practice those procedures until they can be performed without flaw. Unfortunately, it is all too easy to lose focus on principles when the learning emphasis is on procedure. The allure of a procedure guaranteed to yield correct answers overshadows the greater need to master and apply fundamental principles, leading to poor problem-solving ability masked by an illusion of competence.

This is a significant obstacle to deep and significant learning in the sciences, and there is no one solution to it. However, there is a way to identify and self-correct this behavior in some contexts, and it relies on a habit of identifying the real-world relevance of *all* intermediate calculations within a quantitative problem.

To illustrate, let us consider a simple voltage divider circuit comprised of two resistors, where some amount of supply voltage is divided into a smaller proportion to become the output voltage signal. The problem at hand is calculating the output voltage of this divider circuit, knowing the values of supply voltage and resistor resistances:



One of the basic formulae all beginning electronics students learn is the appropriately named *voltage divider formula*, shown here:

$$V_{out} = V_{supply} \left(\frac{R_2}{R_1 + R_2} \right)$$

Calculating V_{out} is a simple matter of substituting known values of voltage and resistance into this formula, then performing the necessary arithmetic. However, the procedure by which a student evaluates this formula – particularly in regard to the understanding of fundamental concepts – matters greatly to that student’s mastery of circuit analysis.

Suppose a student evaluates the voltage divider formula according to the order of operations enforced by the parentheses. Note the sequence of steps in this procedure, and how each step yields a value relevant to the voltage divider circuit:

Step	Calculation	Unit of measurement	Meaning
1	$R_1 + R_2$	Ohms	Total circuit resistance (R_{total})
2	$R_2 \div R_{total}$	(unitless)	Voltage division ratio
3	$V_{supply} \times \text{Ratio}$	Volts	Output voltage (V_{out})

Now let us suppose a student calculates the exact same output voltage for this divider circuit using a modified version of the voltage divider formula. This formula may be derived from the original version by applying the commutative property of multiplication, simply swapping the positions of R_2 and V_{supply} :

$$V_{out} = R_2 \left(\frac{V_{supply}}{R_1 + R_2} \right)$$

The proper order of operations for this modified formula will be different from the original version, but the final result (V_{out}) will be identical. Each step of the evaluation still yields a value relevant and applicable to the voltage divider circuit:

Step	Calculation	Unit of measurement	Meaning
1	$R_1 + R_2$	Ohms	Total circuit resistance (R_{total})
2	$V_{supply} \div R_{total}$	Amps	Circuit current (I)
3	$R_2 \times I$	Volts	Output voltage (V_{out})

Finally, let us suppose a student calculates the exact same output voltage for this divider circuit using another modified version of the voltage divider formula. This formula may be derived from the original version by applying the associative property of multiplication, grouping R_2 and V_{supply} together into the denominator of the fraction:

$$V_{out} = \frac{R_2 V_{supply}}{R_1 + R_2}$$

The proper order of operations for this modified formula will be different from the original version, but the final result (V_{out}) will be identical. Note, however, the irrelevance of the result in step #2:

Step	Calculation	Unit of measurement	Meaning
1	$R_1 + R_2$	Ohms	Total circuit resistance (R_{total})
2	$R_2 \times V_{supply}$	Volt-ohms	???
3	$R_{total} \times \text{previous result}$	Volts	Output voltage (V_{out})

Despite the mathematical equivalence of this last formula to the prior two, step #2 makes no *conceptual* sense whatsoever. The product of resistance and voltage, while mathematically useful in solving for V_{out} , has no practical meaning within this or any other circuit.

Conceptual problem-solving is an important skill, because it is only by mastering fundamental concepts can one become proficient in solving any arbitrary problem involving those concepts. Procedural problem-solving is useful only when applied to the specific type of problem the procedure was developed for, and useless when faced with any other type of problem. A student who understands the meaning of each step as they evaluate a voltage divider circuit will have little problem solving for quantities in other electrical circuits. A student who has merely memorized a step-by-step procedure for evaluating a voltage divider will struggle trying to solve for quantities in other circuits, unless they happen to have memorized step-by-step procedures for all those other circuits as well.

Troubleshooting is also closely linked with conceptual understanding. In my years as an instructor of electronics and instrumentation, I have seen an almost perfect correlation between conceptual understanding and diagnostic proficiency: procedural thinkers are invariably poor troubleshooters, and poor troubleshooters are invariably procedural thinkers.

The practical lesson we may draw from this example of voltage divider circuit evaluation is the importance of identifying the meaning of every intermediate result. If you ever find yourself performing calculations, unable to explain the practical significance of every step you take, it means *you are thinking procedurally rather than conceptually*. Instructors may apply this standard to their students' work by asking students to explain the meaning of each and every calculation they perform.

34.5 Problem-solving by simplification


A whole class of problem-solving techniques focuses on *altering* the given problem into a simpler form that is easier to analyze. Once a solution is found to the simplified problem, fresh ideas for attacking the original problem often become clear. This section will highlight multiple techniques for problem-simplification, as well as other useful techniques for problem-solving.

The first step, however, to problem simplification is to “give yourself the right” to alter the problem into a different form! Many students tend to avoid this, for fear of “getting off track” and losing sight of the original problem. What is needed is a spirit of adventure: a willingness and a confidence to explore the possibilities. Do not think you *must* solve exactly the problem that is given to you at first. Modify the problem, solve the simpler version of that problem, then apply the lessons and patterns obtained from that solution to the original (more complex) problem!

34.5.1 Limiting cases


A powerful method for analyzing the effects of a change in some system is to consider the effects of “extreme” changes, which are often easier to visualize than subtle changes. Such “extreme” changes are examples of what is generally known in science as a *limiting case*: a special case of a more general rule or trend, possessing fewer possibilities. By virtue of the fact that limiting cases have fewer possibilities, applying a limiting case to a given problem generally simplifies that problem.

Consider, for example, this Wheatstone bridge circuit, where changes in the thermistor’s resistance (with temperature) affect the output voltage of the bridge:



A realistic question to ask of this circuit is, “what will happen to V_{out} when the thermistor’s resistance increases?” If our only goal is to arrive at a qualitative answer (e.g. increase/decrease, positive/negative), we may simplify the problem by considering the effects of the thermistor failing completely open, because an “open” fault is nothing more than an extreme example (a *limiting case*) of a resistance increase.

If we perform this “thought experiment” on the bridge circuit, the circuit becomes simpler because we have eliminated one resistor (the thermistor):




Limiting case of increased thermistor resistance

With the thermistor eliminated from the circuit, we see that test point **B** has lost its connection to the negative terminal of the voltage source. This can only mean one thing for the potential at test point **B**: it will become more positive (less negative). If the bridge circuit happened to be balanced prior to the thermistor fault, V_{out} will now be such that **B** is positive and **A** is negative by comparison.

Analyzing the results of this limiting case even further, we can see that resistor R_2 now carries zero current (thanks to the thermistor now being failed open), which means R_2 will now drop zero voltage. If R_2 drops no voltage at all, test point **B** must now be at the exact same potential as the positive terminal of the voltage source. This being the case, measuring V_{out} between test points **A** and **B** will be equivalent¹ to measuring voltage across R_1 . Thus, the limiting case of V_{out} for an increase in thermistor resistance is V_{R1} , with **B** positive and **A** negative.

¹With R_2 dropping zero voltage, test point **B** is now essentially common to the node at the top of the bridge circuit. With test point **A** already common with the lower terminal of R_1 and now test point **B** common to the upper terminal of R_1 , V_{out} is exactly the same as V_{R1} .

Another realistic question to ask of this circuit is, “what will happen to V_{out} when the thermistor’s resistance decreases?” Once again, the problem-solving technique of limiting cases helps us by transforming the four-resistor bridge circuit into a three-resistor bridge circuit. The limiting case of a resistance decrease would be a condition of no resistance: a *shorted* thermistor:




Limiting case of decreased thermistor resistance

With the thermistor shorted in this “thought experiment,” we see that test point **B** now becomes electrically common with the negative terminal of the voltage source. This, of course, has the effect of making test point **B** as negative as it can possibly be. More specifically, by making test point **B** electrically common with the bottom node of the bridge, it makes V_{out} equal² to the voltage drop across R_3 . Thus, the limiting case of V_{out} for a decrease in thermistor resistance is V_{R3} , with **A** positive and **B** negative.


²As before, the limiting case of a thermistor fault causes test points **A** and **B** to become synonymous with the terminals of one of the remaining resistors, in this case R_3 . Since point **A** is already common with the upper terminal of R_3 and the shorted fault has now made point **B** common with the lower terminal of R_3 , V_{out} must be exactly the same as V_{R3} .

Let us consider another application of this problem-solving technique, this time to the analysis of a passive filter circuit:




If the type of filter circuit shown here were unknown (i.e. the student could not identify it as a low-pass, high-pass, band-pass, or band-stop filter circuit at first sight), the technique of limiting cases could be applied to determine its behavior. In this case, the limit to apply is one of frequency: we may perform “thought experiments” whereby we imagine the input frequency being extremely low, versus being extremely high.

We know that the reactance of an inductor is directly proportional to frequency ($X_L = 2\pi fL$) and that the reactance of a capacitor is inversely proportional to frequency ($X_C = \frac{1}{2\pi fC}$). Therefore, at an extremely low frequency ($f = 0$ Hz), the inductor will act like a short while the capacitor acts like an open:




Likewise, at extremely high frequencies ($f = \infty$), the capacitor will act like a short while the inductor acts like an open:



From these two limiting-case “thought experiments” we may conclude that the filter circuit is neither a low-pass nor a high-pass, because it neither passes low-frequency signals nor high-frequency signals. We may also conclude that it is not a band-stop filter, because that would pass both low-frequency and high-frequency signals. This means it must be a band-pass filter, by eliminating the other three alternatives.

If we would wish to confirm the band-pass nature of this filter by a positive experimental result rather than merely by eliminating what it is *not*, we could perform one more limiting-case “thought experiment:” a condition where the signal frequency exactly equals the resonant frequency of the LC network ($f = \frac{1}{2\pi\sqrt{LC}}$). Here, we must recall the principle that a parallel LC network has infinite impedance at its resonant frequency:



In this “thought experiment” we see that the LC network will be completely “open” and allow 100% of the input signal to appear at the output terminals. Thus, it becomes clear that this passive circuit functions as a *band-pass* filter.

As with the Wheatstone bridge circuit, the value of limiting-case analysis is that it acts to simplify the system by effectively eliminating components (replacing them with “shorts” or “opens”). Even in non-electrical problems, limiting cases works the same by simplifying a system’s behavior so that it becomes easier to apprehend, and from these simplified cases we may usually determine behavioral trends of the system (e.g. which way it tends to respond as some variable increases or decreases).

34.6 Scientific system diagnosis

At the root of successful system diagnosis is a rigorous adherence to *scientific reasoning*. There exists no single algorithmic approach to solving problems, but rather a singular *mind-set* characterized by the following traits:

- Curiosity
- Persistence
- Attention to detail
- Diligence in checking conclusions
- Regular checking of assumptions
- A willingness to abandon ideas based on contrary evidence

Science is, at its heart, a methodology useful to identify causes and effects. Thus, it is well-suited to the problem of system diagnosis, where our goal is to quickly and accurately identify the cause(s) behind improper operation (effects).

34.6.1 Scientific method

Although no one technique seems to be universally recognized as “the scientific method,” the following steps are commonly applied in science to determine causes and effects:

- Observe effects, and then create *hypotheses* (explanations accounting for those observations)
- Design a test for one or more of those hypotheses
- Perform the test (experiment), and collect data from it
- Validate or invalidate the hypotheses based on the data
- Repeat

Perhaps the most challenging step in this method is designing a good test for the hypotheses. By “test” I mean a trial that really challenges each hypothesis, and doesn’t just collect more data to support it. A good way to help yourself devise a rigorous test of any hypothesis is to keep these two questions in mind:

“If this hypothesis is true, what other effects should we see if we look for them?”

. . . and . . .

“If this hypothesis is false, what other effects should we *not* see if we look for them?”

An ideal test (experiment) is one that answers both of these questions at once, providing both positive and negative evidence.

In contrast to scientific diagnosis is a technique a colleague of mine refers to as “Easter-egging,” where the troubleshooter tries to find the problem by individually checking every component or possible fault they can think of, in serial fashion. The term “Easter-egging” invokes the image of children hunting for hidden eggs on Easter morning, randomly searching in every place they can think of where an egg might be hidden. There is no logical reasoning to “Easter-egging” and so it is a very inefficient method of solving system problems.

34.6.2 Occam’s Razor

A very helpful principle in scientific testing is something called *Occam’s Razor*, a rule stating that the simplest explanation for any observed effects is usually the most accurate. While not infallible, Occam’s Razor is nevertheless a valid “gambling strategy” based on simple probability. In system troubleshooting, it means that a single fault is more likely to account for the symptoms than a set of coincidental faults. For this reason, it is generally good practice to enter a troubleshooting scenario with the assumption that only one thing is wrong, unless the data conclusively points otherwise.

34.6.3 Diagnosing intermittent problems

Intermittent faults are some of the most challenging to diagnose, for the simple reason that the relevant symptoms come and go. A persistent fault is easier to solve because the data is continuously there for inspection.

The key to troubleshooting intermittent faults is to set up test equipment to capture events that occur when you are not directly observing them. Some suggested methods include:

- Using the “Min/Max” capture mode on a digital multimeter (DMM)
- Using a data recorder or event logger to capture signal history
- Looking for evidence left by certain intermittent faults (e.g. if the suspected fault is high temperature at a certain location, looking for evidence such as charring or discoloration that would be caused by high temperature at some past time)
- Using videorecording equipment to capture events

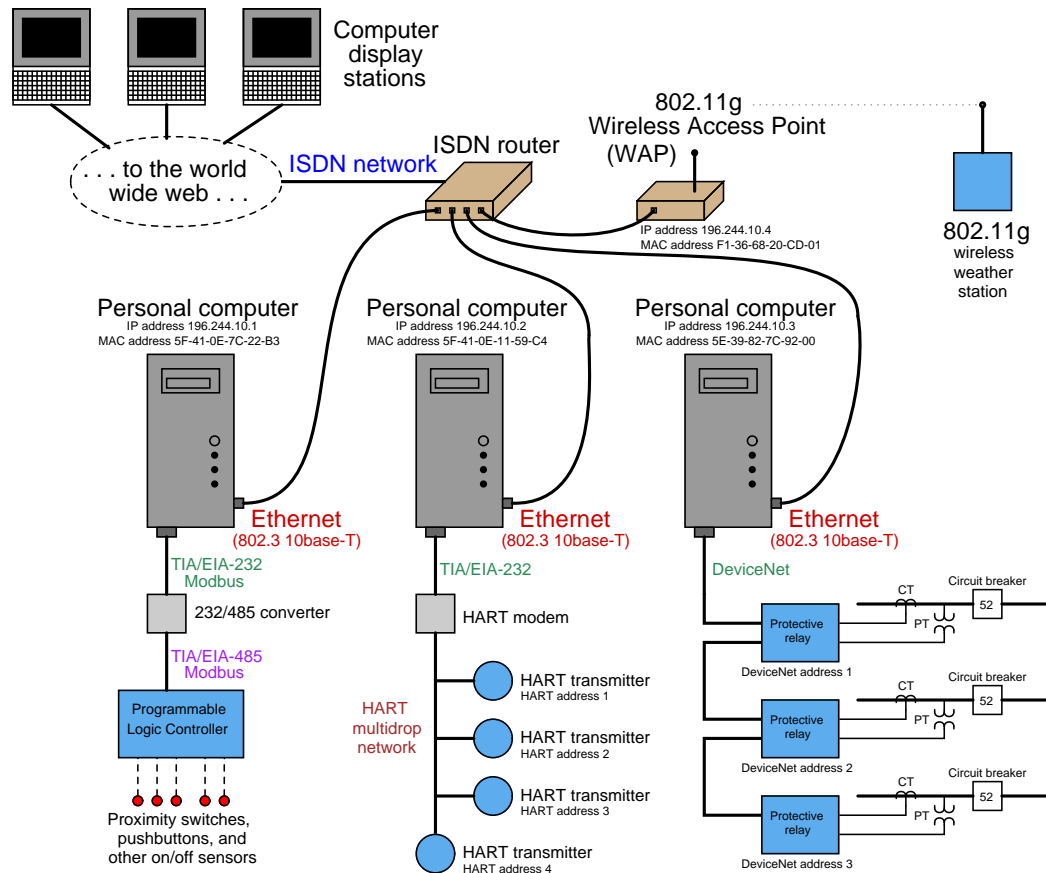
Perhaps one of the most useful features of modern digital multimeters is the ability to capture minimum and maximum signal levels. Many times I have used this feature on my own meter to monitor the highs and lows of some signal in order to capture evidence of an intermittent fault. This is also useful for monitoring signal changes that happen too fast to see on the display of a meter (e.g. detecting the peak pulse amplitude of a fast signal). While limited to the sample rate of the digital meter, it remains a powerful tool in the hands of a knowledgeable technician.

A colleague of mine once diagnosed a complex, intermittent problem on a natural gas compressor unit by setting up a video camera to film the control panel gauges on the compressor, then reviewing the video recording frame-by-frame after the camera had recorded a “trip” event. This kind of creativity is often key to diagnosing intermittent problems.

34.6.4 Strategy: tracing data paths

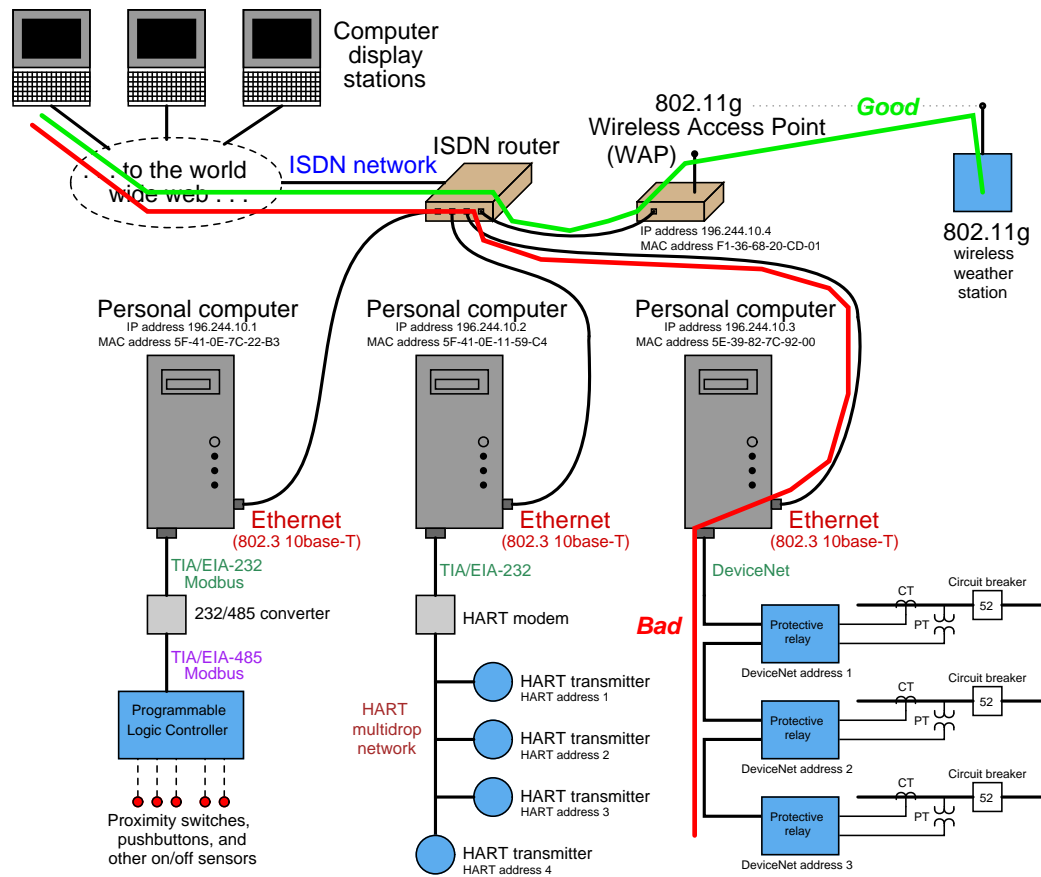
A method often useful for tracing the location of faults in complex systems is to identify where data is coming from (source), where it is going (destination), and all paths taken by the data in between. If we then plot those paths on a one-line diagram of the system, the intersection of paths often tells us where the problem lies.

For example, consider this system of networked devices in a data acquisition system:



Suppose operations personnel noticed they could no longer access any protective relay data from the left-most display station connected to the world-wide web (Internet), but they could still access live weather station data from that same display station. Applying the technique of tracing data paths may be helpful to us in locating the fault in this complex system, and also devising a good test to pinpoint the location.


Here we see the same system with green and red lines overlaid showing good data paths and failed data paths:



Note how these two paths overlap in the display station, the world-wide-web, and the ISDN router. Since data from the weather station is getting through this part of the path just fine, yet protective relay data is not, the most likely location of the fault is in a part of the system where these two data paths are *not* common. In other words, it is unlikely that the problem lies within the display station, the Internet, or the ISDN router, because all those components are proven to work just fine for the weather station data. A more probable location of the fault is in an area of the “bad” data path that is *not* common to the “good” data path. In this particular case, it points to a problem from the ISDN router to the DeviceNet network.

A good test to do at this point is to try “pinging” the right-most personal computer from one of the other two personal computers connected directly to the ISDN router. This would be testing a data path from one PC to the other, thereby testing the integrity of the right-most PC and the cable connecting it to the ISDN router. If this test is successful, the problem likely lies farther beyond the PC (e.g. in the DeviceNet network) ; if this test is unsuccessful, the problem likely lies within that PC or within the cabling connecting it to the router.

Note the careful use of the words “likely” and “probable” in the hypotheses. Hypotheses are uncertain by their very nature, and are never “proven” in any final sense. Even though our initial approach of sketching and comparing data pathways suggests a problem between the ISDN router and the DeviceNet network connecting the protective relays together, it is still possible for the problem to lie somewhere closer to the display station. If, for example, the display station relied on dedicated software “drivers” to properly poll and interpret data from the protective relay network that were not used to do the same for weather station data, the problem could actually lie within the display station! A corrupted relay driver would prevent that station from properly displaying protective relay data and yet permit the display of weather station data. In order for our data path tracing procedure to encompass this possibility, we would need to show the pathways splitting within the display station, so they would no longer be common to each other:



As always, the assumptions we embed into our hypotheses can skew their validity. If we assume completely overlapping data paths for protective relay and weather data within the display station, we would not recognize the possibility of a driver problem.

34.7 Common diagnostic mistakes

Volumes could be written about poor diagnostic technique. The following mistakes are not intended to comprise a comprehensive list, but are merely warnings against errors that are all too common among students and within the profession.

34.7.1 Failing to gather data

Perhaps the most common mistake made by technicians attempting to diagnose a system problem is failing to gather data (i.e. taking measurements and performing simple system tests) during the troubleshooting process. Even a small amount of data gathered from a system may profoundly accelerate the process of diagnosis.

A colleague of mine has a very descriptive term for the poor habit of looking for faults before gathering data: *Easter-Egging*. The idea is that a technician goes about finding the problem in a system the same way they might go about searching for eggs hidden on Easter morning: randomly. With Easter egg hunting, the eggs could literally be hidden *anywhere*, and so there is no rational way to proceed on a search. In like manner, a technician who lacks information about the nature or source of a system problem is likely to hunt in random fashion for its source. Not only will this likely require significant time and effort, but it may very well fail entirely.

A much more efficient way to proceed is to gather new data with each and every step in the troubleshooting process. By “gathering data,” I mean the following:

- Taking measurements with test equipment (multimeter, pressure gauges, etc.)
- Observing equipment indicator lights
- Stimulating the system and observing its response(s)
- Using your other senses (smell, hearing, touch) to gather clues
- Documenting new data in a notepad to help track and analyze the results of your measurements and tests

34.7.2 Failing to use relevant documentation

Diagrams are indispensable “maps” for solving problems in complex systems. A critical first step in diagnosing any system problem is to obtain correct diagrams of the system, so that you may see the pathways of power and signals in the system. Attempting to diagnose a system problem without consulting the relevant diagrams is like trying to find your way around a city without a map.

A corollary to the rule of obtaining relevant diagrams is to *use them* when reasoning through fault scenarios. All too often I see students locate the diagram for a system, glance at it, then set it aside and proceed to stumble through the rest of the diagnosis because they try to trace all the signal paths in the real world as they assess fault possibilities and devise tests. Diagrams are laid out in a clean and logical format for a reason: it is much easier to follow the flow of signals and power in a diagram than it is to follow the same flow through the convoluted paths of real-world wires and cables. If you reject the diagram in favor of tracing all pathways by looking at the real-world system, you are needlessly adding complexity to the problem: not only do you have to reason through the fault hypotheses and diagnostic tests, but you also must mentally “un-tangle” the signal paths as they are laid out in the real world (which can be a daunting task in itself!). *Do all your diagnostic thinking while looking at the diagram, and refer to the real-life system only when the time comes to execute a diagnostic test.* The result will be a much faster and less frustrating experience than if you try to trace everything in real life.


When diagnosing a problem in a system where one or more of the key components are unfamiliar to you, it is important to consult the relevant technical literature on those components. This is especially true if the component in question has been identified as suspect by your diagnostic test(s) and it is quite complex (e.g. loop controller, PLC, motor drive, data acquisition module, etc.). Just a few minutes’ worth of reading the manual may save you hours of fruitless diagnosis.

This point also underscores the necessity of technical reading as a skill to be practiced and honed at every opportunity. Being able to quickly locate pertinent information in a dense technical document is key to fast and efficient troubleshooting!

34.7.3 Fixating on the first hypothesis

When diagnosing a faulted system, an efficient strategy is to brainstorm *multiple* hypotheses accounting for the symptoms, then devise tests to support or discredit those hypotheses in the fewest steps.

Let me illustrate by example. Suppose a pump motor is controlled by a PLC, the PLC sending a command signal to a variable-frequency motor drive (VFD) to command the motor to start and stop:



If the motor refuses to start when the “Pump start” icon is pressed on the HMI screen, a competent troubleshooter will begin to mentally list a range of problems that could prevent the motor from starting:

- Motor faulted
- VFD lacking power
- VFD not configured properly to receive signal from PLC
- PLC output defective, not sending signal to VFD
- PLC program halted or faulty
- HMI not sending signal to PLC
- . . . etc.

After brainstorming such a list, a competent troubleshooter will then devise a simple test to “divide the problem space in half.” One such test³ in this system would be to use a multimeter to measure the electrical signal from the PLC’s output card to the VFD input terminals. If a signal appears when the “Pump start” icon on the HMI is pressed, it means everything in the HMI and PLC is working as it should, and that the problem must lie with the VFD or beyond. If no signal appears, it means the problem lies with the VFD, HMI, or associated cabling. Again, the wise strategy is to brainstorm multiple hypotheses explain why the motor won’t start, then execute simple tests to eliminate most of those hypotheses so you may focus on those that are most likely.

By contrast, a novice might only think of one possibility – such as the VFD being improperly configured – and then immediately fixate on that hypothesis by inspecting the drive parameters looking for one that is improperly set. If the fault lies elsewhere, the novice could spend all day reviewing VFD parameters and never find the problem. Given the wide range of possible faults, fixating on any one fault from the start is very likely a waste of time. I have watched technicians and students alike waste hours of time trying to find a fault that was not where they were looking, simply because that was the first area they thought of to look toward. Only after squandering valuable time on one failed hypothesis will the novice then consider other possibilities and other tests.


³Other possible tests include inspecting the LED status light on that PLC output card channel (a light indicates the HMI and PLC program are working correctly, and that the problem could lie within the output card or beyond to the motor) or measuring voltage at the drive output (voltage there indicates the problem must lie with the motor or the cable to the motor rather than further back).

34.7.4 Failing to build and test a new system in stages

Technicians must sometimes assemble new systems from components. A very common mistake is to assemble the system completely before attempting to test it for proper operation. This is almost always a grievous mistake.

The number of potential mistakes one can make when assembling a brand-new system is quite large. Given this large set of potential mistakes, the probability of making multiple mistakes when assembling the system is very high. Since diagnosis of a system with multiple faults is always more complicated than diagnosing a system with one fault, waiting for the entire system to be assembled before checking it invites multi-fault scenarios.

To illustrate, consider this pressure-control system, where an electronic pressure transmitter sends a 4-20 mA signal to a loop controller, which in turn drives a control valve with another 4-20 mA signal:



Imagine building this system, placing each component in the proper location, connecting all wires together, and testing it for proper operation. If one were to wait until the entire system were assembled before testing, the probability of having to diagnose multiple faults would be great.

A better strategy would be to assemble and test the system in stages. Consider this sequence of steps as a more practical alternative:

1. Install the I/P transducer, connecting air tubes to supply and valve diaphragm.
2. *Test the I/P and control valve operation using a loop calibrator in “source” mode to drive a 4-20 mA signal to the I/P.*
3. Install and wire power to the loop controller, ensuring it powers up properly.
4. Connect cabling between the I/P and the loop controller’s output.
5. *Test the controller’s ability in manual mode to “stroke” the control valve throughout its entire range.*
6. Connect wires between the loop controller’s input and the 250 ohm resistor on the terminal block.
7. *Test the controller’s ability to properly read an input signal by using a loop calibrator to drive 4-20 mA through the 250 ohm resistor.*
8. Install the pressure transmitter, connecting impulse line between it and the process line.
9. *Power the transmitter with a portable DC power supply (or loop calibrator set to the appropriate mode) and check its calibration by applying known pressures to the input tube.*
10. Connect wires between the permanent DC power supply, the transmitter, and the controller’s input.
11. *Apply pressure to the transmitter input and check to see that it reads properly on the controller’s digital display.*
12. *Test the controller’s ability to monitor and control process pressure in manual mode.*
13. Perform manual-mode (open-loop) tests to verify process characteristics and obtain data needed for loop tuning (e.g. lag time, dead time, etc.).
14. Enter preliminary PID tuning parameter values.
15. *Test the controller’s ability to monitor and control process pressure in automatic mode.*
16. Modify PID tuning parameter values and re-test in automatic mode until robust control is obtained.

Note how the pressure control instrumentation is constructed and then immediately tested as a series of sub-systems, rather than assembling the entire thing and testing only at the very end. Although the *built-test-build* sequence shown here may appear to be more time-intensive at first blush, it will actually save a lot of time and confusion over the *build-everything-then-test-last* method favored by novices.

34.8 Helpful “tricks” using a digital multimeter (DMM)

The digital multimeter (DMM) is quite possibly the most useful tool in the instrument technician’s collection⁴. This one piece of test equipment, properly wielded, yields valuable insight into the status and operation of many electrical and electronic systems. Not only is a good-quality multimeter capable of precisely indicating electrical voltage, current, and resistance, but it is also useful for more advanced tests. The subject of this section is how to use a digital multimeter for some of these advanced tests⁵.

For all these tests, I suggest the use of a top-quality field multimeter. I am personally a great fan of *Fluke* brand meters, having used this particular brand for nearly my whole professional career. The ability of these multimeters to accurately measure true RMS amplitude, discriminate between AC and DC signals, measure AC signals over a wide frequency range, and survive abuse both mechanical and electrical, is outstanding.

34.8.1 Recording unattended measurements

Many modern multimeters have a feature that records the highest and lowest measurements sensed during the duration of a test. On Fluke brand multimeters, this is called the *Min/Max* function. This feature is extremely useful when diagnosing intermittent problems, where the relevant voltages or currents indicating or causing the problem are not persistent, but rather come and go. Many times I have used this feature to monitor a signal with an intermittent “glitch,” while I attended to other tasks.

The most basic high-low capture function on a multimeter only tells you what the highest and lowest measured readings were during the test interval (and that only within the meter’s scan time – it is possible for a very brief transient signal to go undetected by the meter if its duration is less than the meter’s scan time). More advanced multimeters actually log the *time* when an event occurs, which is obviously a more useful feature. If your tool budget can support a digital multimeter with “logging” capability, spend the extra money and take the time to learn how this feature works!


⁴As a child, I often watched episodes of the American science-fiction television show *Star Trek*, in which the characters made frequent use of a diagnostic tool called a *tricorder*. Week after week the protagonists of this show would avoid trouble and solve problems using this nifty device. The *sonic screwdriver* was a similar tool in the British science-fiction television show *Doctor Who*. Little did I realize while growing up that my career would make just as frequent use of another diagnostic tool: the electrical multimeter.

⁵I honestly considered naming this section “Stupid Multimeter Tricks,” but changed my mind when I realized how confusing this could be for some of my readers not familiar with colloquial American English.

34.8.2 Avoiding “phantom” voltage readings

My first “trick” is not a feature of a high-quality DMM so much as it is a solution to a common problem *caused* by the use of a high-quality DMM. Most digital multimeters exhibit very high input impedance in their voltage-measuring modes. This is commendable, as an ideal voltmeter should have infinite input impedance (so as to not “load” the voltage signal it measures). However, in industrial applications, this high input impedance may cause the meter to register the presence of voltage where none should rightfully appear.

Consider the case of testing for the absence of AC voltage on an isolated power conductor that happens to lie near other (energized) AC power conductors within a long run of conduit:




With the power switch feeding wire 5 in the open state, there should be no AC voltage measured between wire 5 and neutral (L2), yet the voltmeter registers slightly over 10 volts AC. This “phantom voltage” is due to capacitive coupling between wire 5 and wire 8 (still energized) throughout the length of their mutual paths within the conduit.

Such phantom voltages may be very misleading if the technician encounters them while troubleshooting a faulty electrical system. Phantom voltages give the impression of connection (or at least high-resistance connection) where no continuity actually exists. The example shown, where the phantom voltage is 10.3 volts compared to the source voltage value of 120 volts, is actually quite modest. With increased stray capacitance between the conductors (longer wire runs in close proximity, and/or more than one energized “neighboring” wire), phantom voltage magnitude begins

to approach that of the source voltage⁶.


The equivalent circuit is shown here, with the DMM modeled as a $10\text{ M}\Omega$ resistance:



An analog voltmeter would never have registered 10.3 volts under the same conditions, due to its substantially lower input impedance. Thus, “phantom voltage” readings are a product of modern test equipment more than anything else.

⁶I have personally measured “phantom” voltages in excess of 100 volts AC, in systems where the source voltage was 120 volts AC.

The obvious solution to this problem is to use a different voltmeter – one with a much lesser input impedance. But what is a technician to do if their only voltmeter is a high-impedance DMM? Connect a modest resistance in parallel with the meter input terminals, of course! Fluke happens to market just this type of accessory⁷, the SV225 “Stray Voltage Adapter” for the purpose of eliminating stray voltage readings on a high-impedance DMM:



With the voltmeter’s input impedance artificially decreased by the application of this accessory, the capacitive coupling is insufficient to produce any substantial voltage dropped across the voltmeter’s input terminals, thus eliminating. The technician may now proceed to test for the presence of AC control signal (or power) voltages with confidence.

⁷Before there was such an accessory available, I used a 20 kΩ high-power resistor network connected in parallel with my DMM’s input terminals, which I fabricated myself. It was ugly and cumbersome, but it worked well. When I made this, I took great care in selecting resistors with power ratings high enough that accidental contact with a truly “live” AC power source (up to 600 volts) would not cause damage to them. A pre-manufactured device such as the Fluke SV225, however, is a much better option.

34.8.3 Non-contact AC voltage detection

While the last multimeter “trick” was the elimination of a parasitic effect, this trick is the exploitation of that same effect: “phantom voltage” readings obtained through capacitive coupling of a high-impedance voltmeter to a conductor energized with AC voltage (with respect to ground). You may use a high-impedance AC voltmeter to perform qualitative measurements of ground-referenced AC power voltage by setting the meter to the most sensitive AC range possible, grounding one test lead, and simply touching the other test lead to the insulation of the conductor under test. The presence of voltage (usually in the range of millivolts AC) upon close proximity to the energized conductor will indicate the energization of that conductor.

This trick is useful for determining whether or not particular AC power or control wires are energized in a location where the only access you have to those wires is their insulating sheaths. An example of where you might encounter this situation is where you have removed the cover from a conduit elbow or other fitting to gain access to a wire bundle, and you find those wires labeled for easy identification, but the wires do not terminate to any exposed metal terminals for you to contact with your multimeter’s probe tips. In this case, you may firmly connect one probe to the metal conduit fitting body, while individually touching the other probe tip to the desired conductors (one at a time), watching the meter’s indication in AC millivolts.

Several significant caveats limit the utility of this “trick:”

- The impossibility of quantitative measurement
- The potential for “false negative” readings (failure to detect a voltage that is present)
- The potential for “false positive” readings (detection of a “phantom voltage” from an adjacent conductor)
- The exclusive applicability to AC voltages of significant magnitude (≥ 100 VAC)

Being a qualitative test only, the millivoltage indication displayed by the high-impedance voltmeter tells you nothing about the actual magnitude of AC voltage between the conductor and ground. Although the meter’s input impedance is quite constant, the parasitic capacitance formed by the surface area of the test probe tip and the thickness (and dielectric constant) of the conductor insulation is quite variable. However, in conditions where the validity of the measurement may be established (e.g. cases where you can touch the probe tip to a conductor known to be energized in order to establish a “baseline” millivoltage signal), the technique is useful for quickly checking the energization status of conductors where ohmic (metal-to-metal) contact is impossible.

For the same reason of wildly variable parasitic capacitance, this technique should *never* be used to establish the de-energization of a conductor for safety purposes. The only time you should trust a voltmeter’s non-indication of line voltage is when that same meter is validated against a known source of similar voltage in close proximity, and when the test is performed with direct metal-to-metal (probe tip to wire) contact. A non-indicating voltmeter *may* indicate the absence of dangerous voltage, or it may indicate an insensitive meter.

34.8.4 Detecting AC power harmonics

The presence of *harmonic* voltages⁸ in an AC power system may cause many elusive problems. Power-quality instruments exist for the purpose of measuring harmonic content in a power system, but a surprisingly good qualitative check for harmonics may be performed using a multimeter with a frequency-measuring function.

Setting a multimeter to read AC voltage (or AC current, if that is the quantity of interest) and then activating the “frequency” measurement function should produce a measurement of exactly 60.0 Hz in a properly functioning power system (50.0 Hz in Europe and some other parts of the world). The only way the meter should ever read anything significantly different from the base frequency is if there is significant harmonic content in the circuit. For example, if you set your multimeter to read frequency of AC voltage, then obtained a measurement of 60 Hz that intermittently jumped up to some higher value (say 78 Hz) and then back down to 60 Hz, it would suggest your meter was detecting harmonic voltages of sufficient amplitude to make it difficult for your meter to “lock on” to the fundamental frequency.

It is very important to note that this is a crude test of power system harmonics, and that measurements of “solid” base frequency do not guarantee the absence of harmonics. Certainly, if your multimeter produces unstable readings when set to measure frequency, it suggests the presence of strong harmonics in the circuit. However, the absence of such instability does not necessarily mean the circuit is free of harmonics. In other words, a stable reading for frequency is *inconclusive*: the circuit might be harmonic-free, or the harmonics may be weak enough that your multimeter ignores them and only displays the fundamental circuit frequency.

⁸These are AC voltages having frequencies that are integer-multiples of the fundamental powerline frequency. In the United States, where 60 Hz is standard, harmonic frequencies would be whole-number multiples of 60: 120 Hz, 180 Hz, 240 Hz, 300 Hz, etc.

34.8.5 Identifying noise in DC signal paths

An aggravating source of trouble in analog electronic circuits is the presence of AC “noise” voltage superimposed on DC signals. Such “noise” is immediately evident when the signal is displayed on an oscilloscope screen, but how many technicians carry a portable oscilloscope with them for troubleshooting?

A high-quality multimeter exhibiting good discrimination between AC and DC voltage measurement is very useful as a qualitative noise-detection instrument. Setting the multimeter to read AC voltage, and connecting it to an signal source where pure (unchanging) DC voltage is expected, should yield a reading of nearly zero millivolts. If noise is superimposed on this DC signal, it will reveal itself as an AC voltage, which your meter will display.

Not only is the AC voltage capability of a high-quality (discriminating) multimeter useful in detecting the presence of “noise” voltage superimposed on analog DC signals, it may also give clues as to the source of the noise. By activating the frequency-measuring function of the multimeter while measuring AC voltage (or AC millivoltage), you will be able to track the frequency of the noise to see its value and stability.

Once on a job I was diagnosing a problem in an analog power control system, where the control device was acting strangely. Suspecting that noise on the measurement signal line might be causing the problem, I set my Fluke multimeter to measure AC volts, and read a noise voltage of several tenths of a volt (superimposed on a DC signal a few volts in magnitude). This told me the noise *was* indeed a significant problem. Pressing the “Hz” button on my multimeter, I measured a noise frequency of 360 Hz, which happens to be the “ripple” frequency of a six-pulse (three-phase) AC-to-DC rectifier operating on a base frequency of 60 Hz. This told me where the likely source of the noise was, which led me to the physical location of the problem (a bad shield on a cable run near the rectified power output wiring).

34.8.6 Generating test voltages

Modern digital multimeters are fantastically capable measurement tools, but did you know they are also capable of *generating* simple test signals? Although this is not the design purpose of the resistance and diode-check functions of a multimeter, the meter does output a low DC voltage in each of these settings.

This is useful when qualitatively testing certain instruments such as electronic indicators, recorders, controllers, data acquisition modules, and alarm relays, all designed to input a DC voltage signal from a 250 ohm resistor conducting the 4-20 mA electronic transmitter signal. By setting a multimeter to either the resistance (Ω) or diode check function and then connecting the test leads to the input terminals of the instrument, the instrument's response may be noted.


Of course, this is a *qualitative* test only, since multimeters are not designed to output any precise amount of voltage in either the resistance or diode-check modes. However, for testing the basic response of a process indicator, recorder, controller, data acquisition channel, DCS input, or any other DC-signal-receiving devices, it is convenient and useful. In every multimeter I have ever tried this with, the diode-check function outputs *more* voltage than the resistance measurement function⁹. This gives you two levels of “test signal” generation: a low level (resistance) and a high level (diode check). If you are interested in using your multimeter to generate test voltages, I recommend you take the time to connect your multimeter to a high-impedance voltmeter (such as another digital multimeter set to measure DC volts) and note just how much voltage your meter outputs in each mode. Knowing this will allow you to perform tests that are more quantitative than qualitative.

⁹There is a design reason for this. Most digital multimeters are designed to be used on semiconductor circuits, where the minimum “turn-on” voltage of a silicon PN junction is approximately 500 to 700 millivolts. The diode-check function must output more than that, in order to force a PN junction into forward conduction. However, it is useful to be able to check ohmic resistance in a circuit *without* activating any PN junctions, and so the resistance measurement function typically uses test voltages *less than 500 millivolts*.

34.8.7 Using the meter as a temporary jumper

Often in the course of diagnosing problems in electrical and electronic systems, there is a need to temporarily connect two or more points in a circuit together to force a response. This is called “jumping,” and the wires used to make these temporary connections are called *jumper wires*.

More than once I have found myself in a position where I needed to make a temporary “jumper” connection between two points in a circuit, but I did not have any wires with me to make that connection. In such cases, I learned that I could use my multimeter test leads while plugged into the *current-sensing* jacks of the meter. Most digital multimeters have a separate jack for the red test lead, internally connected to a low-resistance *shunt* leading to the common (black) test lead jack. With the red test lead plugged into this jack, the two test leads are effectively common to one another, and act as a single length of wire.



Touching the meter’s test leads to two points in a circuit will now “jumper” those two points together, any current flowing through the shunt resistance of the multimeter. If desired, the meter may be turned on to monitor how much current goes through the “jumper” if this is diagnostically relevant.

An additional benefit to using a multimeter in the current-measuring mode as a test jumper is that this setting is usually current-protected by a fuse inside the meter. Applying jumper wires to a live circuit may harbor some danger if significant potential and current-sourcing capability exist between those two points: the moment a jumper wire bridges those points, a dangerous current may develop within the wire. Using the multimeter in this manner gives you a *fused* jumper wire: an added degree of safety in your diagnostic procedure.

References

Adler, Mortimer, “How to Mark a Book”, *The McGraw-Hill Reader*, McGraw-Hill Book Company, New York, NY, 1982.

Appendix A

Flip-book animations


This appendix demonstrates certain principles through the use of “flip-book” animation. Each page of these appendix sections forms one frame of the “animation,” viewed by rapidly flipping pages (if the book is printed), or rapidly clicking the “next page” button (if the book is viewed on a computer). While crude, this animation technique enjoys the benefits of low technology (it even works in paper form!) and convenient pausing at critical frames.


Enjoy!


A.1 Polyphase light bulbs animated


The key to understanding how three-phase electric motors work is to have an accurate mental picture of the *rotating magnetic field* created by the stator windings of a polyphase motor. One of the best ways to visualize this phenomenon is to observe a string of “chaser” lights blinking in a polyphase sequence. Just in case you don’t happen to have a string of polyphase lights at your viewing convenience, I have provided a simulation here to demonstrate how the illusion of motion is created by the sequential energization of light bulbs.


Note how the lights appear to “move” from left to right as the energization sequences moves from A to C. After a few cycles of ABC, two of the wires are crossed to reverse phase sequence. This has the result of reversing the apparent direction of motion! It matters not which two phases are reversed. In this animation, I reverse phases A and B, but I could have just as well swapped phases B and C, or phases A and C, and created the exact same effect.














































































A.2 Polyphase induction motor animated


The following animation shows how the “rotating” magnetic field of a three-phase AC induction motor is produced by the interaction of three stator winding sets energized with different phases (A, B, and C) of a three-phase AC power source. A red arrow shows the direction of the *resultant* magnetic field created by the interaction of the three winding sets.




































































A.3 Rotating phasor animated

The following animation shows a rotating phasor in three-dimensional form. The phasor rotates in a complex plane (with real and imaginary axes), but travels linearly along the time axis. In doing so it traces a path that looks like a circle when viewed along the centerline (time axis) but looks like a sinusoidal wave when viewed from above or along the side.

Euler's Relation describes the phasor's position in the complex plane:

$$e^{j\omega t} = \cos \omega t + j \sin \omega t$$

Where,

e = Euler's number (approximately equal to 2.718281828)

ω = Angular velocity, in radians per second

t = Time, in seconds

$\cos \omega t$ = Horizontal projection of phasor (along a real number line) at time t

$j \sin \omega t$ = Vertical projection of phasor (along an imaginary number line) at time t

If you imagine the phasor's length either growing or decaying exponentially over time, the result will be a spiral that either widens like a horn or shrinks like a funnel. This would be a visualization of a *complex* exponential, where the s variable defines both the rate of growth/decay (the envelope of the spiral) and the angular velocity (the pitch of the spiral):

$$e^{st} = e^{(\sigma + j\omega)t} = e^{\sigma t} e^{j\omega t}$$

Where,

s = Complex growth/decay rate and frequency (sec^{-1})


$\sigma = \frac{1}{\tau}$ = Real growth/decay rate (time constants per second, or sec^{-1})


$j\omega$ = Imaginary frequency (radians per second, or sec^{-1})


t = Time (seconds)


The example of a unit-length rotating phasor is nothing more than a special case of the complex exponential, where $\sigma = 0$ (i.e. there is no growth or decay over time):


$$e^{st} = e^{(0 + j\omega)t} = e^0 e^{j\omega t} = e^{j\omega t}$$




















































































































A.4 Differentiation and integration animated

The following animation shows the calculus concepts of differentiation and integration (with respect to time) applied to the filling and draining of a water tank.


The animation shows two graphs relating to the water storage tank: one showing the volume of stored water in the tank (V) and the other showing volumetric flow rate in and out of the tank (Q). We know from calculus that volumetric flow rate is the *time-derivative* of volume:


$$Q = \frac{dV}{dt}$$


We also know that change in volume is the *time-integral* of volumetric flow rate:


$$\Delta V = \int_{t_0}^{t_1} Q dt$$


Thus, the example of a water storage tank filling and draining serves to neatly illustrate both concepts in relation to each other.













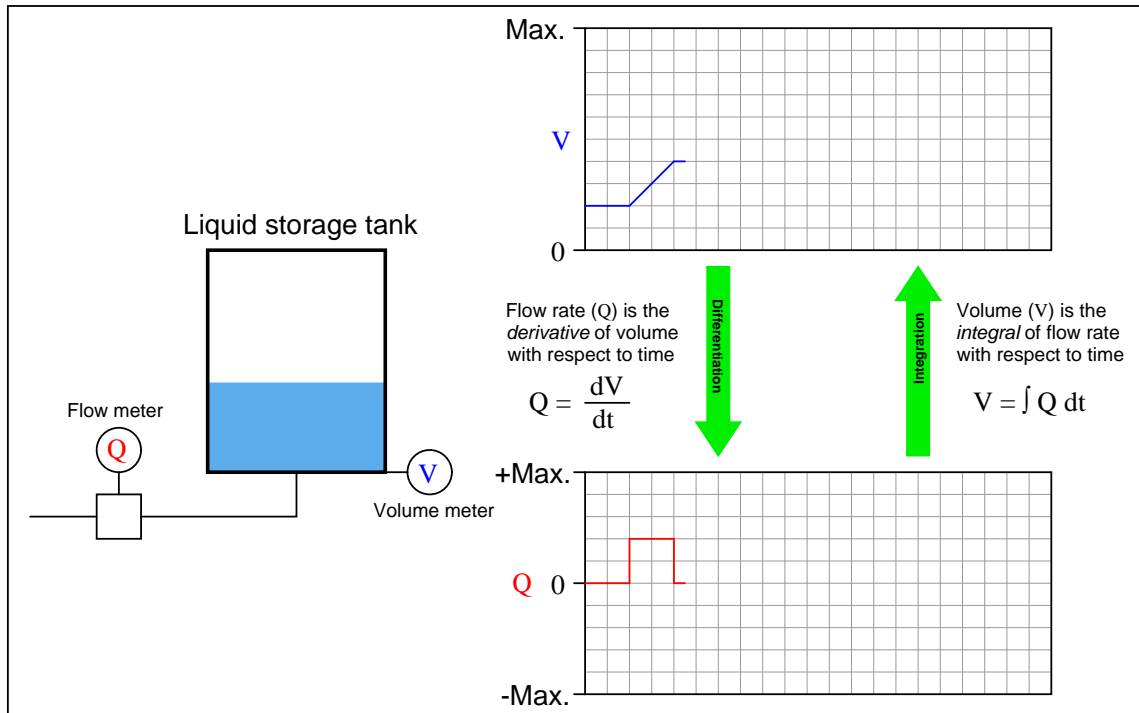

















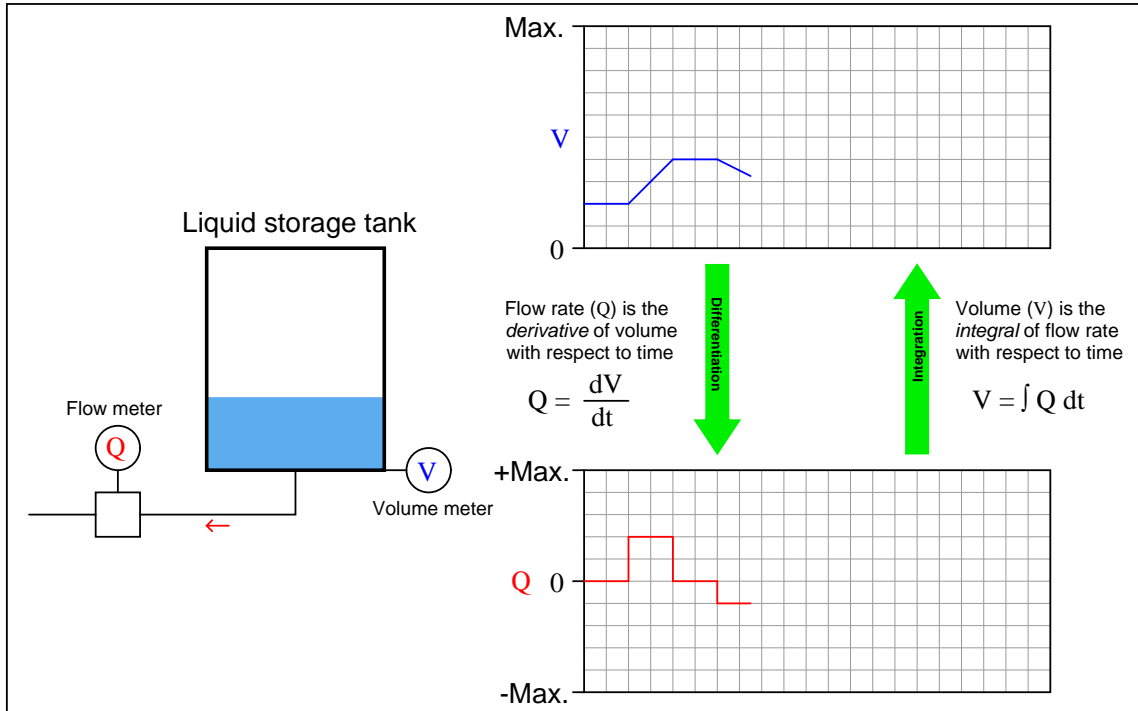


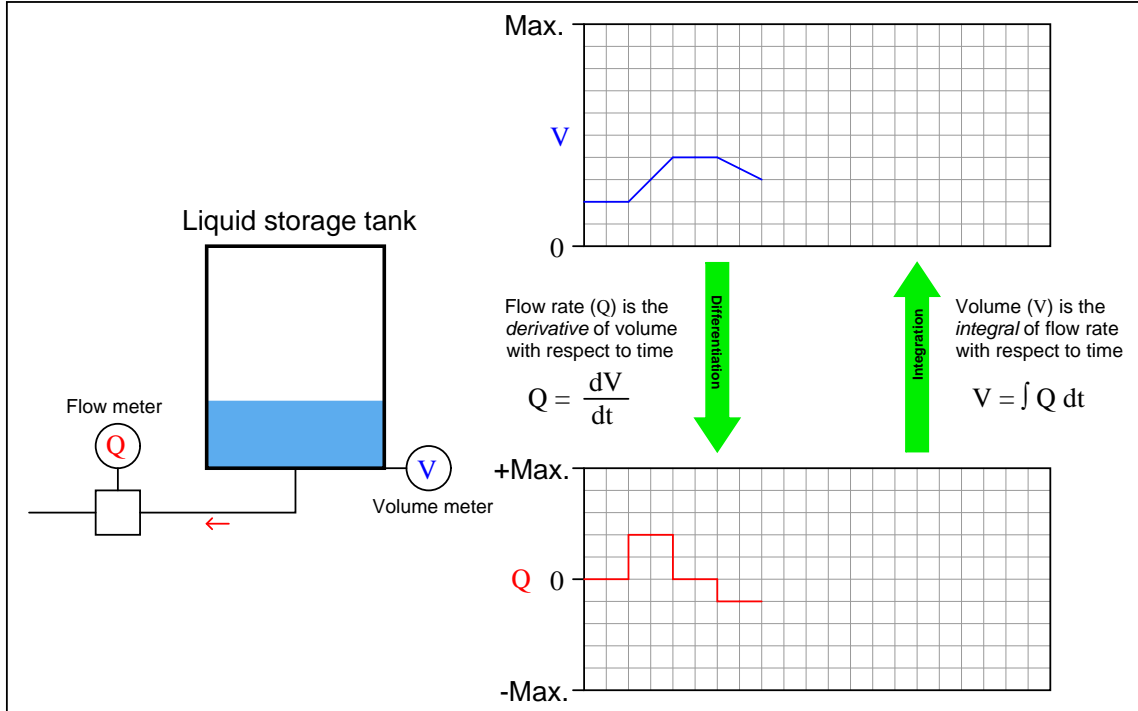




















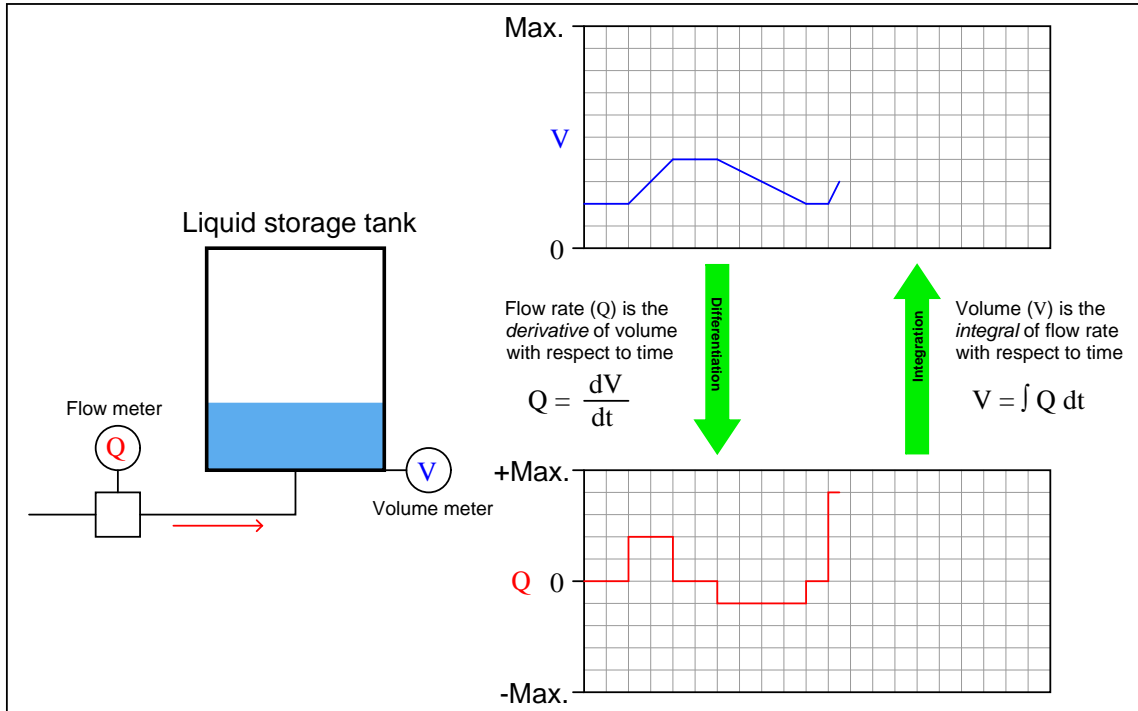











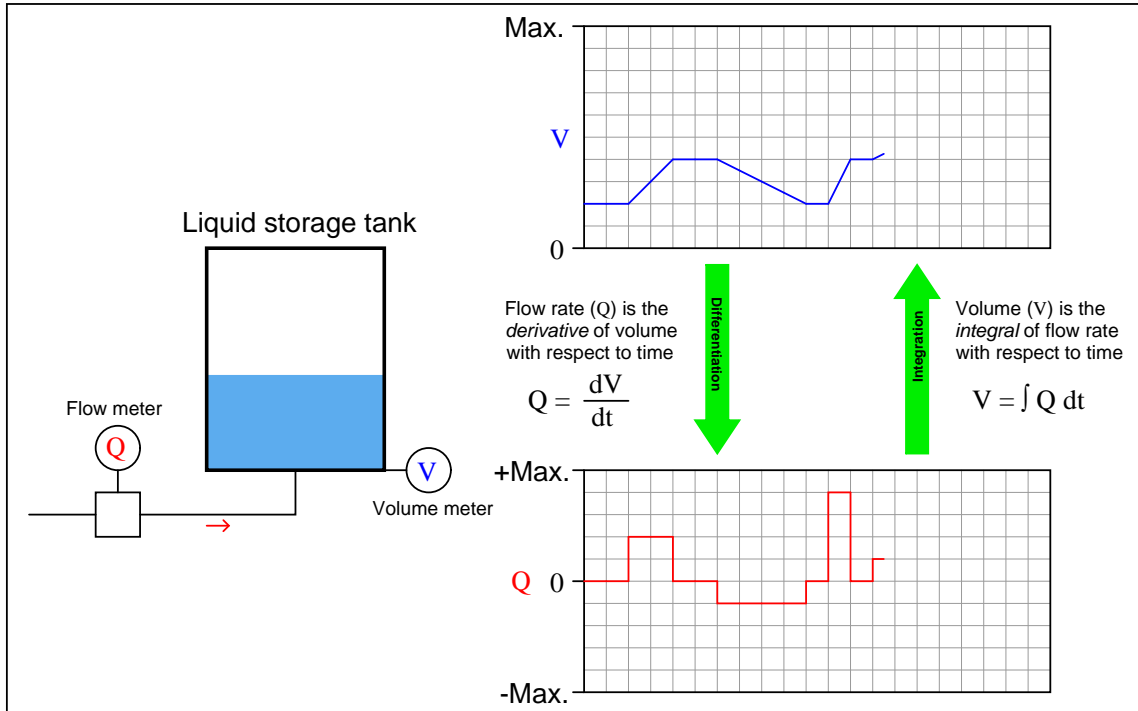


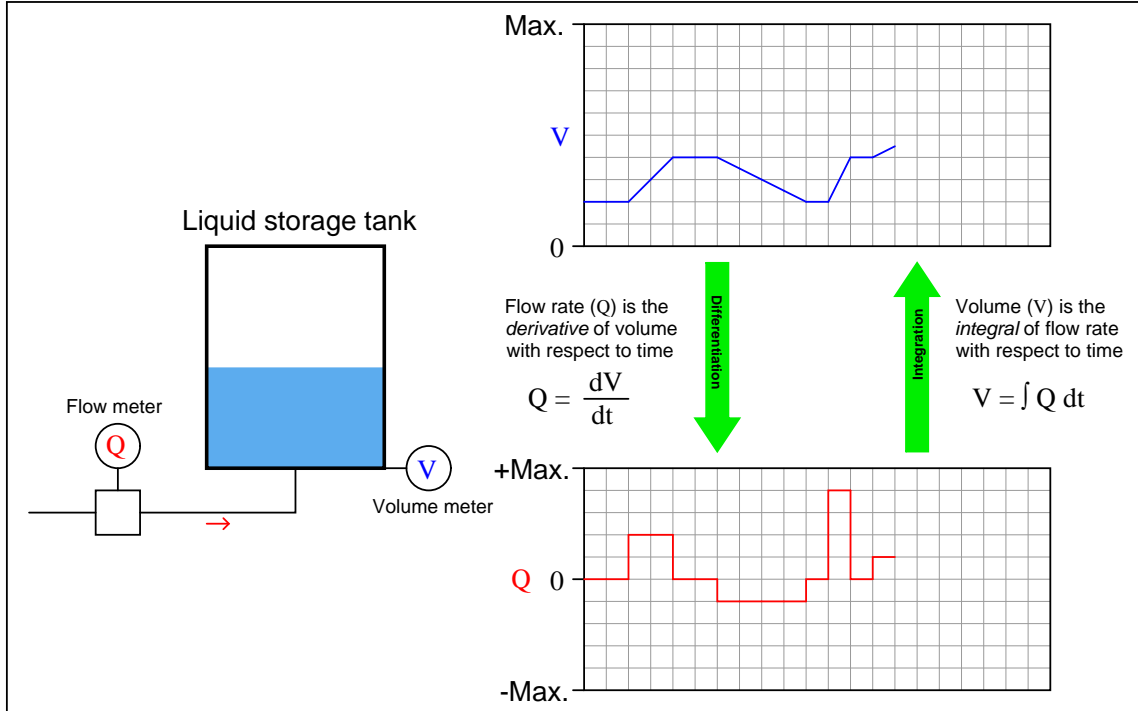


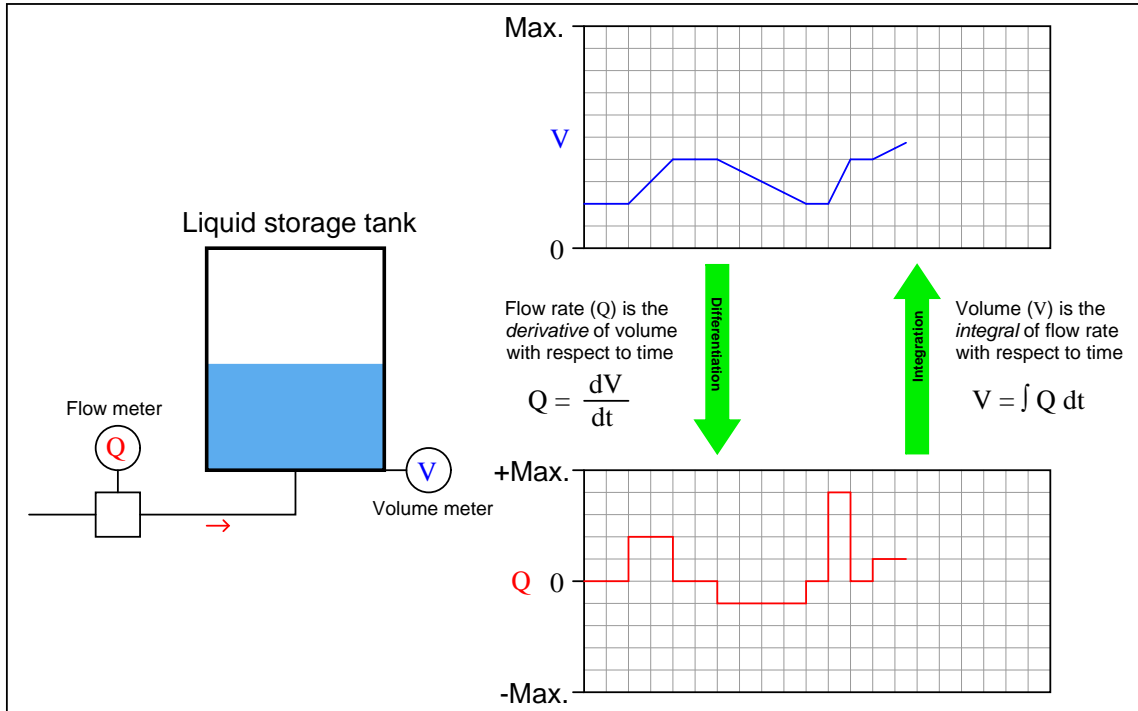














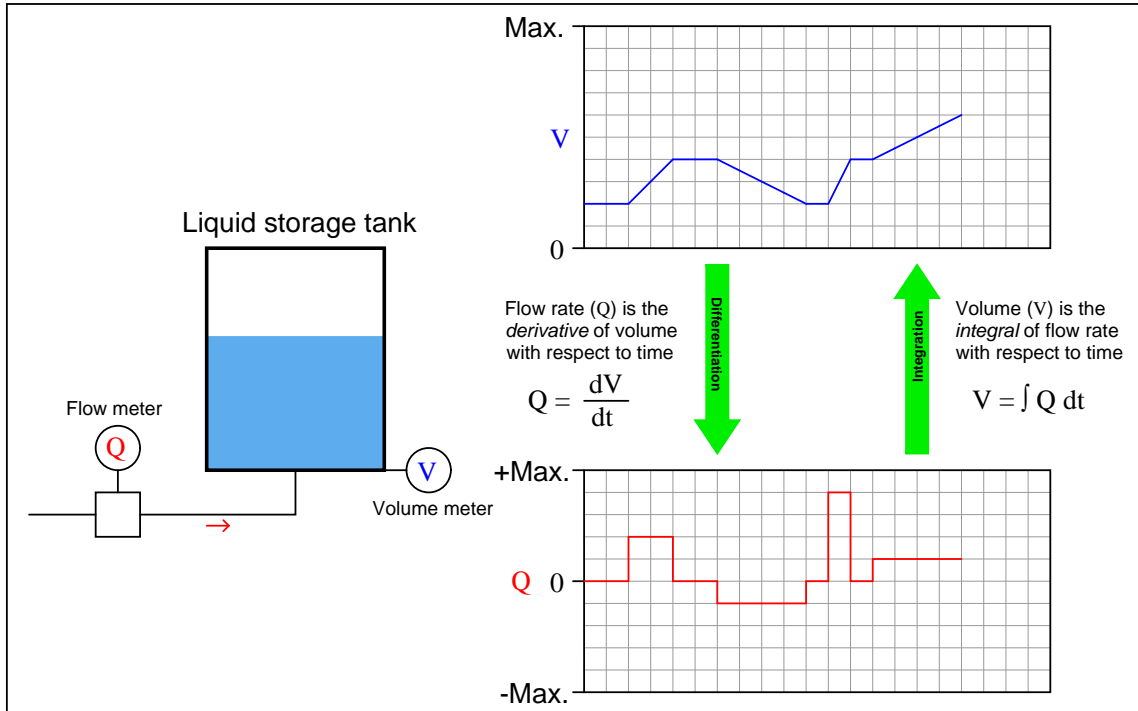














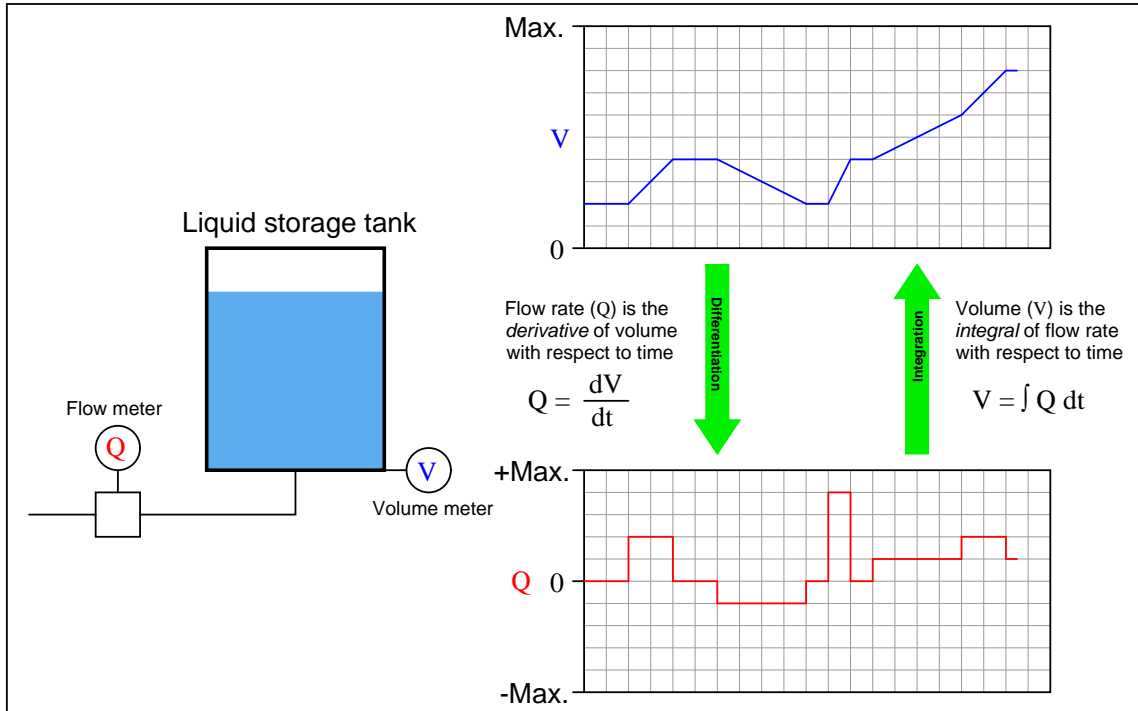


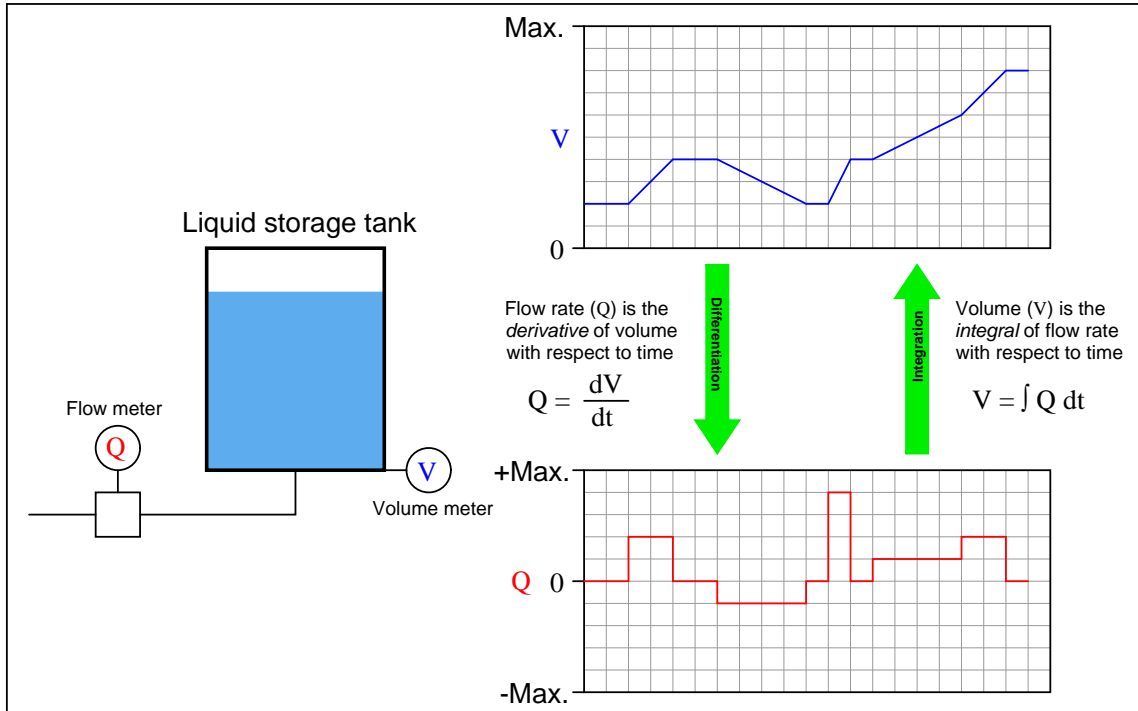





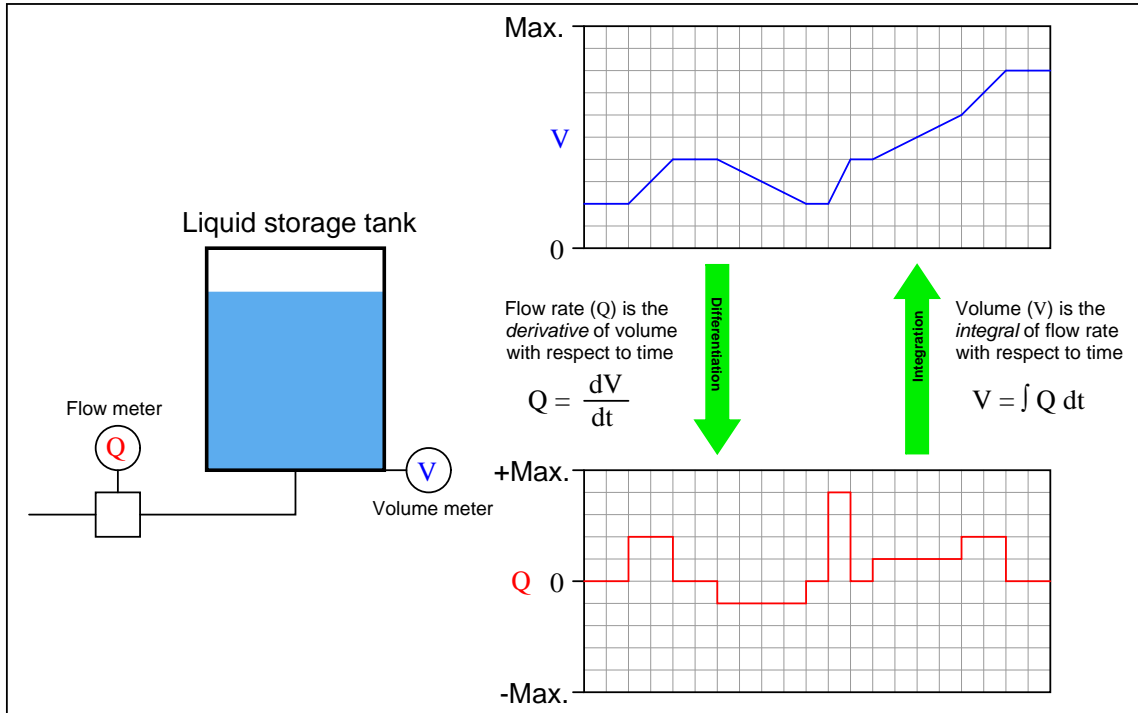


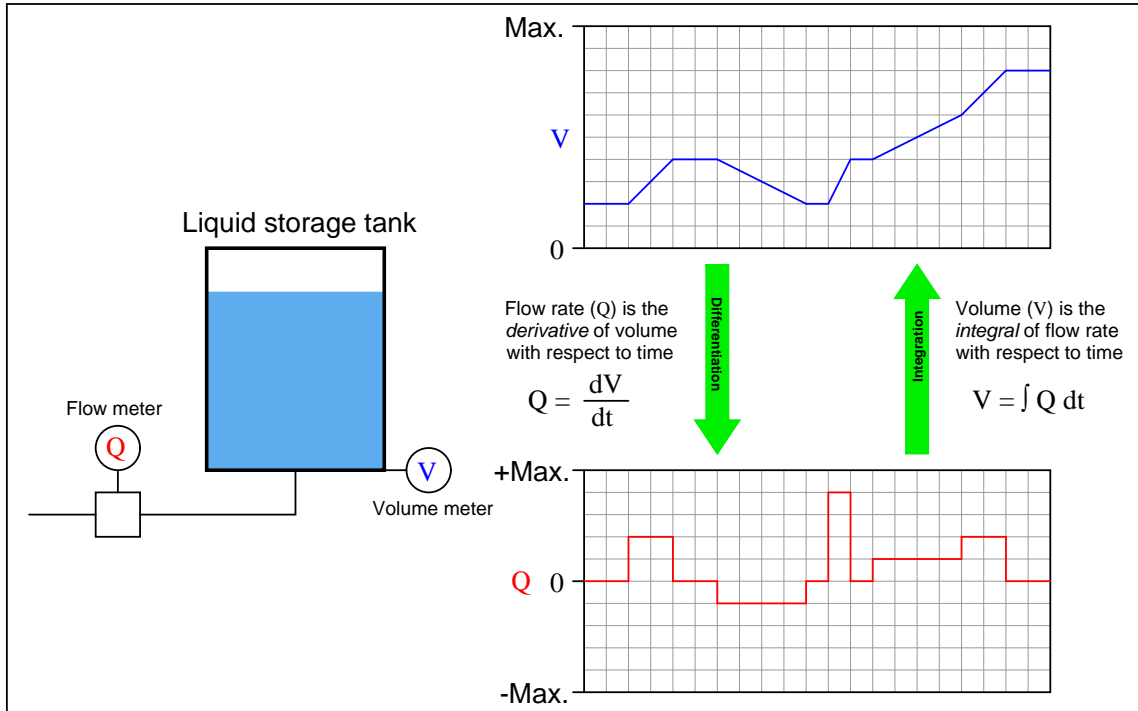













Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time


$$V = \int Q dt$$







Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$





Volume (V) is the integral of flow rate with respect to time


$$V = \int Q dt$$






























Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$





Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$










Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$





Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$










Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height** of the flow graph directly relates to the **slope** of the volume graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$




Volume (V) is the integral of flow rate with respect to time

$$V = \int Q dt$$





Note how the **height increase** of the volume graph directly relates to the **area accumulated** by the flow graph . . .




Flow rate (Q) is the derivative of volume with respect to time


$$Q = \frac{dV}{dt}$$



Volume (V) is the integral of flow rate with respect to time


$$V = \int Q dt$$








A.5 Guided-wave radar level measurement


The following animation shows how a radio-energy pulse travels down and then up the waveguide of a guided-wave radar level instrument, relating the peaks on an *echo curve* to the real-world interfaces inside the process vessel.



























































































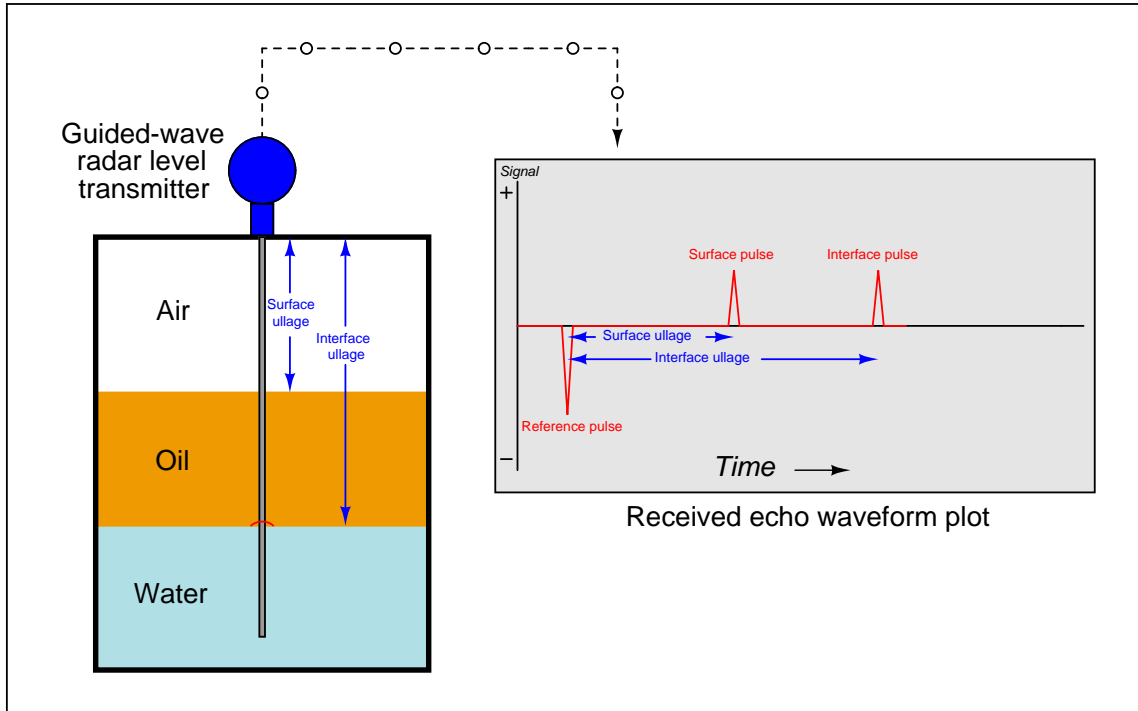



































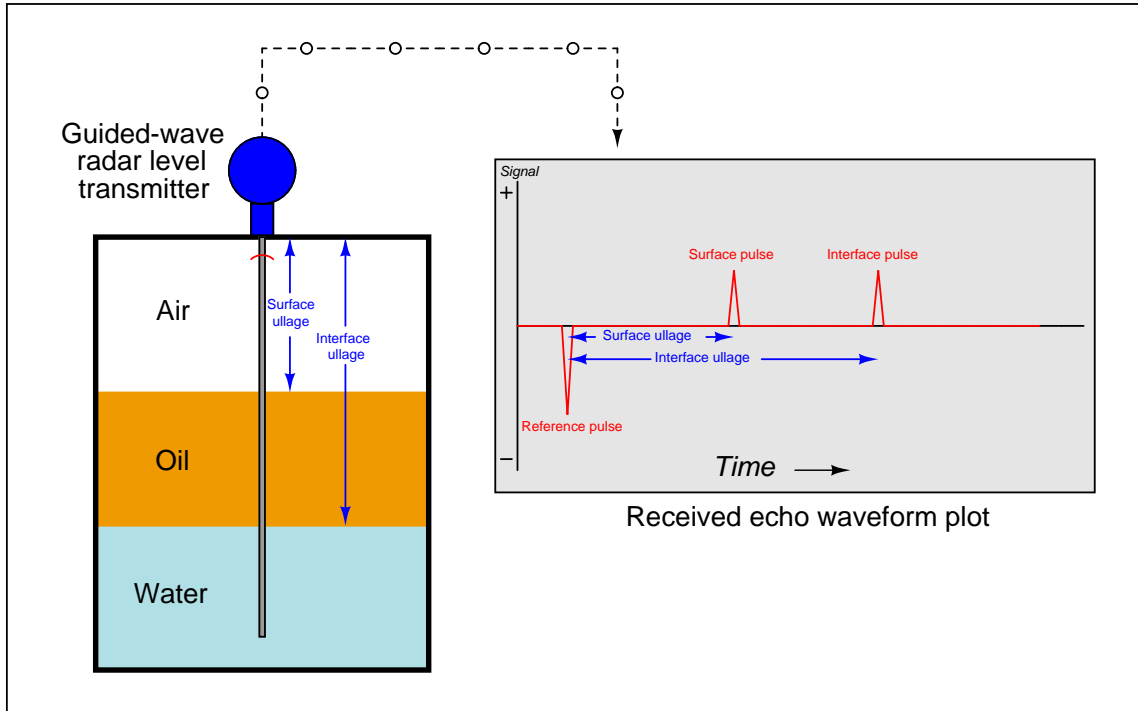


























A.6 Basic chromatograph operation


This animation shows the basic operation of a gas chromatograph, showing the separation of different molecular species in a gas mixture. Each gas type is represented by a different colored dot moving along the tubing.


Carrier gas is represented by orange dots moving constantly through the sample valve and column. Process sample is represented by a cluster of three dots: red (light), green (medium), and blue (heavy) molecules mixed together. These molecules move together at the same rate until they reach the column. There, the light molecules (red) travel fastest, the medium molecules (green) travel slower, and the heavy molecules (blue) travel slowest. Thus, the differing velocities within the chromatograph column performs the task of separation necessary to identify and measure each chemical component in the mixture. All the while, you can see the chromatogram developing, a peak appearing each time one of the components reaches the detector.


Each chemical component (light, medium, heavy) is thus identified by its place in *time* when its peak appears on the chromatogram, while the concentration (quantity) of each component is discernible by the *area* integrated underneath each peak.




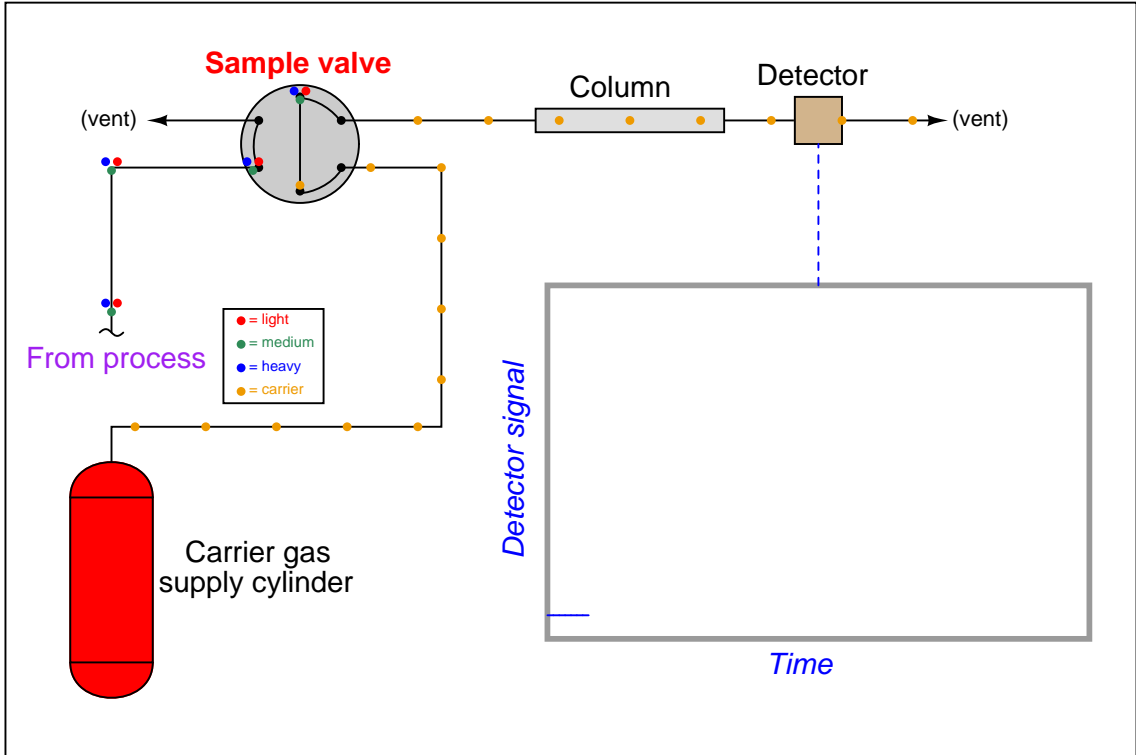




















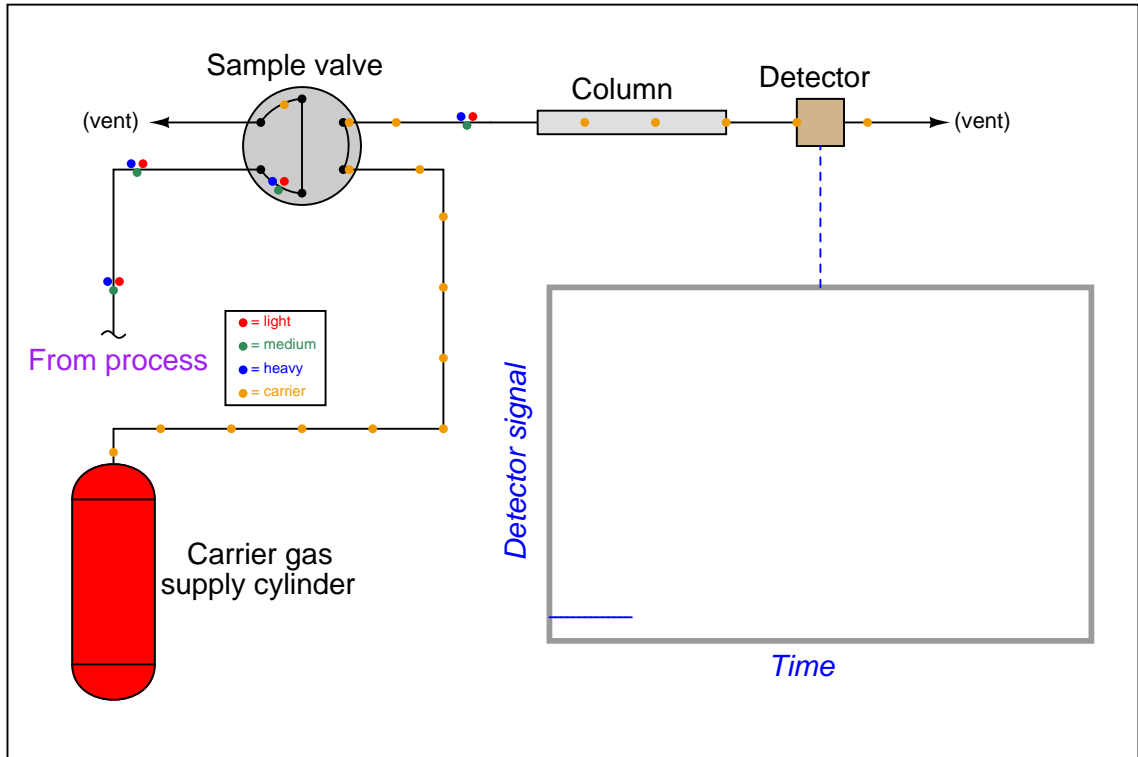
























































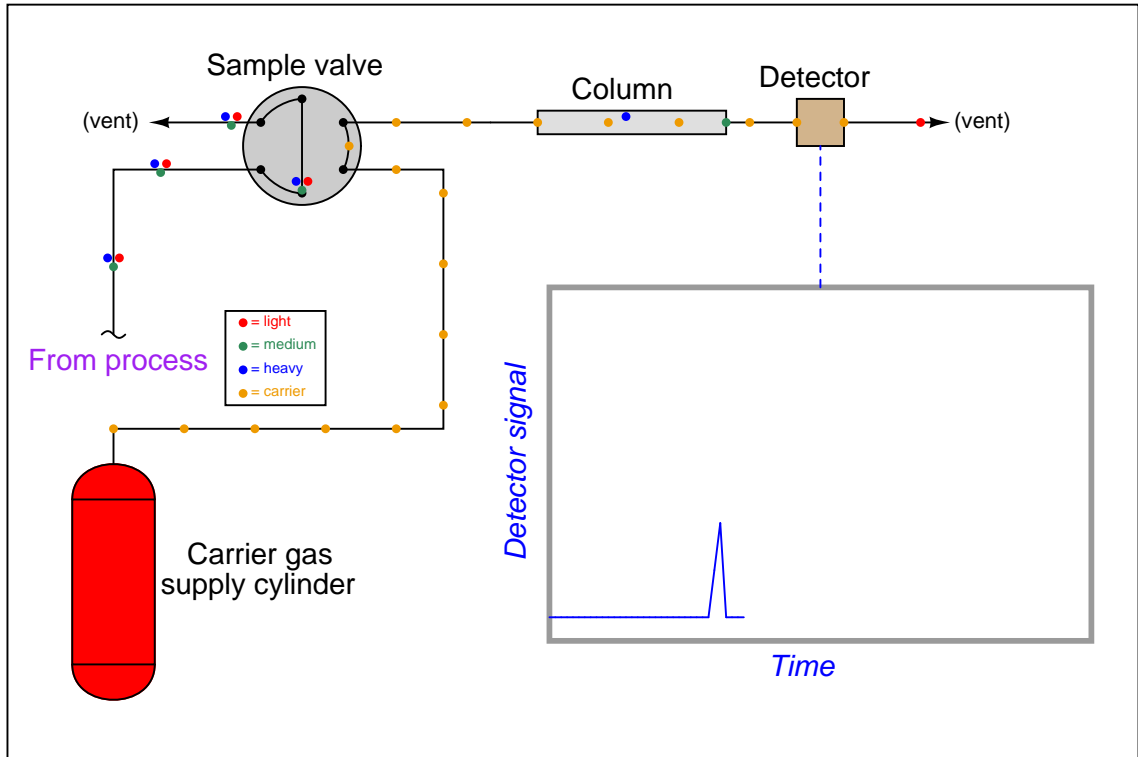





























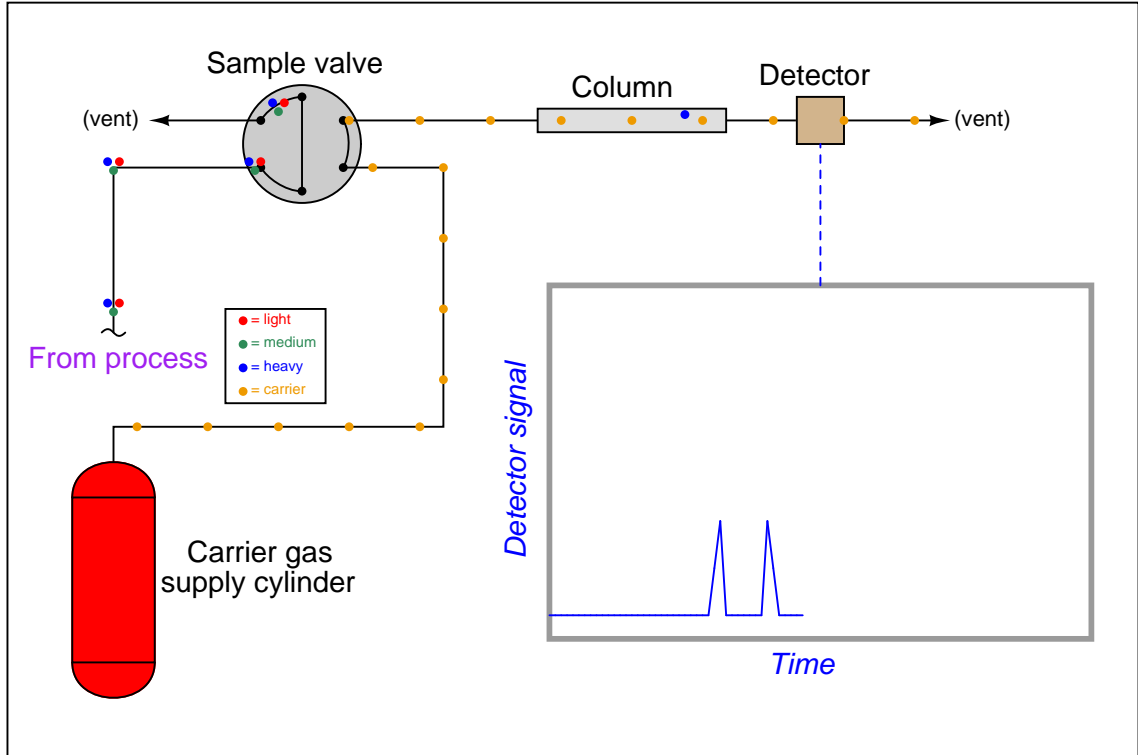


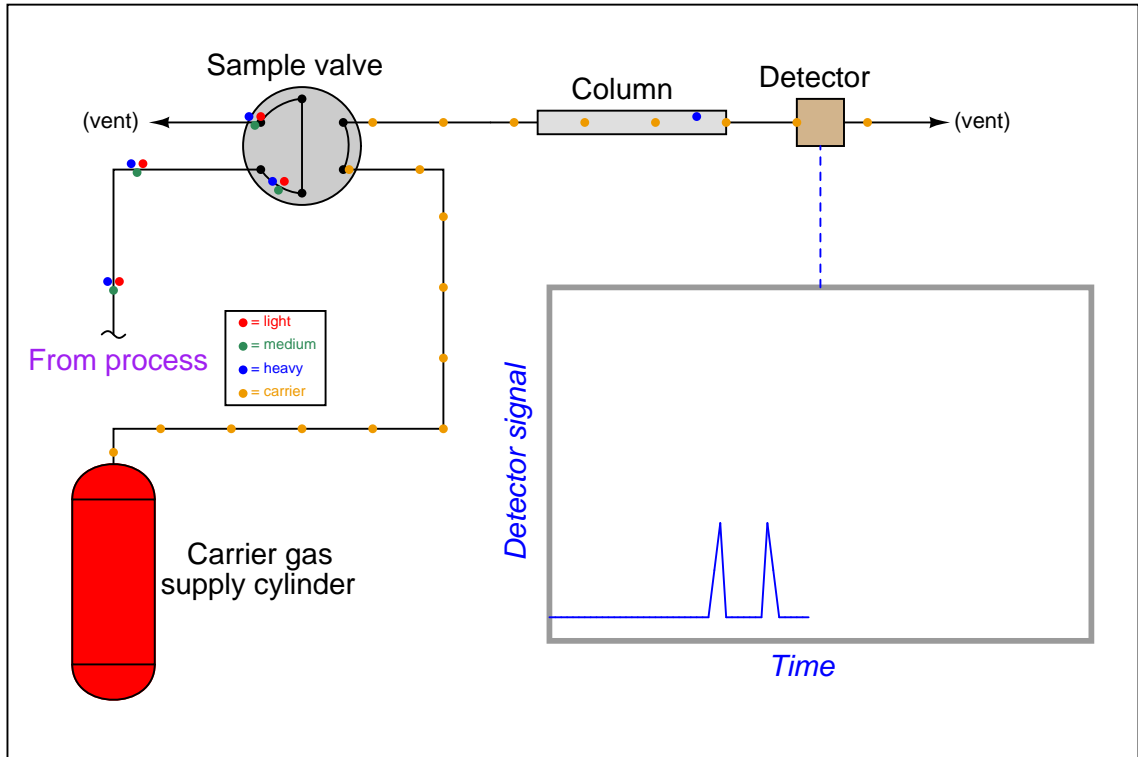





























































Appendix B

Doctor Strangeflow, or how I learned to relax and love Reynolds numbers

Of all the non-analytical (non-chemistry) process measurements students encounter in their Instrumentation training, flow measurement is one of the most mysterious. Where else would we have to *take the square root* of a transmitter signal just to measure a process variable in the simplest case? Since flow measurement is so vital to many industries, it cannot go untouched in an Instrumentation curriculum. Students must learn how to measure flow, and how to do it accurately. The fact that it is a fundamentally complex thing, however, often leads to oversimplification in the classroom. Such was definitely the case in my own education, and it led to a number of misunderstandings that were corrected after a lapse of 15 years, in a sudden “Aha!” moment that I now wish to share with you.

The orifice plate is to flow measurement what a thermocouple is to temperature measurement: an inexpensive yet effective primary sensing element. The concept is disarmingly simple. Place a restriction in a pipe, then measure the resulting pressure drop (ΔP) across that restriction to infer flow rate.

You may have already seen a diagram such as the following, illustrating how an orifice plate works:




Now, the really weird thing about measuring flow this way is that the resulting ΔP signal does not linearly correspond to flow rate. Double the flow rate, and the ΔP quadruples. Triple the flow rate and the ΔP increases by a factor of nine. To express this relationship mathematically:

$$Q^2 \propto \Delta P$$

In other words, differential pressure across an orifice plate (ΔP) is proportional to the *square* of the flow rate (Q^2). To be more precise, we may include a coefficient (k) with a precise value that turns the proportionality into an equality:

$$Q^2 = k(\Delta P)$$

Expressed in graphical form, the function looks like one-half of a parabola:



To obtain a linear flow measurement signal from the differential pressure instrument's output signal, we must "square root" that signal, either with a computer inside the transmitter, with a computer inside the receiving instrument, or a separate computing instrument (a "square root extractor"). We may see mathematically how this yields a value for flow rate (Q), following from our original equation:

$$Q^2 = k(\Delta P)$$

$$\sqrt{Q^2} = \sqrt{k(\Delta P)}$$

$$Q = \sqrt{k(\Delta P)}$$


. . . substituting a new coefficient value k^1 . . .

$$Q = k\sqrt{\Delta P}$$

Students are taught that the differential pressure develops as a consequence of energy conservation in the flowing liquid stream. As the liquid enters a constriction, its velocity must increase to account for the same volumetric rate through a reduced area. This results in kinetic energy increasing, which must be accompanied by a corresponding decrease in potential energy (i.e. pressure) to conserve total fluid energy.

¹Since we get to choose whatever k value we need to make this an equality, we don't have to keep k inside the radicand, and so you will usually see the equation written as it is shown in the last step with k outside the radicand.

Pressure measurements taken in a venturi pipe confirm this:



In all honesty, this did not make sense to me when I heard this. My “common sense” told me the fluid pressure would *increase* as it became crammed into the constriction, not decrease. Even more, “common sense” told me that whatever pressure was lost through the constriction would never be regained, contrary to the pressure indication of the gauge furthest downstream. Accepting this principle was an act of faith on my part, putting preconceived notions aside for something new. A leap of faith, however, is not the same as a leap in understanding. I believed what I was told, but I really didn’t understand *why* it was true.

The problem intensified when my teacher showed a more detailed flow equation. This new equation contained a term for fluid density (ρ):

$$Q = k \sqrt{\frac{\Delta P}{\rho}}$$

What this equation showed us is that orifice plate flow measurement depended on density. If the fluid density changed, our instrument calibration would have to change in order to maintain good accuracy of measurement. Something disturbed me about this equation, though, so I raised my hand. The subsequent exchange between my teacher and I went something like this:

Me: What about viscosity?

Teacher: What?

Me: Doesn’t fluid viscosity have an effect on flow measurement, just like density?

Teacher: You don’t see a variable for viscosity in the equation, do you?

Me: Well, no, but it’s *got* to have some effect on flow measurement!

Teacher: How come?

Me: Imagine clean water flowing through a venturi, or through the hole of an orifice plate. At a certain flow rate, a certain amount of ΔP will develop across the orifice. Now imagine

that same orifice flowing an equal rate of liquid honey: approximately the same density as water, but much thicker. Wouldn't the increased "thickness," or viscosity, of the honey result in more friction through the orifice, and thus more of a pressure drop than what the water would create?

Teacher: I'm sure viscosity has some effect, but it must be minimal since it isn't in the equation.

Me: Then why is honey so hard to suck through a straw?

Teacher: Come again?

Me: A straw is a narrow pipe, similar to the throat of a venturi or the hole of an orifice, right? The difference in pressure between the suction in my mouth and the atmosphere is the ΔP across that orifice. The result is flow through the straw. If viscosity is of such little effect, then why is liquid honey so much harder to suck through a straw than water? The pressure is the same, the density is about the same, then why isn't the flow rate the same according to the equation you just gave us?

Teacher: In industry, we usually don't measure fluids as thick as honey, and so it's safe to ignore viscosity in the flow equation . . .

My teacher's smokescreen – that thick fluid flow streams were rare in industry – did nothing to alleviate my confusion. Despite my ignorance of the industrial world, I could very easily imagine liquids that were more viscous than water, honey or no honey. Somewhere, somehow, someone had to be measuring the flow rate of such liquids, and there the effects of viscosity on orifice ΔP must be apparent. Surely my teacher knew this. But then why did the flow equation not have a variable for viscosity in it? How could this parameter be unimportant? Like most students, though, I could see that arguing would get me nowhere and it was better for my grade to just go along with what the teacher said than to press for answers he couldn't give. In other words, I swept my doubts under the carpet of "learning" and made a leap of faith.

After that, we studied different types of orifice plates, different types of pressure tap locations, and other inferential primary sensing elements (Pitot tubes, target meters, pipe elbows, etc.). They all worked on Bernoulli's principle of decreased pressure through a restriction, and they all required square root extraction of the pressure signal to obtain a linearized flow measurement. In fact, this became the sole criterion for determining whether or not we needed square root extraction on the signal: did the flow measurement originate from a differential pressure instrument? If so, then we needed to "square root" the signal. If not, we didn't. A neat and clean distinction, separating ΔP -based flow measurements from all the others (magnetic, vortex shedding, Coriolis effect, thermal, etc.). Nice, clean, simple, neat, and only 95% correct, as I was to discover later.

Fast-forward fifteen years. I was now a teacher in a technical college, teaching Instrumentation to students just like myself a decade and a half ago. It was my first time preparing to teach flow measurement, and so I brushed up on my knowledge by consulting one of the best technical references I could get my hands on: Béla Lipták's *Process Measurement and Analysis*, third edition. Part of the *Instrument Engineers' Handbook* series, this wonderful work was to be our primary text as we

explored the world of process measurement during the 2002-2003 academic year.

It was in reading this book that I had an epiphany. Section 2.8 of the text discussed a type of flowmeter I had never seen or heard of before: the *laminar* flowmeter. As I read this section of the book, my jaw hit the floor. Here was a differential-pressure-based flowmeter that was linear! That is, there was no square root extraction required at all to convert the ΔP measurement into a flow measurement. Furthermore, its operation was based on some weird equation called the *Hagen-Poiseuille* Law rather than Bernoulli's Law.

Early in the section's discussion of this flowmeter, a couple of paragraphs explained the meaning of something called *Reynolds number* of a flow stream, and how this was critically important to laminar flowmeters. Now, I had heard of Reynolds number before when I worked in industry, but I never knew what it meant. All I knew is that it had something to do with the selection of flowmeter types: one must know the Reynolds number of a fluid before one could properly select which type of flow-measuring instrument to use in a particular application. Since this determination typically fell within the domain of instrument engineers and not instrument technicians (as I was), I gave myself permission to remain ignorant about it and blissfully went on my way. Little did I know that Reynolds number held the key to understanding my "honey-through-a-straw" question of years ago, as well as comprehending (not just believing) how orifice plates actually worked.

According to Lipták, laminar flowmeters were effective only for low Reynolds numbers, typically below 1200. Cross-referencing the orifice plate section of the same book told me that Reynolds numbers for typical orifice-plate flow streams were much greater (10000 or higher). Furthermore, the orifice plate section contained an insightful passage on page 152 which I will now quote here. Italicized words indicate my own emphasis, locating the exact points of my "Aha!" moments:

The basic equations of flow assume that the velocity of flow is uniform across a given cross-section. In practice, flow velocity at any cross section approaches zero in the boundary layer adjacent to the pipe wall, and varies across the diameter. *This flow velocity profile has a significant effect on the relationship between flow velocity and pressure difference developed in a head meter.* In 1883, Sir Osborne Reynolds, an English scientist, presented a paper before the Royal Society, proposing a single, dimensionless ratio now known as Reynolds number, as a criterion to describe this phenomenon. This number, *Re*, is expressed as


$$Re = \frac{VD\rho}{\mu}$$

where *V* is velocity, *D* is diameter, ρ is density, and μ is absolute viscosity. Reynolds number expresses the ratio of inertial forces to viscous forces. At a very low Reynolds number, viscous forces predominate, and the inertial forces have little effect. *Pressure difference approaches direct proportionality to average flow velocity and to viscosity.* At high Reynolds numbers, inertial forces predominate and viscous drag effects become negligible.

What the second paragraph is saying is that for slow-moving, viscous fluids (such as honey in a straw), the forces of friction (fluid "dragging" against the pipe walls) are far greater than the forces of inertia (fluid momentum). This means that the pressure difference required to move such a fluid through a pipe primarily works to overcome the friction of that fluid against the walls of the pipe. For most industrial flows, where the flow velocities are fast and the fluids have little viscosity (like clean water), flow through an orifice plate is assumed to be frictionless. Thus, the pressure dropped across a constriction is *not* the result of friction between the fluid and the pipe, but rather it is a consequence of having to *accelerate* the fluid from a low velocity to a high velocity through the narrow orifice.

My mistake, years ago, was in assuming that water flowing through an orifice generated substantial friction, and that this is what created the ΔP across an orifice plate. This is what my “common sense” told me. In my mind, I imagined the water having to rub past the walls of the pipe, past the face of the orifice plate, and through the constriction of the orifice at a very high speed, in order to make it through to the other side. I memorized what my teacher told us about energy exchange and how pressure had to drop as velocity increased, but I never really internalized it because I still held to my faulty assumption of friction being the dominant mechanism of pressure drop in an orifice plate. In other words, while I could parrot the doctrine of kinetic and potential energy exchange, I was still *thinking* in terms of friction, which is a totally different phenomenon. The difference between these two phenomena is the difference between energy *exchanged* and energy *dissipated*. To use an electrical analogy, it is the difference between *reactance* (X) and *resistance* (R). Incidentally, many electronics students experience the same confusion when they study reactance, mistakenly thinking it is the same thing as resistance where in reality it is quite different in terms of energy, but that is a subject for another essay!

In a frictionless flow stream, fluid pressure decreases as fluid velocity increases in order to conserve energy. Another way to think of this is that a pressure differential must develop in order to provide the “push” needed to *accelerate* the fluid from a low speed to a high speed. Conversely, as the fluid slows back down after having passed through the constriction, a reverse pressure differential must develop in order to provide the “push” needed for that *deceleration*:



A moving mass does not simply slow down on its own! There must be some opposing force to decelerate a mass from a high speed to a low speed. This is where the pressure recovery downstream of the orifice plate comes from. If the pressure differential across an orifice plate originated primarily from friction, as I mistakenly assumed when I first learned about orifice plates, then there would be no reason for the pressure to *ever* recover downstream of the constriction. The presence of friction means energy *lost*, not energy *exchanged*. Although both inertia and friction are capable of creating pressure drops, the lasting effects of these two different phenomena are definitely not the same.

There is a quadratic (“square”) relationship between velocity and differential pressure precisely because there is a quadratic relationship between velocity and kinetic energy as all first-quarter physics students learn ($E_k = \frac{1}{2}mv^2$). This is why ΔP increases with the square of flow rate (Q^2)

and why we must “square-root” the ΔP signal to obtain a flow measurement. This is also why fluid density is so important in the orifice-plate flow equation. The denser a fluid is, the more work will be required to accelerate it through a constriction, resulting in greater ΔP , all other conditions being equal:

$$Q = k \sqrt{\frac{\Delta P}{\rho}} \quad (\text{Our old friend, the “orifice plate” equation})$$

This equation is only accurate, however, when fluid friction is negligible: when the viscosity of the fluid is so low and/or its speed is so high that the effects of potential and kinetic energy exchange completely overshadow² the effects of friction against the pipe walls and against the orifice plate. This is indeed the case for most industrial flow applications, and so this is what students first study as they learn how flow is measured. Unfortunately, this is often the *only* equation two-year Instrumentation students study with regard to flow measurement.

In situations where Reynolds number is low, fluid friction becomes the dominant factor and the standard “orifice plate” equation no longer applies. Here, the ΔP generated by a viscous fluid moving through a pipe really does depend primarily on how “thick” the fluid is. And, just like electrons moving through a resistor in an electric circuit, the pressure drop across the area of friction is directly proportional to the rate of flow ($\Delta P \propto Q$ for fluids, $V \propto I$ for electrons). This is why laminar flowmeters – which work only when Reynolds number is low – yield a nice *linear* relationship between ΔP and flow rate and therefore do not require square root extraction of the ΔP signal. These flowmeters do, however, require temperature compensation (and even temperature *control* in some cases) because flow measurement accuracy depends on fluid viscosity, and fluid viscosity varies according to temperature. The Hagen-Poiseuille equation describing flow rate and differential pressure for laminar flow (low Re) is shown here for comparison:

$$Q = k \left(\frac{\Delta P D^4}{\mu L} \right)$$

Where,

Q = Flow rate (gallons per minute)

k = Unit conversion factor = 7.86×10^5

ΔP = Pressure drop (inches of water column)

D = Pipe diameter (inches)

μ = Liquid viscosity (centipoise) – this is a temperature-dependent variable!

L = Length of pipe section (inches)

Note that if the pipe dimensions and fluid viscosity are held constant, the relationship between flow and differential pressure is a direct proportion:

$$Q \propto \Delta P$$

²In engineering, this goes by the romantic name of *swamping*. We say that the overshadowing effect “swamps” out all others because of its vastly superior magnitude, and so it is safe (not to mention simpler!) to ignore the smaller effect(s). The most elegant cases of “swamping” are when an engineer intentionally designs a system so the desired effect is many times greater than the undesired effect(s), thereby forcing the system to behave more like the ideal. This application of swamping is prevalent in electrical engineering, where resistors are often added to circuits for the purpose of overshadowing the effects of stray (undesirable) resistance in wiring and components.

In reality, there is no such thing as a frictionless flow (excepting superfluidic cases such as helium II which are well outside the bounds of normal experience), just as there is no such thing as a massless flow (no inertia). In normal applications there will always be both effects at work. By not considering fluid friction for high Reynolds numbers and not considering fluid density for low Reynolds numbers, engineers draw simplified models of reality which allow us to more easily measure fluid flow. As in so many other areas of study, we exchange accuracy for simplicity, precision for convenience. Problems arise when we forget that we've made this Faustian exchange and wander into areas where our simplistic models are no longer reasonable.

Perhaps the most practical upshot of all this for students of Instrumentation is to realize exactly why and how orifice plates work. Bernoulli's equation does *not* include any considerations of friction. To the contrary, we must assume the fluid to be completely frictionless in order for the concept to make sense. This explains several things:

- There is pressure recovery downstream of an orifice: most of the pressure lost at the vena contracta is regained further on downstream as the fluid decelerates to its original (slow) speed. Permanent pressure drop will occur only where there is energy *lost* through the constriction, such as in cases where fluid friction is substantial. Where the fluid is frictionless there is no mechanism in an orifice to dissipate energy, and so with no energy lost there must be full pressure recovery as the fluid returns to its original speed.
- Pressure tap location makes a difference: to ensure that the downstream tap is actually sensing the pressure at a point where the fluid is moving significantly faster than upstream (the “vena contracta”), and not just anywhere downstream of the orifice. If the pressure drop were due to friction alone, it would be permanent and the downstream tap location would not be as critical.
- Standard orifice plates have knife-edges on their upstream sides: to minimize contact area (friction points) with the high-speed flow.
- Care must be taken to ensure Reynolds number is high enough to permit the use of an orifice plate: if not, the linear $Q/\Delta P$ relationship for viscous flow will assert itself along with the quadratic potential/kinetic energy relationship, causing the overall $Q/\Delta P$ relationship to be polynomial rather than purely quadratic, and thereby corrupting the measurement accuracy.
- Sufficient upstream pipe length is needed to condition flow for orifice plate measurement, not to make it “laminar” as is popularly (and wrongly) believed, but to allow natural turbulence to “flatten” the flow profile for uniform velocity. *Laminar flow* is something that only happens when viscous forces overshadow inertial forces (e.g. flow at low Reynolds numbers), and is totally different from the *fully developed turbulent flow* that orifice plates need for accurate measurement.

In a more general sense, the lesson we should learn here is that blind faith is no substitute for understanding, and that a sense of confusion or disagreement during the learning process is a sign of one or more misconceptions in need of correction. If you find yourself disagreeing with what you are being taught, either you are making a mistake and/or your teacher is. Pursuing your questions to their logical end is the key to discovery, while making a leap of faith (simply believing what you are told) is an act of avoidance: escaping the discomfort of confusion and uncertainty at the expense of a deeper learning experience. This is an exchange no student should ever feel they must make.

References

Lipták, Béla G. et al., *Instrument Engineers' Handbook – Process Measurement and Analysis Volume I*, Third Edition, CRC Press, New York, NY.

Appendix C

Disassembly of a sliding-stem control valve

The following collection of photographs chronicles the complete disassembly of a Fisher E-body globe valve with pneumatic diaphragm actuator. This control valve design is quite mature, but nevertheless enjoys wide application in modern industrial settings.

An important safety note when disassembling pneumatic control valves is to first relieve all tension from the actuator spring so that its stored energy cannot harm you or anyone else. These springs may be quite large, exerting *thousands of pounds* of force during normal operation.

Spring tension may be relieved by moving the spring adjuster until it turns easily by hand without further aid of tools, or in the procedure shown in the following photographs by loosening the spanner nut attaching the actuator yoke to the valve bonnet.

This is the complete control valve, without a positioner attached. What you see here is the actuator (painted green) and the valve body (painted grey), mounted on a steel plate for student learning in a laboratory setting. The left-hand photograph shows the complete control valve assembly, while the right-hand photograph shows a student loosening the spanner nut holding the valve actuator yoke to the valve body:




The next step is to un-couple the actuator stem from the valve stem. On Fisher sliding-stem valves, this connection is made by a split block with threads matching those on each stem. Removing two bolts from the block allows it to be taken apart (left-hand photograph). Nuts threaded on to the valve stem, jammed up against the coupling block, must also be loosened before the stems may be uncoupled (right-hand photograph):



It is very important that no spring tension exists on the stem prior to disassembly of the stem coupler, or else the two stems will slip past each other with great force once the coupler is removed. Spring tension must be released, either by loosening the spring adjuster or by loosening the spanner nut holding the actuator yoke to the bonnet.

A close-up photograph of this stem connector block, with the front half removed for inspection, shows how it engages both threaded stems (valve and actuator) in a single nut-like assemblage. The solid valve stem (below) slides into the hollow actuator stem (above), while the split connector “nut” engages the threads of both, holding the two stems together so they move up and down as one piece:




Once the actuator and valve body stems have been uncoupled, the actuator may be removed from the valve body entirely:



The bonnet is held to the rest of the valve body (in this case) by four large studs. Removing the nuts on these studs allows the bonnet to be lifted off the body, exposing the valve trim for view:



Seats in Fisher E-body globe valves rest in the bottom of the body, held in place by the cage surrounding the valve plug. Once the bonnet is removed from the body, the seat may be removed without need of any specialized tools (left-hand photograph). A view inside the body shows the place where the seat normally rests (right-hand photograph):



With the bonnet removed, the plug and cage may be easily removed for inspection:



The packing follower (between the student's fingers) has been removed from the valve bonnet, and you can see the upper Teflon packing rings within the bonnet. The student is also holding the packing flange in the same hand as the packing follower (left-hand photograph). In the right-hand photograph, we see the student using a screwdriver to gently push the Teflon packing rings out of the bonnet, from the bottom side. Care should be taken not to gouge or otherwise damage these rings during removal:



The left-hand photograph shows all the packing components stacked on top of each other on the concrete floor, next to the bonnet. From top to bottom you see the following components: a felt wiper, the packing follower, five (5) Teflon packing rings, a coil spring, and the packing box ring. The right-hand photograph shows the same packing components stacked on the valve stem:




Turning to the actuator, we begin disassembly by loosening the diaphragm hold-down bolts (left-hand photograph) and removing the upper half of the diaphragm casing (right-hand photograph). A single bolt secures the upper diaphragm plate to the top of the actuator stem:



In the left-hand photograph you see the student removing the spring seat, having previously loosened the spring adjuster nut. With the spring seat removed, the spring may be removed from the actuator assembly. In the right-hand photograph the spring adjuster and spring seat have been removed from the actuator stem. The student is now pointing at the valve spring, partially removed:



Sliding the actuator diaphragm, plate, and stem out of the actuator assembly from the top of the actuator makes it easy to remove the large actuator spring (left-hand photograph). The right-hand photograph shows all the moving actuator components re-assembled in their proper order outside of the yoke:



The left-hand photograph shows the lower half of the actuator casing, with the student removing six (6) hold-down bolts joining this casing half to the actuator yoke. The right-hand photograph shows the actuator casing half completely removed from the yoke, revealing a gasket and the bronze stem bushing (which serves to both guide the actuator stem and seal air pressure, since this is a reverse-acting actuator):



A circular spring clip holds the stem bushing in the yoke casting. The left-hand photograph shows the student using pliers to squeeze this spring clip and remove it from its groove cut into the metal of the yoke. In the right-hand photograph, we see the student using the wooden handle of a hammer to gently tap the bushing out of the yoke. The bushing has rubber O-ring seals between it and the yoke casting, so a small amount of force will be necessary to dislodge it. Using the hammer's wooden handle to drive the bushing instead of a metal tool protects the relatively soft bronze bushing from impact damage. Note how the student's right hand is waiting to catch the bronze bushing when it emerges from the hole, to protect it from falling against the hard concrete floor:



The final photograph shows the bushing removed from its hole:



Appendix D

How to use this book – some advice for teachers

If you would like to maximize your students' learning in a field of study that emphasizes critical thinking as much as Instrumentation, I have one simple piece of advice: *engage your students, don't just present information to them*. Do not make the mistake so many teachers do, of thinking it is their role in the classroom to provide information in pre-digested form to their students, and that it is each student's responsibility to passively absorb this information.

High achievement happens only in an atmosphere of high expectations. If you design coursework allowing students to expend minimal effort, your students will achieve minimal learning. Alternatively, if you require students to think deeply about their subject of study, challenge them with interesting and relevant assignments, and hold them accountable to rigorous standards of demonstrated competence, your students can and will move mountains.

In this appendix I present to you some concepts and models for achieving high standards of learning in the field of Instrumentation. The ideas documented here have all been proven to work in my own instruction, and I continue to use them on a daily basis. However, this is not a rigid blueprint for success – I invite and encourage others to experiment with variations on the same themes. More than anything else, I hope to encourage educators with examples of unconventional thinking and unconventional curricula, to show what may be accomplished if you allow yourself to be creative and results-driven in your instructional design.

D.1 Teaching technical theory

Learning is not merely a process of information transfer. It is first and foremost a transformation of one's thinking. When we learn something substantial, it alters the way we perceive and interact with the world around us. Learning any subject also involves a substantial accumulation of facts in one's memory, but memorization alone is not really learning (at least it isn't learning at the college level). Transmitting facts into a student's memory is easy, and does not require a live human presenter. A well-written book or a well-edited video can do a far better and far more consistent job of conveying facts and concepts than any live presentation¹. The greatest value of a human teacher is to help students develop higher-order thinking skills. This includes (but is not limited to) problem-solving, logical reasoning, diagnostic techniques, and metacognition (critiquing one's own thinking).

Rather than devote most of your classroom time to lecture-style presentations – where the flow of information goes primarily from you to your students – place the responsibility for fact-gathering on your students. Have them read books such as this² and arrive at the classroom *prepared* to discuss what they have already studied.

When students are with you in the classroom and in the lab, probe their understanding with questions – lots of questions. Give them realistic problems to solve. Challenge them with projects requiring creative thought. Get your students to reveal how they think, both to you and to their peers. This will transform your classroom atmosphere from a monologue into a dialogue, where you engage with the minds of your students as partners in the learning process instead of lecturing to them as subordinates.

A format I have used with great success is to assign homework exploring new topics, so students much research those topics in advance of our coverage of it in class. The pedagogical term for this is an *inverted classroom*, where communication of facts occurs outside of class, and higher-order activities such as problem-solving occur inside the classroom. This stands in contrast to traditional learning structures, where the instructor spends most of the class time transmitting facts and working example problems, while subsequent homework questions (ostensibly completed on the students' own time) stimulate the development of problem-solving skills.

Homework in my “inverted” classroom comes in the form of question sets designed to lead the student on a path to acquiring the necessary facts and exposing them to certain problem-solving techniques. Some of these questions point students directly to specific texts to read, while others allow students to choose their own research material. Many of the questions are simply conceptual or quantitative problems to solve, without reference to source materials. When my students arrive for class, they first take a quiz on the material they should have studied for that day. After the quiz, students working alone or in small groups solving the assigned problems. My role as instructor is to assist students in their problem-solving efforts, observing the students' attempts, offering advice, and helping to identify and correct misconceptions. Students are much more engaged, less distracted,

¹To be sure, there are some gifted lecturers in the world. However, rather than rely on a human being's live performance, it is better to capture the brilliance of an excellent presentation in static form where it may be peer-reviewed and edited to perfection, then placed into the hands of an unlimited number of students in perpetuity. In other words, if you think you're great at explaining things, do us all a favor and translate that brilliance into a format capable of reaching more people!

²It would be arrogant of me to suggest my book is the best source of information for your students. Have them research information on instrumentation from other textbooks, from manufacturers' literature, from whitepapers, from reference manuals, from encyclopedia sets, or whatever source(s) you deem most appropriate. If you possess knowledge that your students need to know that isn't readily found in any book, *publish it for everyone's benefit!*

and comfortable raising questions while working together in these groups than they ever are as one large group.

Students are considered finished with the class session (and free to leave) when they are able to successfully demonstrate to me their grasp of the day's material. This happens in the form of a "summary quiz" which may be done one-on-one (my preference) or as a whole class. If done individually, the summary quizzes must be varied enough so I am adequately challenging each student even though they have overheard classmates answering similar summary quiz questions.

An inverted classroom structure shifts the burden of transmitting facts and concepts from a live teacher to static sources such as textbooks³. This shift in responsibility frees valuable class time for more important tasks, namely the refinement of higher-order thinking skills. It makes no sense to have a subject-matter expert (the instructor) spend most or all of the students' valuable time at school doing what a book or a video could do just as well. Any instructor who can be replaced with a book or a video *should* be replaced by a book or a video! It is far better to apply that same subject-matter expertise to challenges no book or video can meet: actively identifying student misconceptions, dispensing targeted advice for overcoming difficulties, and stimulating students' minds with follow-up questions designed to illuminate concepts in further detail.

Several important advantages are realized by managing a classroom in this way. First, the instructor gets to have a very close-up and personal view of each students' comprehension, struggles, and misconceptions. If you are accustomed to teaching in a lecture or other "stand-up-in-front" classroom format, you will be utterly amazed to see what your students do and do not comprehend when you watch them dialogue and problem-solve in small groups. Much of what goes on inside your students' minds is hidden from you when they are seated in neat rows watching you in the front of the classroom. When students are free to work together in more intimate settings, you get to see how they think, what they understand, and most importantly what they mis-understand.

Another important advantage of an inverted classroom is how it maximizes your contact with those students who need it most. Faster students are able to finish their work quickly and leave (or stay in the class to tutor their peers), while the slower students stay to the very end with you until their work is done at their pace. *When conversing with small groups (3 to 4 students) in the classroom, I spend an average of only 5 to 6 minutes per student to query them on their understanding of the day's material and to challenge them with at least one problem to solve.* The time saved is incredibly useful, as it allows you to re-structure the topic coverage to include more review, cover additional material, or devote more time to hands-on labwork and projects.

Perhaps the most important benefit of an inverted classroom is that students learn how to independently research, which is no small feat. In a complex field where technology advances on a daily basis, your students will need to be able to learn new facts on their own (without your assistance!) after they graduate. Employers have consistently advised me that this is the single most important skill any person can learn in school: how to independently acquire new knowledge and new abilities. Such a skill not only prepares them for excellence in their chosen career, but it also brings great benefit to every other area of life where the acquisition of new information is essential to decision-making (e.g. participation in the democratic process, legal proceedings, medical decision-making, investing, parenting, etc.). In this way, the inverted classroom is not just a "better mousetrap" but in fact is really catching a "better mouse."

³And multimedia resources, too! With all the advances in multimedia presentations, there is no reason why an instructor cannot build a library of videos, computer simulations, and other engaging resources to present facts and concepts to students outside of class time.

D.1.1 The problem with lecture

I speak negatively of lecture as a teaching tool because I have suffered its ineffectiveness from a teacher's perspective for several years. The problem is not that students cannot learn from an instructor's eloquent presentation; it's that following the lead of an expert's presentation obscures students' perception of their own learning. Stated in simple terms, *lecture forces every student into the role of spectator when they should be participants*. Students observing a lecture cannot tell with certainty whether they are actually learning from an expert presentation, or whether they are merely being stimulated. This is not an obvious concept to grasp, so allow me to elaborate in more detail.

When I began my professional teaching career, I did what every other teacher I knew did: I lectured daily to my students. My goal was to transfer knowledge into my students' minds, and so I chose the most direct method I knew for that necessary transference of information.

My first year of teaching, like most teachers' first year, was a trial by fire. Many days I was lecturing to my students on some subject I had just reviewed (for the first time in a long time) no more than a few days before. My lesson plans were chaotic to non-existent. By my second and third years, however, I had developed lesson plans and was thereby able to orchestrate my lectures much more efficiently. These lesson plans were complete enough to support live demonstrations of concepts during almost every lecture, listing all the materials, components, and equipment I would need to set up in preparation. If an extensive amount of set-up was required for some demonstration, instructions would be found in lesson plan(s) multiple days in advance in order to give me adequate time. The result was a very smooth and polished presentation in nearly every one of my lectures. I was quite proud of the work I had done.

However, I noticed a strange and wholly unintended consequence of all this preparation: with each passing year, my students' long-term recall of concepts presented in lecture seemed to grow worse. Even my best students, who demonstrated an obvious commitment to their education by their regular study habits, outstanding attendance, and quality work, would shock me by asking me to re-explain basic concepts we had covered in extensive detail months before. They never complained about the lectures being bad – quite to the contrary, their assessment of my lectures was always “excellent” in my performance surveys.

An increasingly common lament of students as they tried to do the homework was *“I understand things perfectly when you lecture, but for some reason I just can't seem to figure it out on my own.”* This baffled me, since I had made my presentations as clear as I could, and students seemed engaged and attentive throughout. It was clear to me as I later worked with these students that often they were missing crucial concepts and/or harbored severe misconceptions, and that there was *no way* things should have made sense to them during lecture given these misconceptions.

Another detail that caught my attention was the fine condition of their textbooks. In fact, their textbooks were looking better and better with each passing year. At the conclusion of my first teaching year, my students' textbooks looked as though they had been dragged behind a moving vehicle: pages wrinkled, binding worn, and marks scribbled throughout the pages. As my lectures became more polished, the textbooks appeared less and less used. My reading assignments were no less thorough than before, so why should the books be used less?

One day I overheard a student's comment that made sense of it all. I was working in my office, and just outside my door were two students conversing who didn't think I could hear them talking. One of them said to the other, *“Isn't this class the best? The lectures are so good, you don't even*

have to read the book!” At the sound of this, my heart sunk. I began to realize what the problem was, what was needed to fix it, and how I had unwittingly created a poor learning environment despite the best of intentions.

The fundamental problem is this: students observing an expert presentation are fooled into thinking the concepts are easier to grasp and the processes easier to execute than they actually are. The mastery and polish of the lecturer actually hinders student learning by veiling the difficulty of the tasks. Matters are no different watching a professional athlete or musician at work: a master makes any task look effortless. It isn’t until you (the spectator) actually try to do the same thing (as a participant) that you realize just how challenging it is, just how much you have to learn, and how much effort you must invest before you achieve a comparable level of proficiency.

When students told me “for some reason” they just couldn’t seem to solve the same problems I did during lecture “even though they understood it perfectly” as I lectured, they were being honest. This was not some excuse made up to cover a lack of effort. From their perspective, they truly believed they grasped the concepts while watching me work through them in front of class, and were genuinely mystified why it was so hard for them to perform the same problem-solving tasks on their own.

The simple fact of the matter was that my students did *not* actually grasp the concepts as they watched me lecture. If they had, the solution of similar problems after lecture should have presented little trouble for them. Lecture had generated a *false sense of understanding* in their minds. This made my lectures worse than useless, for not only did they fail to convey the necessary knowledge and skill, but they actually *created an illusion of proficiency* in the minds of my students powerful enough to convince them they did not need to explore the concepts further (by reading their textbooks). This served to hinder learning rather than foster learning.

What I needed to do was shatter this illusion if my students were to learn from me more effectively, and especially if they were ever to become independent learners. Thus began my own personal quest of educational reform.

D.1.2 A more accurate model of learning

A humbling fact every teacher eventually learns is that the depth of a student’s learning is primarily a function of the student’s effort and not their own. Even the most dedicated and talented instructor cannot make a student learn if that student does not invest the necessary time and effort. Conversely, even an unmotivated or incompetent instructor cannot prevent a self-dedicated student from learning on their own.

However, a great many students enter college with the belief that learning is a passive activity: “*It’s the instructor’s job to give us information – all we’re supposed to do is attend and observe.*” Unfortunately, this flawed model of learning seems embedded in modern American culture, anchored in students’ minds from years of compulsory lecture-based education. The role of teacher as expert presenter is so relentlessly reinforced that we have difficulty recognizing its flaws, much less conceiving better alternatives. Teachers choosing to depart from this model invite suspicion and even anger from students accustomed to the status quo of lecture.

One way to help see past one’s own biases on a subject is to consider the same (or similar) subject in a different context. Here, the absurdity of passive learning becomes obvious if we simply switch contexts from academic instruction to athletic instruction. It would be laughable if a coach or fitness trainer were the one performing all the weight-lifting, sprints, stretching, and practice movements while the student never did anything but observe. It would be only slightly less humorous if the trainer spent the whole of every session modeling these activities, leaving the student to practice those activities on their own time as “homework.” Instead, effective physical training sessions *always* place the student in an active role as soon and as often as possible, so that the instructor’s valuable expertise may be applied toward identifying errors and recommending corrections. Instructor-led demonstration is minimized in order to maximize time spent with the student *practicing* their craft.

Academic learning really isn’t much different. If we want students to learn new skills and acquire new mental abilities, students must *practice* those skills and mental processes, and it is during this practice time that an instructor’s expertise becomes most valuable. Accurate self-perception is another reason to immediately engage students in practicing the skills and processes they seek to learn. When students directly experience just how challenging any concept or task is to master, they immediately recognize how much they have to learn. This self-awareness is a vital first step to learning, proving the need for committed action on their part. Only after this recognition is the student psychologically prepared for the hard work of learning.

This is why I favor the “inverted classroom” approach. Students must engage with the new subject(s) *prior* to every classroom session. From the very moment they arrive, they recognize the challenges of the subject matter, and where they need help understanding it. With the presentation of facts occurring before class time, the bulk of that time may be spent actually *applying* the concepts instead of *encountering* them. The relatively mundane task of fact-gathering is relegated to students’ time outside of class, while the more challenging and meaningful tasks of problem-solving and analysis happen where the instructor can actually observe and coach.

In order for an inverted classroom structure to work, though, each student must have access to the necessary information in static form (e.g. textbooks, videos, etc.), and be held accountable for doing the necessary preparations. Access and accountability are absolutely essential to ensuring an inverted classroom will work. Without the former, students will become frustrated; without the latter, some students won’t engage.

D.1.3 The ultimate goal of education

As I have transitioned from a traditional lecturer to a “reform-minded” educator, my general philosophy of education has shifted. My teaching techniques and classroom organization changes preceded this philosophical shift, to be sure, but more and more I am realizing just how important it is to have an educational philosophy, and how such a philosophy helps to guide future reforms.

When I began teaching, my belief was that teaching was a matter of knowledge and skills transference: it was my job as an educator first and foremost to transfer information into my students’ minds. Now, it is my belief that my primary task is to help my students become *autonomous*: able to analyze complex data, turn their thoughts into practical action, and continue learning long after they have left my classroom. If all I accomplish is helping my students memorize facts, procedures, and formulae about instrumentation, I have utterly failed them. My real job is to challenge them to become autonomous, critical thinkers and doers, so they will be able to fully take responsibility for their own lives and their own careers.

This shift in philosophy happened as a result of contact with many employers of my students, who told me the most important thing any student could learn in school was *how to learn*. In life, learning is not an option but a necessity, especially in highly complex fields such as instrumentation and control. Any instrument technician or engineer who stagnates in their learning is destined for obsolescence. Conversely, those with the ability and drive to continually learn new things will find opportunities opening for them all throughout their careers.

A former classmate of mine I studied instrumentation with told me of his path to success in this field. While never the smartest person in class (he struggled mightily with math concepts), he was always very determined and goal-oriented. His first job placed him in the field of automotive research and development, where he was responsible for “instrumenting” heavy trucks with sensors to perform both destructive and non-destructive tests on them. The instrumentation he used at this job was often quite different from the industrial instruments he learned in school, and so he found himself having to constantly refer to textbooks and equipment manuals to learn how the technology worked. This was true even when he was “on the road” doing field-service work. He told me of many evenings spent in some tavern or pub, a beer in one hand and an equipment manual in the other, learning how the equipment was supposed to work so he could fix the customer’s problem the next day.

This hard work and self-directed learning paid off handsomely for my friend, who went on to set up an entire testing department for a major motorcycle manufacturer, and then started his own vehicle testing company (specializing in power-sport vehicles) after that. None of this would have been possible for my friend had he relied exclusively on others to teach him what he needed to know, taking the passive approach to learning so many students do. The lesson is very general and very important: *continual learning is a necessary key to success*.

One of the corollaries to my philosophy of education is that individual learning styles are ultimately not to be accommodated. This may come as a shock to many educators, who have learned about the various styles of learning (auditory, kinesthetic, visual, etc.) and how the acquisition of new information may be improved by teaching students according to their favored modes. Please understand that I am not denying the fact different people prefer learning in different ways. What I am saying is that we fail to educate our students (i.e. empower them with *new* abilities) if all we ever do is teach to their preferences, if we never challenge them to do what is novel or uncomfortable.

The well-educated person can learn by listening, learn by watching others, *and* learn by direct hands-on experimentation. A truly educated person may still retain a preference for one of these modes over the others, but that preference does not *constrain* that person to learning in only one way. This is our goal as educators: nothing short of expanding each student's modes of learning.

If a student experiences difficulty learning in a particular way, the instructor needs to engage with that student in whatever mode makes the most sense for them *with the goal of strengthening the areas where that student is weak*. For example, a student who is weak in reading (visual/verbal) but learns easily in a hands-on (kinesthetic) environment should be shown how to relate what they perceive kinesthetically to the words they read in a book. Spending time with such students examining an instrument to learn how it functions, then reading the service manual or datasheet for that instrument to look for places where it validates the same principles, is one example of how an instructor might help a student build connections between their strong and weak modes of learning. Investigating subjects through multiple approaches such as this also shows students the value of each learning mode: a student might find they easily grasp “how” an instrument works by directly observing and experimenting with it, but that they more readily grasp “why” it was built that way by reading the manufacturer's “theory of operation” literature.

Keeping the goal of life-long learning in mind, we must ask ourselves the question of how our students will need to learn new things once they are no longer under our tutelage. The obvious answer to this question is that they will need to be able to learn in any mode available to them, if they are to flourish. Life is indifferent to our needs: reality does not adapt itself to favor our strengths or to avoid challenging our weaknesses. Education must therefore focus on the well-rounded development of learning ability.

By far the greatest amount of resistance I encounter from students in terms of learning styles is learning by reading. It is rare to find a student who reads well, for example, but struggles at learning in a hands-on environment (kinesthetic) or struggles to understand spoken information (auditory). The reason for this, I believe, is that reading is a wholly unnatural skill. Entire cultures exist without a written language, but there is not a culture in the world that lacks a spoken one. Interpreting the written word, to the level of proficiency required for technical learning, is a skill born of much practice.

Unfortunately, the popular application of learning styles in modern education provides students with a ready-made, officially-sanctioned excuse for not only their inability (“That’s just not how my brain works”), but also for continued tolerate of that inability (“I shouldn’t have to learn in a way I’m not good at”). The challenge for the instructor is helping students develop their ability to learn in non-favored modes despite this psychological resistance.


My general advice for educators is to never compromise the “big picture” philosophy of empowering your students’ thinking. Some key points I always try to keep in mind are:

- **Lead by example:** Regularly showcase for your students your own excitement for the subject and your own continual learning adventures. Likewise, you need to model the same learning modes you ask them to develop: let them see you learn new things, demonstrating how multiple modes are necessary to be an effective self-educator.
- **Teach by asking questions:** Socrates had the right idea – if you want to make people examine their assumptions and discover misconceptions, ask lots of challenging questions. This is how I have conditioned myself to respond to student questions: I generally answer with a question of my own seeking the heart of the student’s confusion. Posing “thought experiments” for students to conduct is another form of questioning that not only clarifies concepts, but also builds good critical-thinking habits. Anyone familiar with Socrates’ fate knows, however, that people tend to react defensively when their assumptions are challenged by persistent questions. A helpful hint for avoiding this kind of reaction is to give the student adequate time and personal space to contemplate your questions. If the student ever feels uncomfortable either with your observation of their efforts or the rapidity of your questioning, they will “lock up” and refuse to engage. Sometimes the best way to manage this behavior is to pose a question to the student, then tell them you will get back to them after a few minutes, rather than to watch them struggle answering your question(s).
- **Be willing to provide the help they need:** If students struggle at certain tasks or with thinking in certain ways, devote extra time with them to practice these skills. Let them know in very practical ways how you are willing to work just as hard as you are asking them to work. Note that this does not mean giving in to demands for lecture. That would be giving students what they *want*, rather than what they *need*. Instead, it means focusing directly on whatever weaknesses are hindering their growth as learners, and aggressively working to strengthen those weaknesses. If it means reserving time to read with students who say they can’t understand the text, then that is your job as their teacher.
- **Nothing builds confidence and dissolves apprehension like success:** Remind your students of the challenges they have already overcome, and the progress they have already made.
- **Be patient:** That same student who complains now about having to read, to think independently, and tackle challenging problems will come back to thank you years later. Just as you expected them to think long-term while they were in your class, so you need to think long-term with regard to their appreciation for your standards and efforts. Transformative education is a marathon, not a sprint!

D.2 Teaching technical practices (labwork)

Labwork is an essential part of any science-based curriculum. Here, much improvement may be made over the “standard” educational model to improve student learning. In my students’ Instrumentation courses, I forbid the use of pre-built “trainer” systems and lab exercises characterized by step-by-step instructions. Instead, I have my students construct real working instrumentation systems. The heart of this approach is a “multiple-loop” system spanning as large a geographic area as practically possible, with instruments of all kinds connecting to a centralized control room area. None of the instruments need perform any practical purpose, since the goal of the multiple-loop system is for students to learn about the instruments themselves.

A model for a multiple-loop system might look something like this:



Instruments may or may not be grouped together to form complete control systems, since process control is not necessarily the purpose of this system. The primary purpose of a multiple-loop instrument system is to provide an infrastructure for students to investigate instrumentation apart from the dynamics of a functioning process. The separation of controls from process may seem counter-productive at first, but it actually provides a rich and flexible learning experience. Students are able to measure instrument signals and correlate them with actual physical measurements, take instruments in and out of service, check instrument calibration, see the effects of calibration on measurement accuracy and resolution, practice lock-out and tag-out procedures, diagnose instrument problems introduced by the instructor, practice installing and removing instruments, remove old wire and pull new wire into place, practice sketching and editing loop diagrams, and many other practical tasks without having to balance the needs of a working process. The system may be altered at any time as needed, since there are no process operating constraints to restrict maintenance operations.

The fundamental advantage of a process-less instrument system is there are no process limitations restricting educational objectives. In this sense it is as flexible as a computer simulation, but with the advantage of using real-world components.

The first academic year I attempted to build such a system with my students was 2002-2003. Our system cost almost nothing⁴, with a control panel fabricated from a discarded fiberglass electrical enclosure and 4-20 mA loop wiring salvaged from discarded spools of category-5 data communications cable (four twisted pairs per cable). We stapled the cable runs to the lab room wall, and used cheap terminal block assemblies to provide connection points between the cat-5 trunk cables and individual instrument cables. Our first loops built with this system included the following:

- Air compressor receiver tank pressure measurement – *measurement only*
- Air compressor temperature measurement – *measurement only*
- Regulated (service) air pressure measurement – *measurement only*
- Wash basin water level measurement – *measurement only*
- Water column level and temperature control – *measurement and control*
- Air reservoir pressure control – *measurement and control*

The first four of these instrument loops were “permanent” in that they were never disconnected once installed. The water level and temperature control system was a later addition made toward the end of the academic year. It began as a pneumatic system, then was upgraded to electronic (single-loop digital controller), then as a PLC-controlled process, then finally as a DCS-controlled process. The air pressure control system was much the same. All the time we left the process vessels and field instruments in place, used the same signal tubing and wiring, but merely changed the control instruments at the other end of that tubing and wiring.

In addition to these six permanent and semi-permanent loops, students used the system throughout the year to connect individual instruments for loop calibration. Usually there was no control involved, as they were simply studying individual instruments and were not ready for a complete control system yet. Every time they had a transducer to calibrate, a control valve to test, or a transmitter to configure, I required them to tie it into the loop system and document the loop using ISA standard loop diagrams. Then, I would fault their loops (usually electrically by creating opens or shorts in signal wiring, or pneumatically by creating leaks or by plugging tubes with foam earplugs) and have them troubleshoot the loops using real test equipment, documenting their diagnostic steps for grading purposes. After successful commissioning, calibration, and troubleshooting, students disassembled the loop so the instruments could be used again in a different loop.

Our multiple-loop instrument system – despite its crude appearance and low cost – was extremely successful as an educational tool. My students gained a tremendous amount of practical knowledge and skill in addition to the basic theory. Abstract principles of measurement and instrument

⁴Of course, we had to have plenty of instruments to install in this loop system, and industrial instruments are not cheap. My point is that the *infrastructure* of control panel, trunk cabling, field wiring, terminal blocks, etc. was very low-cost. If an Instrumentation program already has an array of field instruments for students to work with in a lab setting, it will not cost much at all to integrate these instruments into a realistic multi-loop system as opposed to having students work with individual instruments on the benchtop or installed in dedicated “trainer” modules.

application “came alive” for them as they saw the pieces fit together to make a working system. The intentionally distributed nature of the system – with the control panel located in one far corner of the room and field instruments scattered around the rest of the room – forced students to think and work in a manner much more similar to the real work environment. There were days they were so excited about working on this system that I had to coax them out of class when the school day was over!

In the summer of 2006 I upgraded the loop system to include a 12 foot by 8 foot metal control room panel (donated by a local paper mill), a set of computer workstations for DCS and SCADA system consoles, industry-standard terminal block assemblies located in electrical enclosures, with plenty of electrical conduit runs between different locations in the lab facility to allow pulling of new wires and cables. Students still must connect each instrument they learn about into the system, configuring either a panel-mounted or computer-based display to register the measured variable in proper units (or to receive a control signal if the instrument in question is a final control element). Construction of working control systems (transmitter, controller, valve or motor) is quite easy with this infrastructure in place. The geographically distributed nature of the system lends itself well to realistic troubleshooting, with students working in teams (communicating via hand-held radios) to diagnose problems intentionally placed into the system.

A new feature of the 2006 multi-loop system is that it included digital communication as well as analog (4-20 mA) signaling. Multiple Ethernet hubs were installed throughout the lab, interconnected to form a single 10 Mbps network linking personal computers with loop controllers and PLCs. Non-dedicated category 5 cabling was also used for RS-232 and RS-485 communication between serial devices (e.g. data acquisition modules) as needed. FOUNDATION Fieldbus wiring was also installed (twin-lead shielded cable with 100 Ω characteristic impedance) allowing the interconnection of fully digital field instruments such as transmitters and digital valve positioners.

The following photographs show the appearance of the new (2006) multiple loop system, beginning with the control panel and computer workstation cluster. These two elements comprise the “control room area” of the lab:



In another area of the lab room is a pneumatic control panel and a cabinet housing the distributed control system (DCS) I/O rack:



The rest of the lab room is dedicated as a “field area” where field instruments are mounted and wires (or tubes) run to connect those instruments to remote indication and/or control devices:



Note the use of metal strut hardware to form a frame which instruments may be mounted to, and the use of flexible liquid-tight conduit to connect field instruments to rigid conduit pieces so loop wiring is never exposed.

A less expensive alternative⁵ to metal strut is standard *industrial pallet racking*, examples shown here with 2 inch pipe attached for instrument mounting, and enclosures attached for instrument cable routing and termination:




The multiple-loop system is designed to be assembled, disassembled, and reassembled repeatedly as each student team works on a new instrument. As such, it is in a constant state of flux. It is not really a *system* so much as it is an *infrastructure* for students to build working loops and control systems within.

⁵When I built my first fully-fledged educational loop system in 2006 at Bellingham Technical College in Washington state (I built a crude prototype in 2003), I opted for Cooper B-Line metal strut because it seemed the natural choice for the application. It wasn't until 2009 when I needed to expand and upgrade the loop system to accommodate more students that I happened to come up with the idea of using pallet racking as the framework material. Used pallet racking is plentiful, and very inexpensive compared to building a comparable structure out of metal strut. As these photographs show, I still used Cooper B-Line strut for some portions, but the bulk of the framework is simply pallet racking adapted for this unconventional application.

In addition to the multiple-loop system, my students' lab contains working processes (also student-built!) which we improve upon every year. One such process is a water flow/level/temperature control system, shown here:



Another is a turbocompressor system, built around a diesel engine turbocharger (propelled by the discharge of a 2 horsepower air blower) and equipped with a pressurized oil lubrication system and temperature/vibration monitor:



Yet another permanent process is this electrical power monitoring unit, where protective (overcurrent) relay operation may be demonstrated:



Measurements of voltage and current in this particular system may be integrated into the rest of the multi-loop system by using voltage and current transducers with 4-20 mA output signals. Digital protective relays may be connected to the multi-loop system using serial data communication (RS-232, RS-485) signals.

The process piping and equipment on these permanent systems are altered only when necessary, but the control systems on these processes may undergo major revisions each year when a new group of students takes the coursework relevant to those systems. Having a set of functioning process systems present in the lab at all times also gives students examples of working instrument systems to study as they plan construction of their temporary loops in the multiple-loop system.

D.3 Teaching diagnostic principles and practices

Diagnostic ability is arguably the most difficult skill to develop within a student, and also the most valuable skill a working technician can possess⁶. In this section I will outline several principles and practices teachers may implement in their curricula to teach the science and art of troubleshooting to their students.

First, we need to define what “troubleshooting” is and what it is not. It is *not* the ability to follow printed troubleshooting instructions found in equipment user’s manuals⁷. It is *not* the ability to follow one rigid sequence of steps ostensibly applicable to any equipment or system problem⁸. Troubleshooting is first and foremost the practical application of *scientific thinking* to repair of malfunctioning systems. The principles of hypothesis formation, experimental testing, data collection, and re-formulation of hypotheses is the foundation of any detailed cause-and-effect analysis, whether it be applied by scientists performing primary research, by doctors diagnosing their patients’ illnesses, or by technicians isolating problems in complex electro-mechanical-chemical system. In order for anyone to attain mastery in troubleshooting skill, they need to possess the following traits:

- A rock-solid understanding of relevant, fundamental principles (e.g. how electric circuits work, how feedback control loops work)
- Close attention to detail
- An open mind, willing to pursue actions led by data and not by preconceived notions

The first of these points is addressed by any suitably rigorous curriculum. The other points are habits of thought, best honed by months of practice. Developing diagnostic skill requires much time and practice, and so the educator must plan for this in curriculum design. It is not enough to sprinkle a few troubleshooting activities throughout a curriculum, or (worse yet!) to devote an isolated course to the topic. Troubleshooting should be a topic tested on every exam, present in every lab activity, and (ideally) touched upon in every day of the student’s technical education.

Scientific, diagnostic thinking is characterized by a repeating cycle of *inductive* and *deductive* reasoning. Inductive reasoning is the ability to reach a general conclusion by observing specific details. Deductive reasoning is the ability to predict details from general principles. For example, a student engages in deductive reasoning when they conclude an “open” fault in a series DC circuit will cause current in that circuit to stop. That same student would be thinking inductively if they measured zero current in a DC series circuit and thus concluded there was an “open” fault somewhere in it. Of these two cognitive modes, inductive is by far the more difficult because multiple solutions exist for any one set of data. In our zero-current series circuit example, inductive reasoning might lead the troubleshooter to conclude an open fault existed in the circuit. However, an unpowered source could also be at fault, or for that matter a malfunctioning ammeter falsely registering zero

⁶One of the reasons diagnostic skill is so highly prized in industry is because so few people are actually good at it. This is a classic case of supply and demand establishing the value of a commodity. Demand for technicians who know how to troubleshoot will always be high, because technology will always break. Supply, however, is short because the skill is difficult to teach. This combination elevates the value of diagnostic skill to a very high level.

⁷Yes, I have actually heard people make this claim!

⁸The infamous “divide and conquer” strategy of troubleshooting where the technician works to divide the system into halves, isolating which half the problem is in, is but *one particular procedure: merely one tool in the diagnostician’s toolbox*, and does not constitute the whole of diagnostic method.

current when in fact there is current. Inductive conclusions are *risky* because the leap from specific details to general conclusions always harbor the potential for error. Deductive conclusions are *safe* because they are as secure as the general principles they are built on (e.g. *if* an “open” exists in a series DC circuit, there will be *no* current in the circuit, guaranteed). This is why inductive conclusions are always validated by further deductive tests, not vice-versa. For example, if the student induced that an unpowered voltage source might cause the DC series circuit to exhibit zero current, they might elect to test that hypothesis by measuring voltage directly across the power supply terminals. If voltage is present, then the hypothesis of a dead power source is incorrect. If no voltage is present, the hypothesis is provisionally true⁹.

Scientific method is a cyclical application of inductive and deductive reasoning. First, an hypothesis is made from an observation of data (inductive). Next, this hypothesis is checked for validity – an experimental test to see whether or not a prediction founded on that hypothesis is correct (deductive). If the data gathered from the experimental test disproves the hypothesis, the scientist revises the hypothesis to fit the new data (inductive) and the cycle repeats.

Since diagnostic thinking requires both deductive and inductive reasoning, and deductive is the easier of the two modes to engage in, it makes sense for teachers to focus on building deductive skill first. This is relatively easy to do, simply by adding on to the theory and practical exercises students already engage in during their studies.

Both deductive and inductive diagnostic exercises lend themselves very well to Socratic discussions in the classroom, where the instructor poses questions to the students and the students in turn suggest answers to those questions. The next two subsections demonstrate specific examples showing how deductive and inductive reasoning may be exercised and assessed, both in a classroom environment and in a laboratory environment.

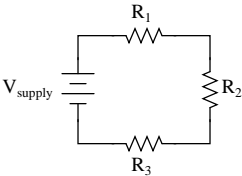
D.3.1 Deductive diagnostic exercises

Deductive reasoning is where a person applies general principles to a specific situation, resulting in conclusions that are logically necessary. In the context of instrumentation and control systems, this means having students predict the consequence(s) of specified faults in systems. The purpose of building this skill is so that students will be able to quickly and accurately test “fault hypotheses” in their minds as they analyze a faulted system. If they suppose, for example, that a cable has a break in it, they must be able to deduce what effects a broken cable will have on the system in order to formulate a good test for proving or disproving that hypothesis.

⁹Other things could be at fault. An “open” test lead on the multimeter for example could account for both the zero-current measurement and the zero-voltage measurement. This scientific concept eludes many people: it is far easier to *disprove* an hypothesis than it is to *prove* one. To quote Albert Einstein, “No amount of experimentation can ever prove me right; a single experiment can prove me wrong.”

Example: predicting consequence of a single fault

For example, consider a simple three-resistor series DC circuit, the kind of lab exercise one would naturally expect to see within the first month of education in an Instrumentation program. A typical lab exercise would call for students to construct a three-resistor series DC circuit on a solderless breadboard, predict voltage and current values in the circuit, and validate those predictions using a multimeter. A sample exercise is shown here:


Competency: Series DC resistor circuit		Version:	
Schematic			
			
Given conditions			
$V_{\text{supply}} =$	$R_1 =$	$R_2 =$	$R_3 =$
Parameters			
	Predicted	Measured	
I_{supply}	<input type="text"/>	<input type="text"/>	I_{R1}
V_{R1}	<input type="text"/>	<input type="text"/>	I_{R2}
V_{R2}	<input type="text"/>	<input type="text"/>	I_{R3}
V_{R3}	<input type="text"/>	<input type="text"/>	<input type="text"/>
Analysis			
Relationship between resistor voltage drops and total voltage:			
Fault analysis			
Suppose component <input type="text"/> fails <input type="checkbox"/> open <input type="checkbox"/> other _____			
<input type="checkbox"/> shorted			
What will happen in the circuit?			

Note the **Fault Analysis** section at the end of this page. Here, after the instructor has verified the correctness of the student's mathematical predictions and multimeter measurements, he or she would then challenge the student to predict the effects of a random component fault (either quantitatively or qualitatively), perhaps one of the resistors failing open or shorted. The student makes their predictions, then the instructor simulates that fault in the circuit (either by pulling the resistor out of the solderless breadboard to simulate an "open" or placing a jumper wire in parallel with the resistor to simulate a "short"). The student then uses his or her multimeter to verify the predictions. If the predicted results do not agree with the real measurements, the instructor

works with the student to identify why their prediction(s) were faulty and hopefully correct any misconceptions leading to the incorrect result(s). Finally, a different component fault is chosen by the instructor, predictions made by the student, and verification made using a multimeter. The actual amount of time added to the instructor’s validation of student lab completion is relatively minor, but the benefits of exercising deductive diagnostic processes are great.

Example: predicting consequences of multiple faults

An example of a more advanced deductive diagnostic exercise appropriate to later phases of a student’s Instrumentation education appears here. A loop diagram shows a pressure recording system for an iso-butane distillation column:



A set of questions accompanying this diagram challenge each student to predict effects in the instrument system resulting from known faults, such as:


- PT-325 block valve left shut and bleed valve left open (*predict voltage between TB27-16 and TB27-17*)
- Loose wire connection at TB64-9 (*predict pressure indication at PIR-325*)
- Circuit breaker #5 shut off (*predict loop current at applied pressure of 50 PSI*)

Given each hypothetical fault, there is only one correct conclusion for any given question. This makes deductive exercises unambiguous to assess.

Example: identifying possible faults

A more challenging type of deductive troubleshooting problem easily given in homework or on exams appears here. It asks students to examine a list of potential faults, marking each one of them as either “possible” or “impossible” based on whether or not each fault is independently capable of accounting for all symptoms in the system:

Suppose a voltmeter registers 6 volts between test points **C** and **B** in this series-parallel circuit:




Fault	Possible	Impossible
R_1 failed open		
R_2 failed open		
R_3 failed open		
R_1 failed shorted		
R_2 failed shorted		
R_3 failed shorted		
Voltage source dead		

This is still a *deductive* thinking exercise because each of the faults is given to the student, and it is a matter of deduction to determine whether or not each one of these proposed faults is capable of accounting for the symptoms. Students need only apply the general rules of electric circuits to tell whether or not each of these faults would cause the reported circuit behavior.

True to form for any deductive problem, there can only be one correct answer for each proposed fault. This makes the exercise easy and unambiguous to grade, while honing vitally important diagnostic skills.

One of the benefits of this kind of fault analysis problem is that it requires students to consider *all* consequences of a proposed fault. In order for one of the faults to be considered “possible,” it must account for all symptoms, not just one symptom. An example of this sort of problem is seen here:

Suppose the voltmeter in this circuit registers a strong *negative* voltage. A test using a digital multimeter (DMM) shows the voltage between test points **D** and **B** to be 6 volts:




Fault	Possible	Impossible
R_1 failed open		
R_2 failed open		
R_3 failed open		
R_4 failed open		
R_1 failed shorted		
R_2 failed shorted		
R_3 failed shorted		
R_4 failed shorted		
Voltage source dead		

Several different faults are capable of causing the meter to read strongly negative (R_1 short, R_2 open, R_3 open, R_4 short), but only two are capable of this while not affecting the normal voltage (6 volts) between test points D and B: R_3 open or R_4 short. This simple habit of checking to see that the proposed fault accounts for *all* apparent conditions and not just some of them is essential for effective troubleshooting.

This same question format may be easily applied to most any system, not just electrical circuits. Consider this example, determining possible versus impossible faults on an exhaust scrubber system:

After years of successful operation, the level control loop in this exhaust scrubbing system begins to exhibit problems. The liquid level inside the scrubbing tower mysteriously drops far below setpoint, as indicated by the level gauge (LG) on the side of the scrubber. The operators have tried to rectify this problem by increasing the setpoint adjustment on the level controller (LC), to no avail. The level transmitter (LT) is calibrated 3 PSI at 0% (low) level and 15 PSI at 100% (high) level:



Fault	Possible	Impossible
Air supply to LT shut off		
Air supply to LC shut off		
Pump shut off		
Broken air line between LT and LC		
Broken air line between LC and LV		
Plugged nozzle inside LC		
Plugged orifice inside LC		
Leak in bottom of scrubber		


This exercise is particularly good because it requires the student to determine the action of the level controller (LC) before some of the proposed faults may be analyzed. In this case, the level controller must be *direct-acting*, so that an increasing liquid level inside the scrubber will cause an increasing air signal to the air-to-open (ATO) valve, letting more liquid out of the scrubber to

stabilize the level. Without knowing that the level controller is direct-acting, it would be impossible to conclude the effect of a failed air supply to the level transmitter (LT), the first fault proposed in the table.

Example: assessing value of multiple diagnostic tests

A variation on this theme of determining the possibility of proposed faults is to assess the usefulness of proposed diagnostic tests. In other words, the student is presented with a scenario where something is amiss with a system, but instead of selecting a set of proposed faults as being either possible or impossible, the student must determine whether or not a set of proposed *tests* would be diagnostically relevant. An example of this in a simple series-parallel resistor circuit is shown here:

Suppose a voltmeter registers 0 volts between test points **E** and **F** in this circuit. Determine the diagnostic value of each of the following tests. Assume only one fault in the system, including any single component or any single wire/cable/tube connecting components together. If a proposed test could provide new information to help you identify the location and/or nature of the one fault, mark “yes.” Otherwise, if a proposed test would not reveal anything relevant to identifying the fault (already discernible from the measurements and symptoms given so far), mark “no.”



Diagnostic test	Yes	No
Measure V_{AC} with power applied		
Measure V_{JK} with power applied		
Measure V_{CK} with power applied		
Measure I_{R1} with power applied		
Measure I_{R2} with power applied		
Measure I_{R3} with power applied		
Measure R_{AC} with source disconnected from R_1		
Measure R_{DF} with source disconnected from R_1		
Measure R_{EG} with source disconnected from R_1		
Measure R_{HK} with source disconnected from R_1		

This form of diagnostic problem tends to be much more difficult to solve than simply determining the possibility of proposed faults. To solve this form of problem, the student must first determine all possible component faults, and then assess whether or not each proposed test would provide *new information* useful in identifying which of these possible faults is the actual fault.

In this example problem, there are really only a few possible faults: a dead source, an open resistor R_1 , a shorted resistor R_2 , a shorted resistor R_3 , or a broken wire (open connection) somewhere in the loop E-B-A-C-D-F.

The first proposed test – measuring voltage between points A and B – would be useful because it would provide different results given a dead source, open R_1 , shorted R_2 , or shorted R_3 versus an open between A-E or between C-F. Any of the former faults would result in 0 volts between A and B, while any of the latter faults would result in full source voltage between A and B.

The next proposed test – measuring voltage between points J and K – would be useless because we already know what the result will be: 0 volts. This result of this proposed test will be the same no matter which of the possible faults causing 0 voltage between points E and F exists, which means it will shed no new light on the nature or location of the fault.

Despite being very challenging, this type of deductive diagnostic exercise is nevertheless easy to administer and unambiguous to grade, making it very suitable for written tests.


D.3.2 Inductive diagnostic exercises

Inductive reasoning is where a person derives general principles from a specific situation. In the context of instrumentation and control systems, this means having students propose faults to account for specific symptoms and data measured in systems. This is actual troubleshooting, as opposed to deductive diagnosis which is an enabling skill for effective troubleshooting.

While real hands-on exercises are best for developing inductive diagnostic skill, much learning and assessment may be performed in written form as well.

Example: proposing faults in loop diagram

This exam question is a sample of an inductive diagnosis exercise presented in written form:



This system used to work just fine, but now it has a problem: the controller registers zero flow, and its output signal (to the valve) is saturated at 100% (wide open) as though it were trying to “ask” the valve for more flow. Your first diagnostic step is to check to see if there actually is gasoline flow through the flowmeter and valve by looking at the rotameter. The rotameter registers a flow rate in excess of 2 gallons per hour.

Identify possible faults in this system that could account for the controller’s condition (no flow registered, saturated 100% output), depending on what you find when you look at the rotameter:

- Possible fault:
- Possible fault:


Here, the student must identify two probably faults to account for all exhibited symptoms. More than two different kinds of faults are possible¹⁰, but the student need only identify two faults independently capable of causing the controller to register zero flow when it should be registering more than 2 GPH.

¹⁰ Jammed turbine wheel in flowmeter, failed pickup coil in flowmeter, open wire in cable FT-112 or pair 1 of cable 3 (assuming the flow controller’s display was not configured to register below 0% in an open-loop condition), etc.

Example: “virtual troubleshooting”

An excellent supplement to any hands-on troubleshooting activities is to have students perform “virtual troubleshooting” with you, the instructor. This type of activity cannot be practiced alone, but requires the participation of someone who knows the answer. It may be done with individual students or with a group.

A “virtual troubleshooting” exercise begins with a schematic diagram of the system such as this, containing clearly labeled test points (and/or terminal blocks) for specifying the locations of diagnostic tests:



Each student has their own copy of the diagram, as does the instructor. The instructor has furthermore identified a realistic fault within this system, and has full knowledge of that fault’s effects. In other words, the instructor is able to immediately tell a student how much voltage will be read between any two test points, what the effect of jumpering a pair of test points will be, what will happen when a pushbutton is pressed, etc.

The activity begins with a brief synopsis of the system’s malfunction, narrated by the instructor. Students then propose diagnostic tests to the instructor, with the instructor responding back to each student the results of their tests. As students gather data on the problem, they should be able to narrow their search to find the fault, choosing appropriate tests to identify the precise nature and location of the fault. The instructor may then assess each student’s diagnostic performance based on the number of tests and their sequence.


When performed in a classroom with a large group of students, this is actually a lot of fun!

Example: realistic faults in solderless breadboards

Solderless breadboards are universally used in the teaching of basic electronics, because they allow students to quickly and efficiently build different circuits using replaceable components. As wonderful as breadboards are for fast construction of electronic circuits, however, it is virtually impossible to create a realistic component fault without the fault being evident to the student simply by visual inspection. In order for a breadboard to provide a realistic *diagnostic* scenario, you must find a way to hide the circuit while still allowing access to certain test points in the circuit.

A simple way to accomplish this is to build a “troubleshooting harness” consisting of a multi-terminal block connected to a multi-conductor cable. Students are given instructions to connect various wires of this cable to critical points in the circuit, then cover up the breadboard with a five-sided box so that the circuit can no longer be seen. Test voltages are measured between terminals on the block, not by touching test leads to component leads on the breadboard (since the breadboard is now inaccessible).

The following illustration shows what this looks like when applied to a single-transistor amplifier circuit:



If students cannot visually detect a fault, they must rely on voltage measurements taken from terminals on the block. This is quite challenging, as not even the shapes of the components may be seen with the box in place. The only guide students have for relating terminal block test points to points in the circuit is the schematic diagram, which is good practice because it forces students to interpret and follow the schematic diagram.

Example: realistic faults in a multi-loop instrument system

Whole instrumentation systems may also serve to build and assess individual diagnostic competence. In my lab courses, students work in teams to build functioning measurement and control loops using the infrastructure of a multiple-loop system (see Appendix section D.2 beginning on page 3186 for a detailed description). Teamwork helps expedite the task of constructing each loop, such that even an inexperienced team is able to assemble a working loop (transmitter connected to an indicator or controller, with wires pulled through conduits and neatly landed on terminal blocks) in just a few hours.

Each student creates their own loop diagram showing all instruments, wires, and connection points, following ISA standards. These loop diagrams are verified by doing a “walk-through” of the loop with all student team members present. The “walk-through” allows the instructor to inspect work quality and ensure any necessary corrections are made to the diagrams. After each team’s loop has been inspected and all student loop diagrams edited, the diagrams are placed in a document folder accessible to all students in the lab area.

Once the loop is wired, calibrated, inspected, and documented, it is ready to be faulted. When a student is ready to begin their diagnostic exercise, they gather their team members and approach the instructor. The instructor selects a loop diagram from the document folder *not* drawn by that student, ideally of a loop constructed by another team. The student and teammates leave the lab room, giving the instructor time to fault the loop. Possible faults include:

- Loosen wire connections
- Short wire connections (loose strands of copper strategically placed to short adjacent terminals together)
- Cut cables in hard-to-see locations
- Connect wires to the wrong terminals
- Connect wire pairs backward
- Mis-configure instrument calibration ranges
- Insert square root extraction where it is not appropriate
- Mis-configure controller action or display
- Insert unrealistically large damping constants in either the transmitter, indicator, or final element
- Plug pneumatic signal lines with foam earplugs
- Turn off hand valves
- Trip circuit breakers

After the fault has been inserted, the instructor calls the student team back into the lab area (ideally using a hand-held radio, simulating the work environment of a large industrial facility where technicians carry two-way radios) to describe the symptoms. This part of the exercise works best

when the instructor acts the part of a bewildered operator, describing what the system is not doing correctly, without giving any practical advice on the location of the problem or how to fix it¹¹. An important detail for the instructor to include is the “history” of the fault: is this a new loop which has never worked correctly, or was it a working system that failed? Faults such as mis-connected wires are realistic of improper installation (new loop), while faults such as loose connections are perfectly appropriate for previously working systems. Whether the instructor freely offers this “history” or waits for the student to ask, it is important to include in the diagnostic scenario because it is an extremely useful piece of information to know while troubleshooting actual systems in industry. Virtually anything may be wrong (including multiple faults) in a brand-new installation, whereas previously working systems tend to fail in fewer ways.

After this introduction, the one student begins his or her diagnosis, with the other team members acting as scribes to document the student’s steps. The diagnosing student may ask a teammate for manual assistance (e.g. operating a controller while the student observes a control valve’s motion), but no one is allowed to help the one student diagnose the problem. The instructor observes the student’s procedure while the student explains the rationale motivating each action, with only a short time given (typically 5 minutes) to determine the general location of the fault causing the problem (e.g. located in the transmitter, control valve, wiring, controller/DCS, tubing, etc.). If after that time period the student is unable to correctly isolate the general location, the exercise is aborted and the instructor reviews the student’s actions (as documented by the teammates) to help the student understand where they went wrong in their diagnosis. Otherwise, the student is given more time¹² to pinpoint the nature of the fault.

Depending on the sequencing of your students’ coursework, some diagnostic exercises may include components unfamiliar to the student. For example, a relatively new student familiar only with the overall function of a control loop but intimately familiar with the workings of measurement devices may be asked to troubleshoot a loop where the fault is in the control valve positioner rather than in the transmitter. I still consider this to be a fair assessment of the student’s diagnostic ability, so long as the expectations are commensurate with the student’s knowledge. I would not expect a student to precisely locate the nature of a positioner fault if they had never studied the function or configuration of a valve positioner, but I would expect them to be able to broadly identify the location of the fault (e.g. “it’s somewhere in the valve”) so long as they knew how a control signal is supposed to command a control valve to move. That student should be able to determine by manually adjusting the controller output and measuring the output signal with the appropriate loop-testing tools that the valve was not responding as it should despite the controller properly performing its function. The ability to diagnose problems in instrument systems where some components of the system are mysterious “black boxes” is a very important skill, because your students *will* have to do exactly that when they step into industry and work with specific pieces of equipment they never had time to learn about in school¹³.

¹¹I must confess to having a lot of fun here. Sometimes I even try to describe the problem incorrectly. For instance, if the problem is a huge damping constant, I might tell the student that the instrument simply does not respond, because that is what it looks like if you do not take the time to watch it respond *very slowly*.

¹²The instructor may opt to step away from the group at this time and allow the student to proceed unsupervised for some time before returning to observe.

¹³I distinctly remember a time during my first assignment as an industrial instrument technician that I had to troubleshoot a problem in a loop where the transmitter was an oxygen analyzer. I had no idea how this particular analyzer functioned, but I realized from the loop documentation that it measured oxygen concentration and output a signal corresponding to the percentage concentration (0 to 21 percent) of O₂. By subjecting the analyzer to known concentrations of oxygen (ambient air for 21%, inert gas for 0%) I was able to determine the analyzer was responding

I find it nearly impossible to fairly assign a letter or percentage grade to any particular troubleshooting effort, because no two scenarios are quite the same. Mastery assessment (either pass or fail, with multiple opportunities to re-try) seems a better fit. Mastery assessment with no-penalty retries also enjoys the distinct advantage of directing needed attention toward and providing more practice for weaker students: the more a student struggles with troubleshooting, the more they must exercise those skills.

Successfully passing a troubleshooting exercise requires not only that the fault be correctly identified and located in a timely manner, but that all steps leading to the diagnosis are logically justified. Random “trial and error” tests by the student will result in a failed attempt, even if the student was eventually able to locate the fault. A diagnosis with no active tests such as multimeter or test gauge measurements, or actions designed to stimulate system components, will also fail to pass. For example, a student who successfully locates a bad wiring connection by randomly tugging at every wire connection should *not* pass the troubleshooting exercise because such actions do not demonstrate diagnostic thinking¹⁴.

To summarize key points of diagnostic exercises using a multiple-loop system:

- Students work in teams to build each loop
- Loop inspection and documentation finalized by a “walk-through” with the instructor
- Instructor placement of faults (it is important no student knows what is wrong with the loop!)
- Each student individually diagnoses a loop, with team members acting merely as scribes
- Students must use loop diagrams drawn by someone else, ideally diagnosing a loop built by a different team
- Brief time limit for each student to narrow the scope of the problem to a general location in the system
- Passing a diagnostic exercise requires:
 - Accurate identification of the problem
 - Each diagnostic step logically justified by previous results
 - Tests (measurements, component response checks) performed before reaching conclusions
- Mastery (pass/fail) assessment of each attempt, with multiple opportunities for re-tries if necessary

quite well, and that the problem was somewhere else in the system. If the analyzer had failed my simple calibration test, I would have known there was something wrong with it, which would have led me to either get help from other technicians working at that facility or simply replace the analyzer with a new unit and try to learn about and repair the old unit in the shop. In other words, my ignorance of the transmitter’s specific workings did not prevent me from diagnosing the loop in general.

¹⁴Anyone can (eventually) find a fault if they check every detail of the system. Randomly probing wire connections or aimlessly searching through a digital instrument’s configuration is not troubleshooting. I have seen technicians waste incredible amounts of time on the job randomly searching for faults, when they could have proceeded much more efficiently by taking a few multimeter measurements and/or stimulating the system in ways revealing what and where the problem is. One of your tasks as a technical educator is to discourage this bad habit by refusing to tolerate random behavior during a troubleshooting exercise!

D.4 Practical topic coverage

Nearly every technical course teaches and tests students on definitions, basic concepts, and at least some form of quantitative analysis. If you really intend to prepare your students for the challenges of a career like instrumentation, however, you must cover far more than this. The following is a list of topics that should be represented in your curriculum every bit as prevalently as definitions, basic concepts, and math:

- *Qualitative* analysis of instrument systems (e.g. “Predict how the control system will respond if the flow rate *increases*”)
- *Qualitative* analysis of processes (e.g. “Predict what will happen to the pressure in reactor vessel R-5 if valve LV-21 *closes*”)
- Spatial relations (e.g. mapping wires in a schematic diagram to connection points in a pictorial diagram)
- Evaluating the validity of someone else’s diagnosis of a problem (e.g. “The last instrument technician to examine this system concluded the problem was a shorted cable. Based on the data presented here, do you agree or disagree with that conclusion?”)
- Identification of safety hazards and possible means of mitigation
- Documentation, both creating it and interpreting it
- Basic project management principles (e.g. scheduling of time, budgeting material and fiscal resources, limiting project scope, following through on “loose ends”)
- Mental math (e.g. approximate calculation without the use of computing equipment)
- Evaluation of real-life case studies (e.g. students read and answer questions on industry accident reports such as those published by the US Chemical Safety and Hazard Investigation Board)

These topics can and should be an explicit – not implicit – part of theory and lab (practical) instruction alike. I do not recommend teaching these topics in separate courses, but rather embedding them within each and every course taught in an Instrumentation program. By “explicit” I mean that these topics should be scheduled for discussion within lesson plans, included within student homework questions, appear as actual questions on exams, and individually demonstrated during labwork.

D.5 Principles, not procedures

One of the marks of a successful problem-solver is a habit of applying general principles to every new problem encountered. One of the marks of an ineffective problem-solver is a fixation on procedural steps. Sadly, most of the students I have encountered as a technical college teacher fall into this latter category, as well as a number of working instrument technicians.

Teachers share a large portion of the blame for this sad state of affairs. In an effort to get our students to a place where they are able to solve problems on their own, there is the temptation to provide them with step-by-step procedures for each type of problem they encounter. This is a fundamentally flawed approach to teaching, because a set of rigid procedures only works on a very specific set of problems. To be sure, your students might learn how to solve problems falling within this narrow field by following your algorithmic procedures, but they will be helpless when faced with problems not precisely fitting that same mold. In other words, they might be able to pass your exams but they will flounder when faced with real-world challenges, and you are utterly wasting their time if you are not preparing them for real-world challenges.

I am as guilty of this as any other teacher. When I first began teaching (the subject of electronics), I was dismayed at how difficult it was for students to grasp certain fundamental concepts, such as the analysis of series-parallel resistor circuits. Knowing that I had a very limited amount of time to get my students ready to pass the upcoming exam on series-parallel circuits, I decided to make things simpler for my students by repeatedly demonstrating a set of simple steps by which one could analyze and solve any series-parallel resistor circuit. Fellow instructors did the same thing, and gladly shared their procedures with me, including tips such as the use of different pen colors (black for drawing wires and components, red for writing current values and directional arrows, and blue for writing voltage values and braces) to help organize all the work. The procedure could be long-winded depending on how many nested levels of series-parallel resistors were in the circuit, but precisely followed it would never fail to yield the correct answers. Students greatly appreciated me giving them a set of step-by-step instructions they could follow.

The fallacy of this approach became increasingly evident to me as students would request repeated demonstrations on more and more example problems. I remember one particular classroom session, after having applied this procedure to at least a half-dozen example problems, that one of the students asked me to do one more example. “*Are you kidding?*” was the unspoken thought rushing through my mind, “*Just how many times must I demonstrate the same steps over and over before you can do it on your own?*” It suddenly occurred to me that my students were not learning how to solve problems – instead, they were merely memorizing a sequence of steps including keystrokes on their calculators. Despite all my effort, the only thing I was preparing them to do successfully was pass the upcoming exam, and that was only because the exam contained exactly the same types of problems I was beating to death on the whiteboard in front of class.

What I should have been doing instead was presenting to my students *only* the general principles of resistor circuits, which may be neatly summarized as such:

- Ohm’s Law ($V = IR$, where V , I , and R must all refer to the same resistor or same subset of resistors)
- Resistances in series add to make a larger total resistance ($R_{series} = R_1 + R_2 + \cdots R_n$)
- Resistances in parallel diminish to make a smaller total resistance ($R_{parallel} = \frac{1}{\frac{1}{R_1} + \frac{1}{R_2} + \cdots \frac{1}{R_n}}$)
- Current is the same through all series-connected components
- Voltage is the same across all parallel-connected components

It was not as though I had failed to present these principles often enough, nor that I had failed to demonstrate where these principles applied in the procedure. My fault was in giving students a comprehensive procedure in the first place, which had the unintended consequence of drawing their attention away from the fundamental principles. *The simple reason why a step-by-step procedure makes any problem easier to solve is because it eliminates the need for the student to apply general principles to that problem, which is the very thing I my students actually needed to learn.* To put it bluntly, a comprehensive procedure “does the thinking” for the student, because the application of general principles is already pre-determined and encoded into the steps of the procedure itself. What we get by robotically following the procedure is only an illusion of problem-solving competence. The real test of whether or not students have mastered the principles (rather than the procedure) is to check their performance on solving similar problems of different form, where the rote procedure is not applicable.

In order to teach students to approach problem-solving from a conceptual rather than procedural perspective, you must insist students show you how they make the links between general principles and the specifics of given problems. A useful tool for doing this is to have students maintain a notebook identifying and explaining general principles in their own words. You may choose to allow students the use of their own notepage or notecard on exams, as an incentive to tersely summarize all the major principles they will need to solve problems on exams.

An inverted classroom structure is well-suited for the encouragement of principle-based problem solving, in that it affords you the opportunity to see how students approach problems and to continually emphasize principles over procedures.

D.6 Assessing student learning

As a general rule, *high achievement only takes place in an atmosphere of high expectations*. Sometimes these expectations come from within: a self-motivated individual pushes himself or herself to achieve the extraordinary. Most often, the expectations are externally imposed: someone else demands extraordinary performance. One of your responsibilities as a teacher is to hold the standard of student performance high (but reasonable!), and this is done through valid, rigorous assessment.

When the time comes to assess your students' learning, prioritize performance assessment over written or verbal response. In other words, require that your students *demonstrate* their competence rather than merely explain it. Performance assessment takes more time than written exams, but the results are well worth it. Not only will you achieve a more valid measurement of your students' learning, but they will experience greater motivation to learn because they know they must put their learning into action.

Make liberal use of *mastery* assessments in essential knowledge and skill domains, where students must repeat a demonstration of competence as many times as necessary to achieve perfect performance. Not only does this absolutely guarantee students will learn what they should, but the prospect of receiving multiple opportunities to demonstrate knowledge or skill has the beneficial effect of relieving psychological stress for the student. Mastery assessment lends itself very well to the measurement of diagnostic ability.

An idea I picked up through a discussion on an online forum with someone from England regarding engineering education is the idea of breaking written exams into two parts: a *mastery* exam and a *proportional* exam. Students must pass the mastery exam(s) with 100% accuracy in order to receive a passing grade for each course, while the proportional exam is graded like any regular exam (with a score between 0% and 100%) and contributes to their letter grade. Students are given multiple opportunities to pass each mastery exam, with different versions of the mastery exam given at each re-take. Mastery exams cover all the basic concepts, with very straight-forward questions (no tricks or ambiguous wording). The proportional exam, by contrast, is a single-effort test filled with challenging problems requiring high-level thinking. By dividing exams into two parts, it is possible to guarantee the entire class has mastered basic concepts while challenging even the most capable students.

Another unconventional assessment strategy is to create multi-stage exams, where the grade or score received for the exam depends on the highest level passed. I have applied this to the subject of PLC programming: a large number of programming projects are provided as examples, each one fitting into one of four categories of increasing difficulty. The first level is the minimum required to pass the course, while the fourth level is so challenging that only a few students will be able to pass it in the time given. For each of these levels, the student is given the design parameters (e.g. "program a motor start-stop system with a timed lockout preventing a re-start until at least 15 seconds has elapsed"); a micro-PLC; a laptop computer with the PLC programming software; the necessary switches, relays, motors, and other necessary hardware; and 1 hour of time to build and program a working system. There are too many example projects provided for any student to memorize solutions to them all, especially when no notes are allowed during the assessment (only manufacturer's documentation for the PLC and other hardware). This means the student must demonstrate both mastery of the basic PLC programming and wiring elements, as well as creative design skills to arrive at their own solution to the programming problem. There is no limit to the number of attempts a student may take to pass a given level, and no penalty for failed efforts. Best

of all, this assessment method demands little of the instructor, as the working project “grades” itself.

My philosophy on assessment is that good assessment is actually more important than good instruction. If the assessments are valid and rigorous, student learning (and instructor teaching!) will rise to meet the challenge. However, even the best instruction will fail to produce consistently high levels of student achievement if students know their learning will never be rigorously assessed. In a phrase, *assessment drives learning*.

For those who might worry about an emphasis on assessment encouraging teachers to “teach to the test,” I offer this advice: there is nothing wrong with teaching to the test so long as the test is valid! Educators usually avoid teaching to the test out of a fear students might pass the test(s) without actually learning what they are supposed to gain from taking the course. If this is even possible, it reveals a fundamental problem with the test: it does not actually measure what you want students to know. A valid test is one that cannot be “foiled” by teaching in any particular way. Valid tests challenge students to think, and cannot be passed through memorization. Valid tests avoid asking for simple responses, demanding students articulate reasoning in their answers. *Valid tests are passable only by competence.*

Another important element of assessment is long-term review. You should design the courses in such a way that important knowledge and skill areas are assessed on an ongoing basis up through graduation. Frequent review of foundational concepts is a best practice for attaining mastery in any domain.

D.7 Summary

To summarize some of the key points and concepts for teaching:

- Do not waste class time transmitting facts to students – let the students research facts outside of class
- Use class time to develop high-level thinking skills (e.g. problem-solving, diagnostic techniques, metacognition)
- Use Socratic dialogue to challenge each and every student on the subject matter
- Focus on general principles, not specific procedures
- Make labwork as realistic as possible
- Build diagnostic skill by first exercising deductive reasoning, as a prelude to inductive reasoning
- Incorporate frequent troubleshooting exercises in the lab, with students diagnosing realistic faults in instrument systems
- Include a broad range of practical topics and aspects in all coursework rather than fall into the convention of focusing on memorizing definitions, stating concepts, and performing quantitative calculations
- Assess student learning validly and rigorously
- Review important knowledge and skill areas continually until graduation – build this review into the program courses themselves (homework, quizzes, exams) rather than relying on ad hoc review

One final piece of advice for educators at every level: *it is better to teach a few things well than to teach many things poorly!* If external constraints force you to “cover” too much material in too little time, focus on making each learning exercise as integrative as possible, so students will experience different topics in ways that reinforce and give context to each other.

Appendix E

Contributors

This is an open-source book, which means everyone has a legal write to modify it to their liking. As the author, I freely accept input from readers that will make this book better. This appendix exists to give credit to those readers who have made substantial contributions to this book.

Sadly, this list does not show the names of *every* person who has helped me identify and correct minor typographical and grammatical errors. The list of names and errors would be quite substantial, I must admit. Those persons who are listed for their identification of typographical errors have earned a place on the list through sheer volume of errors found. I am indebted to my students, and to readers around the world for their careful reading of the text and their considerate feedback.

If you find this textbook exceeds your expectations, know that it is primarily because of this detailed reader feedback. When a book is edited by its author, based on daily feedback from the people using it, errors as well as unclear explanations get identified and corrected at a much more rapid pace than what is typical during a traditional (published) textbook editing cycle. As an author, you might think you have explained something perfectly well, only to find your readers mystified. It is only by discovering your readers are baffled by your “good” explanation that you will know to edit that explanation, or to take a completely different approach in explaining the concept. Then, you need to get more feedback after the edit(s), to see whether or not you’ve made an improvement.

This kind of tight feedback loop between authors and readers simply doesn’t happen in a traditional textbook publishing model. In an open-source, self-published model, however, this kind of feedback is not only possible, but is also typical. For this reason, I strongly encourage teachers everywhere to consider writing their own open-source textbooks, using their students as editors to prove whether or not their writing is effective. The world is unfortunately filled with poor-quality texts. Let’s change this sad state of affairs, one open-source textbook at a time!

E.1 Error corrections

Abdelaziz, Ahmed (June 2014)

- Identified numerous typographical errors throughout the book, mostly *repeated repeated* words.

Ames, Ben (January 2015)

- Identified typographical error in the *Problem-Solving and Diagnostic Strategies* chapter.

Balbaa, Amr (May 2015)

- Identified mathematical error in the *Continuous Fluid Flow Measurement* chapter.

Barton, Bruce (January 2015, March 2015, April 2015)

- Identified typographical errors in the *Discrete Control Elements* chapter regarding solenoid valve construction, in the *Process Dynamics and PID Controller Tuning* chapter, and also in the *Digital Data Acquisition and Networks* chapter.

Brierly, Ryan (April 2016)

- Identified typographical error in the *Physics* chapter regarding transition points between laminar and turbulent flow regimes.

Brown, Kevin (February 2011)

- Identified typographical error in the *Control Valves* chapter.

Bultsma, Brent (September 2011-June 2012)

- Identified multiple typographical errors throughout the book.

Chong, Yi-Vonne (August 2017)

- Identified typographical error in a protective relay wiring diagram.

Douglas, Jon (Spring 2017)

- Suggested the inclusion of *self-purged impulse lines* to the book and provided documentation of this practice from industry.

Edwards, William (February 2013)

- Identified error in the *Calculus* chapter, where Δt was confused with Δm .

Esher, Cynthia A. (December 2009)

- Identified error of referring to “SAMA” diagrams. I changed these references to “functional” diagrams instead (*Instrumentation Documents* chapter).

Floyd, Wade (November 2015)

- Identified an error the distance between Earth and the sun *Chemistry* chapter: the distance is 93 *million* miles, not 93 *billion* miles as I originally had it.

Gassman, Jessica (February 2013)

- Suggested clarification in the *Closed-Loop Control* chapter regarding the auto/manual transfer of pneumatic controllers.

Glundberg, Blake (February 2011, April 2011)

- Identified typographical errors (*Control Valves* chapter, *Digital Data Acquisition and Networks* chapter).

Ireland, Chandler (February 2015)

- Identified typographical error (*Closed-Loop Control* chapter).

Joe, Savio (February 2013)

- Identified typographical error (*Instrument Connections* chapter).

Ketteridge, Tim and Dieckman, Lloyd (April 2013)

- Identified conceptual errors in the IP and TCP sections of the *Digital Data Acquisition and Network* chapter. The task of portioning large data blocks into segments and packets is the job of TCP (or Application-layer protocols), not IP. IP does, however, support fragmentation of large packets into smaller packets and reassembly into large packets again if needed for communication along network types that cannot support large packet sizes. I had these concepts conflated.

Ketteridge, Tim (December 2013)

- Identified an error in measurement units for the Ideal Gas Law ($PV = NkT$). When using this form of the Ideal Gas Law, P needs to be in units of Pascals (not atmospheres), and V needs to be in units of cubic meters (not liters).

Marques, Tiago (July 2012)

- Identified lack of origin lines in linear graphs where 4-20 mA signals are related to practical measurement ranges (*Analog Electronic Instrumentation* chapter).

Mhyre, Phil (January 2012)

- Identified spelling error on the word “desiccant” (*Discrete Control Elements* chapter).

Nobiling, David (July 2016)

- Identified typographical error in the (*Digital data acquisition and networks* chapter), where the terms “PLC” and “VFD” were interchanged.

Pulido, Alberto (May 2012)

- Identified conceptual error regarding Fieldbus terminator resistor packages (*FOUNDATION Fieldbus Instrumentation* chapter).

Sangani, Champa (September 2009) and Brainard, Ben (January 2012)

- Identified calculation error in milliamp-to-pH scaling problem (*Analog Electronic Instrumentation* chapter).

Saxton, Trevor (August 2015)

- Identified mathematical errors in the “Phasor Arithmetic” section of the (*AC Electricity* chapter).

Schultz, Steven (February 2011)

- Identified error in a pilot-loaded pressure regulator design depicted in the *Control Valves* chapter.

Sloan, Wyatt (October 2012)

- Identified typographical error in the *Continuous Temperature Measurement* chapter. Later (January 2013) identified multiple places where the word “identify” was used instead of “identity.”

Sobczak, Jacek (October 2015)

- Identified missing files in the .tar archive, suggested changes to the L^AT_EX source code producing better page numbering in the book's table of contents, and identified broken links on the website.

Thompson, Brice (June 2009)

- Identified errors in high/low select and high/low limit function illustrations (*Basic Process Control Strategies* chapter).

TooToonchy, Hossein (June 2014)

- Identified typographical errors in the (*FOUNDATION Fieldbus Instrumentation* chapter).

Tsiporenko, Michael (June 2009)

- Identified typographical error in the *Introduction to Industrial Instrumentation* chapter.

Villajulca, Jose Carlos (August 2011)

- Identified error in PID tuning response for a “generic” simulated process in the *Process Dynamics and PID Controller Tuning* chapter.

Wingerter, John (January 2017)

- Identified an oversight in explaining safety considerations during control valve disassembly, in the *Disassembly of a Sliding-Stem Control Valve* appendix.

Young, Jeanne (April-May 2012)

- Identified errors in proximity switch wire color codes and suggested correction in the *Programmable Logic Controllers* chapter.
- Identified pattern inconsistent with industry convention in motor control circuit schematic diagrams within the *Discrete Control Elements* chapter.

E.2 New content

Coy, John J. et. al. (December 2012)

- Sampled graphic illustrations from his NASA Reference Publication “Gearing” (1152), showing different types of gear trains for inclusion in the *Simple Machines* section of the *Physics* chapter. This NASA document, being a US government publication unclassified and approved for “unlimited” distribution, lies within the public domain and so may be freely sampled into other works such as this one.

Dennis, Japheth (2011-2012 academic year)

- Suggested additional examples of PID controller responses to graph, to help illustrate the unique features of each action.

Faydor, L. Litvin et. al. (December 2012)

- Sampled 5-planet planetary gear set illustration from his NASA publication “New Design and Improvement of Planetary Gear Trains”, for inclusion in the *Simple Machines* section of the *Physics* chapter. This NASA document, being a US government publication unclassified and approved for “unlimited” distribution, lies within the public domain and so may be freely sampled into other works such as this one. Also sampled helicopter transmission gear set illustration from his NASA publication “Handbook on Face Gear Drives With a Spur Involute Pinion”.

Goertz, Kevin (2006-2007 academic year)

- Took photographs of various flowmeters, control valves, and an insertion pH probe assembly.

Marshall, Travis (2012-2013 academic year)

- Contributed analogy of a “shadow” to the explanation of absorption spectroscopy, after helping to install a laser stack gas analyzer during a summer internship and helping to explain that instrument’s operating principle to himself.

Poelma, John (2010-2011 academic year)

- Took photographs of various pressure vessels, instruments, control valves, and other process hardware at NASA’s Stennis Space Center in Mississippi.

Toner, Aaron (August 2012)

- Took photographs of his self-constructed MicroLogix 1100 / C-More Micro PLC trainer for inclusion in this book.

Appendix F

Creative Commons Attribution License

F.1 A simple explanation of your rights

This is an “open-source” textbook, which means the entirety of it is freely available for public perusal, reproduction, distribution, and even modification. All digital “source” files comprising this textbook reside at the following website:

<http://ibiblio.org/kuphaldt/socratic/sinst/book/>

The Creative Commons Attribution license grants you (the recipient), as well as anyone who might receive my work from you, the right to freely use it. This license also grants you (and others) the right to modify my work, so long as you properly credit my original authorship and declare these same rights (to my original work) for your own readers. My work is copyrighted under United States law, but this license grants everyone else in the world certain freedoms not customarily available under full copyright. This means no one needs to ask my permission, or pay any royalties to me, in order to read, copy, distribute, publish, or otherwise use this book.

If you choose to modify my work, you will have created what legal professionals refer to as a *derivative work*. The Creative Commons license broadly groups derivative works under the term *adapted material*. In simple terms, the fundamental restriction placed on you when you do this is you must properly credit me for the portions of your adaptation that are my original work. Otherwise, you may treat your adaptation the same way you would treat a completely original work of your own. This means you are legally permitted to enjoy full copyright protection for your adaptation, up to and including exclusive rights of reproduction and distribution. In other words, this license does *not* bind your derivative work under the same terms and conditions I used to release my original work, although it does require you notify your readers of this free source text and the License under which it is offered.

The practical upshot of this is you may modify my work and re-publish it as you would any other book, with the full legal right to demand royalties, restrict distributions, etc. This does not compromise the freedom of my original work, because that is still available to everyone under the terms and conditions of the Attribution license¹. It does, however, protect the investment(s) you make in creating the adaptation by allowing you to release the adaptation under whatever terms you see fit (so long as those terms comply with current intellectual property laws, of course).

In summary, the following “legalese” is actually a very good thing for you, the reader of my book. It grants you permission to do so much more with this text than what you would be legally allowed to do with any other (traditionally copyrighted) book. It also opens the door to open collaborative development, so it might grow into something far better than what I alone could create.

¹You *cannot* pass my original work to anyone else under different terms or conditions than the Attribution license. That is called *sublicensing*, and the Attribution license forbids it. In fact, any re-distribution of my original work must come with a notice to the Attribution license, so anyone receiving the book through you knows their rights.

F.2 Legal code

Creative Commons Attribution 4.0 International Public License

By exercising the Licensed Rights (defined below), You accept and agree to be bound by the terms and conditions of this Creative Commons Attribution 4.0 International Public License ("Public License"). To the extent this Public License may be interpreted as a contract, You are granted the Licensed Rights in consideration of Your acceptance of these terms and conditions, and the Licensor grants You such rights in consideration of benefits the Licensor receives from making the Licensed Material available under these terms and conditions.

Section 1 – Definitions.

a. **Adapted Material** means material subject to Copyright and Similar Rights that is derived from or based upon the Licensed Material and in which the Licensed Material is translated, altered, arranged, transformed, or otherwise modified in a manner requiring permission under the Copyright and Similar Rights held by the Licensor. For purposes of this Public License, where the Licensed Material is a musical work, performance, or sound recording, Adapted Material is always produced where the Licensed Material is synched in timed relation with a moving image.

b. **Adapter's License** means the license You apply to Your Copyright and Similar Rights in Your contributions to Adapted Material in accordance with the terms and conditions of this Public License.

c. **Copyright and Similar Rights** means copyright and/or similar rights closely related to copyright including, without limitation, performance, broadcast, sound recording, and Sui Generis Database Rights, without regard to how the rights are labeled or categorized. For purposes of this Public License, the rights specified in Section 2(b)(1)-(2) are not Copyright and Similar Rights.

d. **Effective Technological Measures** means those measures that, in the absence of proper authority, may not be circumvented under laws fulfilling obligations under Article 11 of the WIPO Copyright Treaty adopted on December 20, 1996, and/or similar international agreements.

e. **Exceptions and Limitations** means fair use, fair dealing, and/or any other exception or limitation to Copyright and Similar Rights that applies to Your use of the Licensed Material.

f. **Licensed Material** means the artistic or literary work, database, or other material to which the Licensor applied this Public License.

g. **Licensed Rights** means the rights granted to You subject to the terms and conditions of this Public License, which are limited to all Copyright and Similar Rights that apply to Your use of the Licensed Material and that the Licensor has authority to license.

h. **Licensor** means the individual(s) or entity(ies) granting rights under this Public License.

i. **Share** means to provide material to the public by any means or process that requires permission under the Licensed Rights, such as reproduction, public display, public performance, distribution, dissemination, communication, or importation, and to make material available to the

public including in ways that members of the public may access the material from a place and at a time individually chosen by them.

j. **Sui Generis Database Rights** means rights other than copyright resulting from Directive 96/9/EC of the European Parliament and of the Council of 11 March 1996 on the legal protection of databases, as amended and/or succeeded, as well as other essentially equivalent rights anywhere in the world.

k. **You** means the individual or entity exercising the Licensed Rights under this Public License. **Your** has a corresponding meaning.

Section 2 – Scope.

a. License grant.

1. Subject to the terms and conditions of this Public License, the Licensor hereby grants You a worldwide, royalty-free, non-sublicensable, non-exclusive, irrevocable license to exercise the Licensed Rights in the Licensed Material to:

A. reproduce and Share the Licensed Material, in whole or in part; and

B. produce, reproduce, and Share Adapted Material.

2. Exceptions and Limitations. For the avoidance of doubt, where Exceptions and Limitations apply to Your use, this Public License does not apply, and You do not need to comply with its terms and conditions.

3. Term. The term of this Public License is specified in Section 6(a).

4. Media and formats; technical modifications allowed. The Licensor authorizes You to exercise the Licensed Rights in all media and formats whether now known or hereafter created, and to make technical modifications necessary to do so. The Licensor waives and/or agrees not to assert any right or authority to forbid You from making technical modifications necessary to exercise the Licensed Rights, including technical modifications necessary to circumvent Effective Technological Measures. For purposes of this Public License, simply making modifications authorized by this Section 2(a)(4) never produces Adapted Material.

5. Downstream recipients.

A. Offer from the Licensor – Licensed Material. Every recipient of the Licensed Material automatically receives an offer from the Licensor to exercise the Licensed Rights under the terms and conditions of this Public License.

B. No downstream restrictions. You may not offer or impose any additional or different terms or conditions on, or apply any Effective Technological Measures to, the Licensed Material if doing so restricts exercise of the Licensed Rights by any recipient of the Licensed Material.

6. No endorsement. Nothing in this Public License constitutes or may be construed as permission

to assert or imply that You are, or that Your use of the Licensed Material is, connected with, or sponsored, endorsed, or granted official status by, the Licensor or others designated to receive attribution as provided in Section 3(a)(1)(A)(i).

b. Other rights.

1. Moral rights, such as the right of integrity, are not licensed under this Public License, nor are publicity, privacy, and/or other similar personality rights; however, to the extent possible, the Licensor waives and/or agrees not to assert any such rights held by the Licensor to the limited extent necessary to allow You to exercise the Licensed Rights, but not otherwise.

2. Patent and trademark rights are not licensed under this Public License.

3. To the extent possible, the Licensor waives any right to collect royalties from You for the exercise of the Licensed Rights, whether directly or through a collecting society under any voluntary or waivable statutory or compulsory licensing scheme. In all other cases the Licensor expressly reserves any right to collect such royalties.

Section 3 – License Conditions.

Your exercise of the Licensed Rights is expressly made subject to the following conditions.

a. Attribution.

1. If You Share the Licensed Material (including in modified form), You must:

A. retain the following if it is supplied by the Licensor with the Licensed Material:

i. identification of the creator(s) of the Licensed Material and any others designated to receive attribution, in any reasonable manner requested by the Licensor (including by pseudonym if designated);

ii. a copyright notice;

iii. a notice that refers to this Public License;

iv. a notice that refers to the disclaimer of warranties;

v. a URI or hyperlink to the Licensed Material to the extent reasonably practicable;

B. indicate if You modified the Licensed Material and retain an indication of any previous modifications; and

C. indicate the Licensed Material is licensed under this Public License, and include the text of, or the URI or hyperlink to, this Public License.

2. You may satisfy the conditions in Section 3(a)(1) in any reasonable manner based on the medium, means, and context in which You Share the Licensed Material. For example, it may be

reasonable to satisfy the conditions by providing a URI or hyperlink to a resource that includes the required information.

3. If requested by the Licensor, You must remove any of the information required by Section 3(a)(1)(A) to the extent reasonably practicable.

4. If You Share Adapted Material You produce, the Adapter's License You apply must not prevent recipients of the Adapted Material from complying with this Public License.

Section 4 – Sui Generis Database Rights.

Where the Licensed Rights include Sui Generis Database Rights that apply to Your use of the Licensed Material:

a. for the avoidance of doubt, Section 2(a)(1) grants You the right to extract, reuse, reproduce, and Share all or a substantial portion of the contents of the database;

b. if You include all or a substantial portion of the database contents in a database in which You have Sui Generis Database Rights, then the database in which You have Sui Generis Database Rights (but not its individual contents) is Adapted Material; and

c. You must comply with the conditions in Section 3(a) if You Share all or a substantial portion of the contents of the database.

For the avoidance of doubt, this Section 4 supplements and does not replace Your obligations under this Public License where the Licensed Rights include other Copyright and Similar Rights.

Section 5 – Disclaimer of Warranties and Limitation of Liability.

a. Unless otherwise separately undertaken by the Licensor, to the extent possible, the Licensor offers the Licensed Material as-is and as-available, and makes no representations or warranties of any kind concerning the Licensed Material, whether express, implied, statutory, or other. This includes, without limitation, warranties of title, merchantability, fitness for a particular purpose, non-infringement, absence of latent or other defects, accuracy, or the presence or absence of errors, whether or not known or discoverable. Where disclaimers of warranties are not allowed in full or in part, this disclaimer may not apply to You.

b. To the extent possible, in no event will the Licensor be liable to You on any legal theory (including, without limitation, negligence) or otherwise for any direct, special, indirect, incidental, consequential, punitive, exemplary, or other losses, costs, expenses, or damages arising out of this Public License or use of the Licensed Material, even if the Licensor has been advised of the possibility of such losses, costs, expenses, or damages. Where a limitation of liability is not allowed in full or in part, this limitation may not apply to You.

c. The disclaimer of warranties and limitation of liability provided above shall be interpreted in a manner that, to the extent possible, most closely approximates an absolute disclaimer and waiver of all liability.

Section 6 – Term and Termination.

a. This Public License applies for the term of the Copyright and Similar Rights licensed here. However, if You fail to comply with this Public License, then Your rights under this Public License terminate automatically.

b. Where Your right to use the Licensed Material has terminated under Section 6(a), it reinstates:

1. automatically as of the date the violation is cured, provided it is cured within 30 days of Your discovery of the violation; or
2. upon express reinstatement by the Licensor.

For the avoidance of doubt, this Section 6(b) does not affect any right the Licensor may have to seek remedies for Your violations of this Public License.

c. For the avoidance of doubt, the Licensor may also offer the Licensed Material under separate terms or conditions or stop distributing the Licensed Material at any time; however, doing so will not terminate this Public License.

d. Sections 1, 5, 6, 7, and 8 survive termination of this Public License.

Section 7 – Other Terms and Conditions.

a. The Licensor shall not be bound by any additional or different terms or conditions communicated by You unless expressly agreed.

b. Any arrangements, understandings, or agreements regarding the Licensed Material not stated herein are separate from and independent of the terms and conditions of this Public License.

Section 8 – Interpretation.

a. For the avoidance of doubt, this Public License does not, and shall not be interpreted to, reduce, limit, restrict, or impose conditions on any use of the Licensed Material that could lawfully be made without permission under this Public License.

b. To the extent possible, if any provision of this Public License is deemed unenforceable, it shall be automatically reformed to the minimum extent necessary to make it enforceable. If the provision cannot be reformed, it shall be severed from this Public License without affecting the enforceability of the remaining terms and conditions.

c. No term or condition of this Public License will be waived and no failure to comply consented to unless expressly agreed to by the Licensor.

d. Nothing in this Public License constitutes or may be interpreted as a limitation upon, or waiver of, any privileges and immunities that apply to the Licensor or You, including from the legal processes of any jurisdiction or authority.

Creative Commons is not a party to its public licenses. Notwithstanding, Creative Commons may elect to apply one of its public licenses to material it publishes and in those instances will be considered the “Licensor.” Except for the limited purpose of indicating that material is shared under a Creative Commons public license or as otherwise permitted by the Creative Commons policies published at creativecommons.org/policies, Creative Commons does not authorize the use of the trademark “Creative Commons” or any other trademark or logo of Creative Commons without its prior written consent including, without limitation, in connection with any unauthorized modifications to any of its public licenses or any other arrangements, understandings, or agreements concerning use of licensed material. For the avoidance of doubt, this paragraph does not form part of the public licenses.

Creative Commons may be contacted at creativecommons.org.

Index

- C_d factor, 2187
- C_v factor, 2179
- K_v factor, 2179
- Q (quality factor of RLC filter circuit), 436
- ΔT , meaning of, 128
- α particle radiation, 1497
- β particle radiation, 1497
- γ ray radiation, 1497
- λ (failure rate), 2628
- μ metal, 623
- $\times 10$ oscilloscope probe, 2557
- s variable, 404, 406, 437
- “Desktop Process” PID tuning tool, 2491
- 10 to 50 mA, 893
- 2-wire RTD circuit, 1520
- 2-wire transmitter, 311, 891
- 21 relay, distance protection, 2025
- 3 to 15 PSI, 916
- 3-element boiler feedwater control, 2529
- 3-valve manifold, 1359
- 3-way globe valve, 2088, 2161
- 3-wire RTD circuit, 1522, 1523
- 4 to 20 mA, 861, 915
- 4-wire RTD circuit, 1521
- 4-wire resistance measurement circuit, 1521, 1756
- 4-wire transmitter, 889
- 4-wire transmitter, active output, 893
- 4-wire transmitter, passive output, 893
- 4-wire transmitter, sinking output, 893
- 4-wire transmitter, sourcing output, 893
- 5-point calibration, 1278
- 5-valve manifold, 1362
- 50 relay, instantaneous overcurrent protection, 1996
- 51 relay, time overcurrent protection, 1996
- 67 relay, directional current protection, 2023
- 802.11, 1045
- 802.3, 1032, 1038
- 802.4, 1042
- 802.5, 1042
- 86 relay, auxiliary/lockout, 2049
- 87 relay, differential protection, 2005
- ABB 800xA distributed control system (DCS), 2376
- Absolute addressing, 801
- Absolute pressure, 186, 1312
- Absolute viscosity, 199
- Absolute zero, 123
- Absorption spectroscopy, 256
- AC, 355
- AC excitation, magnetic flowmeter, 1686
- AC generator synchronization, 1914
- AC induction motor, 722, 2813
- Accelerometer, 2730
- Acceptance testing, 644
- Acid, 286
- Acid, strong, 286
- Acid, weak, 286
- Acre, defined, 57
- Action, controller, 508, 2278, 2585
- Activation energy, 275
- Active intrinsic safety barrier, 2608
- Active reading, 2744
- Actuator, 2108
- Actuator, valve, 2079
- Acyclic communication, Fieldbus, 1155, 1159
- Adaptive gain controller, 2428
- ADC, 784, 977, 991
- Additive transformer windings, 1959
- Address, 1104
- Address Resolution Protocol, 1074, 1083
- Administrator privilege, computer, 2728
- Admittance, 360

- Aerosol, [229](#)
- Aeroswage SX-1 tubing tool, [590](#)
- AGA Report #11, [1725](#)
- AGA Report #3, [1640](#), [1647](#)
- AGA Report #7, [1666](#)
- AGA Report #9, [1693](#)
- Air dryer, [698](#)
- Air-to-close valve, [2128](#)
- Air-to-open valve, [2128](#)
- Alarm, annunciator, [526](#)
- Alarm, process, [2386](#)
- Algorithm, [2272](#)
- Algorithm, control, [540](#)
- Aliasing, [996](#)
- Alkaline, [286](#)
- Alkaline, strong, [287](#)
- Alkaline, weak, [287](#)
- Allen-Bradley ControlLogix 5000 PLC, [771](#)
- Allen-Bradley Data Highway (DH) network, [1058](#)
- Allen-Bradley GuardLogix DCSRT safety instruction, [2681](#)
- Allen-Bradley GuardLogix PLC, [2681](#)
- Allen-Bradley MicroLogix 1000 PLC, [855](#)
- Allen-Bradley PLC-5, [768](#)
- Allen-Bradley SLC 500 PLC, [769](#), [878](#)
- Alpha particle radiation, [1497](#)
- Altek model 334A loop calibrator, [906](#)
- Alternating current, [355](#)
- Alternating motor control, [833](#)
- American alpha value, RTD, [1519](#)
- American Gas Association, [1640](#), [1647](#), [1666](#), [1693](#), [1725](#)
- Ametek model 1500 induction relay, [680](#)
- Ammeter, clamp-on, [341](#)
- Ammonia slip, NOx control, [1849](#)
- Amp-turn, [340](#)
- Ampère, André, [299](#)
- Ampere, [299](#)
- Amplifier, instrumentation, [1006](#)
- amu, [262](#)
- Analog input, single-ended, [1013](#)
- Analog-to-digital converter, [784](#), [977](#), [991](#)
- Analyzer, [1302](#)
- Anderson, Norman A., [3](#)
- Andrew, William G, [3](#)
- Angle of repose, [1472](#), [2072](#)
- Anion, [228](#), [1751](#)
- Annubar, [1626](#), [1737](#)
- Annunciator, [526](#)
- Anode, [1751](#)
- ANPT pipe threads, [572](#)
- ANSI codes for power systems, [1992](#)
- ANSI pressure classes for flanges, [567](#)
- ANSI/IEEE standard 754-1985 for floating-point numbers, [983](#), [1126](#)
- Antenna, half-wave, [491](#), [1208](#)
- Antenna, quarter-wave, [492](#), [1208](#)
- Antenna, whip, [492](#), [1208](#)
- Antenna, Yagi, [493](#), [1208](#)
- Anti-aliasing filter, [998](#)
- Anti-resonance resistor, [431](#)
- Anti-virus, computer, [2728](#)
- API, degrees, [180](#)
- Application blacklisting, computer, [2728](#)
- Application whitelisting, computer, [2728](#)
- Arbitration, channel, [1040](#)
- Arc blast, [371](#), [729](#)
- Arc chute, circuit breaker, [1929](#)
- Arc flash, [635](#), [729](#)
- Archimedes' Principle, [191](#), [1453](#)
- Area, defined, [57](#)
- Armature, [341](#)
- ARP, [1074](#), [1083](#), [1084](#)
- As-found calibration, [1260](#)
- As-left calibration, [1260](#)
- Asbestos valve stem packing, [2104](#)
- ASCII, [989](#)
- ASCII Modbus frames, [1111](#)
- ASCO model NH90 "Hydramotor" linear actuator/positioner, [935](#)
- ASCO solenoid valve, [708](#)
- Ashcroft temperature switch, [683](#)
- Assignment, computer programming, [2389](#)
- Asynchronous data transfer, [1026](#), [1032](#)
- Atmospheres, [188](#)
- Atom, [228](#)
- Atomic clock, [1282](#), [1285](#)
- Atomic mass, [228](#), [230](#), [237](#)
- Atomic mass units, [262](#)
- Atomic number, [228](#), [230](#), [236](#)
- Atomic weight, [228](#), [230](#), [237](#)
- Aufbau order, [247](#)

- Auto-tuning PID controller, 2456
- Automatic mode, 500, 2383
- Autotransformer, 472
- Autotransformer, boost, 472
- Autotransformer, buck, 473
- Auxiliary contact, 739, 794
- Averaging Pitot tube, 1625
- Avogadro's number, 262

- B.I.F. Universal Venturi tube, 1629
- B/W Controls model 1500 induction relay, 680
- Babbage, Charles, 1178
- Back-calculation variable, FOUNDATION
Fieldbus function block programming,
1164, 1170, 2570, 2573
- Background substances, 1810
- Backpressure, nozzle, 921
- Backup (remote), protective relay, 2046
- BACnet, 2364
- Baffle, 921
- Bailey Infi90 distributed control system, 2376
- Bailey Net90 distributed control system, 2376
- Balance beam scale, 923
- Ball valve, 691, 2092
- Ball valve, characterized, 2093
- Ball valve, segmented, 2093
- Balling, degrees, 181
- Balmer series, 255
- Bang-bang control, 2275
- Bar, 188
- Bara, 188
- Barg, 188
- Bark, degrees, 181
- Barkhausen criterion (loop oscillation), 2439,
2443
- Barometer, 1312
- Barrier circuit, intrinsic safety, 2607
- Base, 286
- Base unit, 72
- Base, strong, 287
- Base, weak, 287
- basisk Siemens PLC exploit, 2729
- Bastard voltage, 456
- Bathtub curve, 2634
- Baud rate, 1031
- Baudot code, 988, 1025, 1027

- Baumé, degrees, 180
- Bazovsky, Igor, 2613, 2618
- Beam splitter, 1833, 1836
- Beer-Lambert Law, 1810
- Bel, 1210
- Bell 202 FSK standard, 1097
- Bell 202 standard, 1029
- Belleville washer spring, 1937, 2101
- Bellingham Technical College, 3191
- Bellows, 924, 1314
- Bellows packing seal, 2105
- Belt, 117
- Bench set, control valve, 2137
- Bend radius, optical fiber, 642
- Bently-Nevada model 1701 FieldMonitor
vibration monitor, 1890
- Bently-Nevada model 3300 vibration monitor,
1888
- Bently-Nevada vibration monitoring equipment,
1884
- BER (bit error rate), 1226
- Beresford, Dillon, 2729
- Bernoulli's equation, 208, 1590
- Bernoulli, Daniel, 208
- Beta particle radiation, 1497
- Beta ratio of flow element, 1609, 1646
- Bethlehem flow tube, 1629
- Bevel gear, 111
- Bi-metal strip, 1511
- Biddle Versa-Cal loop calibrator, 911
- Binary numeration, 978
- Biological oxygen demand, 2272
- Bit, 777, 978
- Bit error rate, 1226
- Bit rate, 1030
- Black, Harold, 942
- Blackbody, 1570
- Blackbody calibrator, 1295
- Blacklisting, application, 2728
- Bleed port adapter test accessory, 1367
- Bleed valve fitting, 1365
- Bleeding pneumatic relay, 937
- BLEVE, 161
- Blind, pipe, 569
- Blocking, protective relay, 2019, 2040, 2048, 2050
- Blowdown pressure, safety valve, 2658, 2665

- Bluetooth, [2736](#)
- Bluff body, [1671](#)
- BMS, [840](#), [2687](#)
- BNC style connector, [629](#), [1889](#)
- BNU, Fieldbus, [1161](#)
- BOD, [2272](#)
- Bode plot, [416](#)
- Body, valve, [2079](#)
- Boiling Liquid Expanding Vapor Explosion (BLEVE), [161](#)
- Boiling point of water, [1291](#)
- Boiling Water Reactor (BWR), [2698](#)
- Bolted fault, [370](#)
- Bonneville Power Administration, [1910](#)
- Boolean variable, [2389](#)
- Booster relay, [2170](#)
- Boosting autotransformer, [472](#)
- Bourdon tube, [575](#), [1314](#), [1316](#)
- Boyle's Law, [198](#)
- bps, [1030](#)
- Branch tee fitting, [586](#)
- Branching, computer programming, [2391](#)
- Brazosport College, [1621](#)
- Brick, Fieldbus coupler, [1145](#)
- Bridge circuit, [327](#)
- British Thermal Unit, [125](#)
- Brix, degrees, [181](#)
- Broadcast address, IP, [1080](#)
- Brush, DC motor, [2237](#)
- Brute force attack, password, [2727](#)
- BSPP pipe threads, [574](#)
- BSPT pipe threads, [573](#)
- BTU, [125](#)
- Bubble tube, [1424](#)
- Bubble-tight valve shut-off, [2108](#)
- Bucket, motor control, [738](#)
- Bucking autotransformer, [473](#)
- Buffer solution, [1302](#), [1765](#), [1781](#)
- Bulk modulus, [1467](#), [1469](#), [1487](#), [1644](#), [1689](#), [1690](#)
- Bulkhead, tube, [585](#)
- Buoyancy, [191](#), [1453](#)
- Buoyant test of density, [192](#)
- Burden, instrument transformer, [1973](#)
- Burgess, George, [2746](#)
- Burn-in period, [2634](#), [2641](#)
- Burner Management System (BMS), [840](#), [2687](#)
- Burnout, thermocouple, [1560](#)
- Burst mode, HART, [1105](#)
- Bus duct, [611](#)
- Bus, electrical power, [1898](#)
- Bus, power system, [1917](#)
- Busway, [611](#)
- Butterfly valve, [691](#), [2092](#)
- BWR, [2698](#)
- Byte, [979](#)
- Cable pulling, [606](#)
- Cable tray, [608](#)
- Cache, ARP, [1083](#)
- Cage-guided globe valve, [2086](#)
- Cageless displacer level instrument, [1452](#)
- Calculus, Fundamental Theorem of, [26](#)
- Calibration, [1251](#), [1302](#)
- Calibration gas, [264](#), [1262](#), [1304](#), [1835](#), [1845](#), [1859](#)
- Calibration management software, [1264](#)
- Calibration, dry versus wet, [1453](#)
- Calibration, pH instrument, [1781](#)
- Calibrator, loop, [906](#)
- Callendar-van Dusen formula, [1518](#)
- calorie, [125](#)
- Calorie, dietary, [125](#)
- Cam, [2150](#)
- CAN network, [1044](#)
- CANDU nuclear reactor, [2707](#)
- Cap, tube, [587](#)
- Capacitance, [346](#)
- Capacitive level switch, [679](#)
- Capacitively-coupled voltage transformer (CCVT), [1948](#)
- Capacitor, [346](#)
- Capacitor-start AC induction motor, [724](#)
- Capacity tank, [1267](#), [2328](#)
- Capillary tube, [1372](#), [1375](#), [1649](#)
- Captive flow, [2513](#), [2575](#)
- Captive variable, [2513](#)
- Carrier Sense Multiple Access (CSMA), [1044](#)
- Cascade control strategy, [2383](#), [2502](#)
- Cassier's Magazine, [31](#), [516](#)
- Cathode, [1751](#)
- Cation, [228](#), [1751](#)
- Caustic, [286](#)

- Caustic, strong, 287
Caustic, weak, 287
Cavitation corrosion, 2215, 2229
Cavitation, control valve, 2213
CCVT, capacitively-coupled voltage transformer, 1948
Cell constant, conductivity sensor, 1755
Celsius, 123, 1291
CEMS, 1853
Centigrade, 1291
Centrifugal force, 1712
Centrifuge, 2708
Centripetal force, 1712
cgs, 72
Chain, 120
Chain reaction, nuclear, 2707
Channel arbitration, 1040
Channel, FOUNDATION Fieldbus AI block parameter, 1196
Characteristic impedance, 476
Characteristic, control valve, 2190
Characterized ball valve, 2093
Characterizing process dynamics, 2408, 2442
Charging current, power line, 2012
Charles's Law, 198
Chart recorder, 514, 2456
Chemical seal, 1372
Chemical versus nuclear reaction, 265
Chemiluminescence, 1847
Chemistry, 226
Chernobyl nuclear reactor accident, 2423
Choked flow, 2220
Chopper valve, 2682
Chord, ultrasonic flowmeter, 1692
Chromatogram, 1787
Chromatograph, dual-column, 1801
Chromatography, 1784
CIP, 575, 1371
Cippoletti weir, 1653, 2064
Circuit breaker, 1922
Circuit, electric, 299
Circular chart recorder, 514, 517
Cistern manometer, 1310
Clamp-on ammeter, 341
Clamp-on milliammeter, 898
Clamp-on ultrasonic flowmeter, 1694
Class I filled system, 1513
Class II filled system, 162, 1515
Class III filled system, 1514
Class V filled system, 1513
Class, hazardous area, 2599
Classified location, 2598
Clean-In-Place, 575, 1371
Clutch mechanism, electric valve actuator, 2123
Coherent light, 1834
Cold junction compensation, 1540
Cold junction, thermocouple, 1531, 1540
Collision, 1044
Collision domain, Ethernet, 1064
Colloid, 229
Column, chromatograph, 1784
Combination electrode, 1768
Combustion, 273
Common logarithm, 1210
Common-mode rejection, 1345
Common-mode voltage, 622, 1003, 1054
Communicator, HART, 1095
Commutating diode, 761, 2492
Commutator, DC motor, 2237
Compel Data (CD) token, Fieldbus, 1159
Compensating leg, 1431
Compensating probe, capacitive level instrument, 1496
Complex number, 371, 377
Complex plane, 375
Composition probe, capacitive level instrument, 1496
Compositional chemical formula, 235, 270
Compound, 228
Compressed air dryer, 2643
Compressibility, 198
Compression fitting, 1349
Compression terminal, 599
Compression terminal "crimping" tool, 601
Computer, mainframe, 989
Concentration cell, 1763
Condensate boot, 1451
Conductance, 318
Conduction, heat, 1510
Conduction, heat transfer, 128
Conductivity cell, 1753
Conductivity sensor, 1753

- Conductor, electrical, 298
- Confidence, of scientific data, 1177
- Confined space, 1856
- Conical-entrance orifice plate, 1617
- Connector, tube, 585
- Conservation of Electric Charge, 205, 316
- Conservation of electric Charge, 323
- Conservation of Energy, 73, 75, 89, 99, 172, 208, 226, 248, 278, 317, 319, 452, 1220, 1544, 1588, 2209, 2251, 2415, 2530
- Conservation of Mass, 73, 205, 226, 265, 322, 1588, 1708, 2415, 2530
- Constant of proportionality, 1593
- Constraint, hard versus soft, 2584
- Contact, 465, 728, 762, 766, 767, 794, 810, 830, 1926
- Continental Code, 1024
- Continuous Emissions Monitoring System (CEMS), 1853
- Control algorithm, 2272
- Control characters, 989
- Control valve, 2079
- Control valve noise, 2179, 2186, 2188, 2222
- Controlled rectifier, 2241
- Controller, 499
- Controller action, direct vs. reverse, 508, 2278, 2585
- Controller gain, 2279
- Convection, 1510
- Convection, heat transfer, 130
- Conventional flow, 664
- Converter, 499
- Coolants, 138
- Cooling tower, evaporative, 168
- Cooling tower, forced-draft, 169
- Cooling tower, induced-draft, 169
- Coordination, protective relay, 2003
- Coplanar DP sensor, 1332
- Coriolis force, 1713
- Coriolis mass flowmeter, 1714
- Corner taps (orifice plate), 1619
- Coulomb, 293, 299
- Count value, ADC, 977, 992
- Counter instruction, PLC programming, 822
- Counter, Geiger, 1500
- Counterpropagation ultrasonic flowmeter, 1690
- Coupling device, Fieldbus, 2380
- Coupling device, fieldbus, 1145
- Covalent substance, 282
- cps, 1291
- Crest, weir, 1654
- Crimp terminal, 599
- Crimping tool, 601
- Critical damping, 430
- Critical flow, 2220
- Critical flow nozzle, 2221
- Critical pressure, 153
- Critical speed, rotating machine, 2713
- Critical temperature, 153, 1392
- Crosby pressure safety relief valve, 2666
- Cross product, 93, 304, 1713
- Cross-torquing, 568
- Crossover cable, 1066
- Cryogenic, 1391, 2104
- CS Fieldbus function block, 2573
- CSMA, 1044
- CSMA/BA channel arbitration, 1044
- CSMA/CA, 1045
- CSMA/CD channel arbitration, 1044, 1061
- CT, 736, 1900, 1949
- Current, 298, 299
- Current sinking, 304, 778, 781
- Current sourcing, 304, 778, 781
- Current transformer (CT), 736, 1900, 1917, 1949
- Current transformer (CT) safety, 1964
- Curved arrow notation, showing DC voltage, 308
- Curved manometer, 2060
- Custody transfer, 1441, 1597, 1638, 1666, 1699, 1708, 1725
- Cycles per second, 1291
- Cyclic communication, Fieldbus, 1155, 1159
- DAC, 784, 977, 991
- Dall flow tube, 1629
- Daltons, 262
- Damper, 2095
- Damping, 1265
- Damping, anti-resonance resistor, 431
- Damping, automobile suspension, 431
- Damping, critical, 430
- Damping, digital algorithm, 1267
- Damping, feedback control loop, 431

- Danfoss pressure switch, 671
- Daniel four-beam ultrasonic flowmeter, 1692
- DAQ, 520, 1007
- DAQ input, single-ended, 1013
- Data acquisition module, 520, 1007
- Data Communications Equipment, 1051, 1063
- Data diode, 2737
- Data Terminal Equipment, 1051, 1063
- dB, 484, 1210, 1220
- dBm, 1214, 1222
- dBW, 1215, 1222
- DC, 355
- DC excitation, magnetic flowmeter, 1686
- DC injection, AC motor braking, 2250, 2251
- DCE, 1051, 1063
- DCS, 2374
- DD file, 1093
- DDC, 2361
- DDL, 1181
- de Broglie, Louis, 252
- Dead leg, 575
- Dead time, 519, 996, 1800, 2442, 2463, 2542
- Dead time function, 2543, 2552
- Dead time function used for dynamic compensation, 2543, 2545
- Dead-test unit, 1296
- Deadband setting, flow switch, 687
- Deadband setting, pressure switch, 671
- Deadband setting, temperature switch, 683
- Deadband, control valve response, 2206
- Deadband, integral, 2294, 2450
- Deadband, reset, 2294, 2450
- Deadweight tester, 1296
- Deadweight tester, pneumatic, 1298
- Dean effect, 204
- Decade box, resistance, 1289
- Decibel, 484, 1210, 1220
- Decimal numeration, 978
- Decrement (counter), 822
- Decryption, 2736
- Defense-in-depth, 2597, 2725
- Degrees API, 180
- Degrees Balling, 181
- Degrees Bark, 181
- Degrees Baumé, 180
- Degrees Brix, 181
- Degrees Oleo, 181
- Degrees Plato, 181
- Degrees Soxhlet, 181
- Degrees Twaddell, 180
- Delta model DSC-1280 DDC controller, 2364
- Denial-of-service attack, 2722, 2733
- Density, influence on hydrostatic level measurement accuracy, 1419
- Dependability, 2610, 2611, 2685
- Dependability, versus security, 2670
- Dependent current source, 886
- Derivative control, 2295
- Derivative control action, 2299
- Derivative notation, calculus, 1734, 1873
- Derivative, calculus, 8, 12
- Derived unit, 72
- Desiccant, 698
- desiccant, 2643
- Destination host unreachable, error message, 1079
- Desuperheater, 2519
- Determinism, network, 1156, 1236
- Determinism, network communication, 1045
- Deviation alarm, 2386
- Device Description file, 1093
- Device Description Language, 1181
- Device Join Key, WirelessHART, 1246
- DeviceNet, 1044
- Dew point, 699
- DHCP, 1076
- Diaphragm, 932, 1314
- Diaphragm valve, 691, 2080
- Diaphragm, isolating, 576, 1324, 1326, 1329, 1371
- Dictionary attack, password, 2727
- Dielectric constant, 1475
- Dielectric constant, influence on radar level measurement accuracy, 1480
- Dielectric heating, RF cable, 1216
- Dietary Calorie, 125
- Differential, 2295
- Differential capacitance pressure sensor, 1326
- Differential current protection, 2005
- Differential equation, 27, 2431

- Differential measurement mode on an oscilloscope, 1203
- Differential notation, calculus, 1734
- Differential pressure, 186, 1312
- Differential pressure switch, 672
- Differential setting, flow switch, 687
- Differential setting, pressure switch, 671
- Differential setting, temperature switch, 684
- Differential temperature sensing circuit, 334
- Differential voltage protection, 2005
- Differential voltage signal, 622
- Differential voltmeter, 339
- Differential, calculus, 8
- Differentiation, applied to capacitive voltage and current, 406, 407
- Differentiator circuit, 37
- Diffraction grating, 1812
- Digital damping, 1267
- Digital multimeter, 520, 1289, 2803
- Digital-to-analog converter, 784, 977, 991
- Dimensional analysis, 61, 71, 90, 129, 179, 192, 202, 206, 209, 312
- DIN rail, "top hat", 603
- DIN rail, G, 603
- Diode, commutating, 761, 2492
- Diode, in current loop circuit, 899
- Dip tube, 1424
- Dipole antenna, half-wave, 491, 1208
- Direct addressing, 801
- Direct current, 355
- Direct digital control (DDC), 2361
- Direct valve actuator, 2128
- Direct valve body, 2128
- Direct-acting controller, 508, 2278, 2585
- Direct-acting pneumatic relay, 933
- Direct-acting transmitter, 539
- Direct-acting valve body, 2080
- Discharge coefficient, 1638
- Disconnect, electrical power system, 1922
- Discrete, 653, 689, 747, 1280
- Discrete control valve, 2079
- Disk valve, 691, 2092
- Dispersive chemical analyzer, 1812
- Displacement, 191
- Displacer, 1450
- Displacer level instrument, 1449
- Displayed chemical formula, 233
- Dissociation, 1752
- Dissolved oxygen measurement, 634
- Distech model EC-RTU-L DDC controller, 2366
- Distech model ECP-410 DDC controller, 2365, 2494
- Distillation, 133, 1794
- Distortion, valve performance, 2196
- Distributed control system (DCS), 2374
- Distribution power line, 1898
- Diverting valve, 2088, 2161
- Division, hazardous area, 2599
- DIX Ethernet, 1061
- DMM, 520, 996, 1289, 2803
- DNR, 1084
- DNS, 1084
- Documenting calibrator, 1263
- Domain Name Resolver, 1084
- Domain Name Server, 1084
- Domain Name System, 1084
- Domain, frequency, 1880
- Domain, time, 1880
- Doping optical glass, 637
- Doppler effect, 1687
- Doppler ultrasonic flowmeter, 1687
- Dot product, 94
- Double-helical gear, 107
- Double-ported globe valve, 2084
- Double-precision floating-point number, 983
- DP cell, 953
- DP transmitter calibration fitting, 1367
- DPDT switch contacts, 750
- Drag disk, 1628
- Drain hole, orifice plate, 1614
- Drift, 1260
- Drift test, analyzer, 1855
- Drive, electric motor, 84, 2236
- Droop, 2286, 2410
- Drop, 1062, 1141
- Drum, computer virus, 2716
- Drum instruction, Koyo PLC programming, 841
- Drum sequencer, 838
- Dry calibration, 1453
- Dry contact, 663
- Dry leg, 1432
- Dry-block temperature calibrator, 1295

- Dryer, compressed air, 2643
Dryer, instrument air, 698
Dryseal pipe threads, 572
DSO, 996
DTE, 1051, 1063
Dual-column chromatograph, 1801
Dump solenoid, 2683
Duplex, 1040
Dynamic braking, AC motor, 2250, 2252
Dynamic compensation, 2539, 2543
Dynamic friction, 2204
Dynamic Host Configuration Protocol, 1076
Dynamic luminescence quenching, 634
Dynode, 1843
- EBCDIC, 989
Eccentric disk valve, 2092
Eccentric orifice plate, 1612
Echo Request, ICMP, 1074
Eductor, 216
Effective Isotropic Radiated Power (EIRP), 1222
Effective Radiated Power (ERP), 1222
EIA/TIA-232 serial communication, 1050, 1066, 1074, 1089, 1106, 2737
EIA/TIA-422 serial communication, 1054, 2737
EIA/TIA-485 serial communication, 1054, 1066, 1106, 2737
Einstein, Albert, 73, 250, 1805
EIRP, 1222
EIV, 2682
Ejector, 216
Electric circuit, 299
Electric motor valve actuator, 2122
Electrical conductor, 298
Electrical heat tracing, 1389
Electrical insulator, 298
Electrical interlock, reversing motor starter, 742
Electrodeless conductivity cell, 1759
Electrolysis, 273
Electromagnetic force, 230
Electromagnetic induction, 1675
Electromagnetic wave, 488
Electromagnetism, 340
Electron, 228
Electron capture detector, GC, 1789
Electron flow, 664
Electron orbital, 241, 242
Electron shell, 242
Electron subshell filling order, 247
Electronic manometer, 1300
Element, 228
Elevated voltage signal, 1011
Emergency isolation valve (EIV), 2682
Emergency Shutdown (ESD) system, 2669
Emerson AMS software, 975, 2071, 2156
Emerson DeltaV distributed control system (DCS), 975, 1183, 2071, 2376
Emerson model 1420 WirelessHART gateway, 1240
Emerson model 375 HART communicator, 1095
Emerson Ovation distributed control system (DCS), 2376
Emerson Smart Wireless Gateway, 1114, 1240
Emerson THUM WirelessHART adapter, 1238
Emerson, Ralph Waldo, 2744
Emission spectroscopy, 253
Emissivity, thermal, 1570
Emittance, thermal, 1570
Emulsion, 229
Encryption, 2736
Endothermic, 273, 278, 2173
Endress+Hauser magnetic flowmeter, 1683
Energy, 75
Energy balance, 73, 2415, 2530
Energy control procedure, 86
Energy in chemical reactions, 273, 278
Energy loss, flowmeter, 1632, 1742
Energy, in chemical bonds, 248
Energy, kinetic, 88
Energy, potential, 80, 278
Engine, internal combustion, 53
Engineering units, 1196
Enriched uranium, 2707
Enthalpy, 143, 144, 148
Enthalpy of formation, 278
Entropy, 75, 226
Epistemology, 1282
ePiX C++ library, 417
Equal percentage valve characterization, 2198
Equipotential, 317, 326, 617
Equivalent circuits, series and parallel AC, 361
Erosion, 2224

- ERP, 1222
 Error, controller, 2278, 2291
 ESD, 2669
 Ethernet, 1028, 1061, 1106, 1112, 1140, 1156, 1158, 1189, 1240, 2359, 2373, 2375, 2500, 2729, 2737
 Euler's relation, 373, 390
 Euler, Leonhard, 208, 373
 European alpha value, RTD, 1519
 Evaporative cooling tower, 168
 Examine if closed (XIC), 814
 Examine if open (XIO), 814
 Exchanger, heat, 130
 Excitation source, for bridge circuit, 327
 Excitation wires, 4-wire RTD circuit, 1521
 Exergy, 75
 Exothermic, 273, 278, 2173
 Expanded structural formula, 233
 Explosion-proof enclosure, 2605
 Extended floating-point number, 983
 Extension grade thermocouple wire, 1553
 External fault, power system protection zone, 2010
 External reset, integral control, 2338, 2569
 Externally-loaded pressure regulator, 2120

 Fabry-Perot interferometry, 633
 Fade margin, RF, 1228
 Fade, RF, 1227
 Fahrenheit, 123
 Fail closed, 2130
 Fail locked, 2130
 Fail open, 2130
 Fail-safe mode for a control valve, 2127
 Fail-safe mode for split-ranged control valves, 2172
 Failure rate, λ , 2628
 Failures In Time (FIT), 2628
 False positive, 2653
 False state, PLC programming, 804
 Farad, 346
 Faraday's Law of Electromagnetic Induction, 362, 681, 723, 1975
 Fast Fourier Transform (FFT), 1880
 Fault tolerance, 2671

 FCC – Federal Communications Commission, 1236
 Federal Communications Commission, 1236
 Feedback control system, 2270, 2521
 Feedforward control strategy, 2523
 Feedforward with trim, 2528
 Ferrule, electrical connection, 593, 599
 FF (FOUNDATION Fieldbus), 1133, 2379
 FFT algorithm, 1880
 Fiber optic cable, 627
 Fiducial pulse, radar, 1482
 Fieldbus, 507, 626, 974, 1096, 2379
 Fieldbus coupling device, 2380
 Fieldbus Foundation, 2379
 FIFO shift register, 1267, 2552
 Fill fluid, 576, 1324, 1326, 1328, 1329, 1335, 1363, 1370, 1513
 Fillage, 1410, 1468
 Filled bulb, 174, 1513
 Filled impulse line, 1381
 Filtering, negative (spectroscopy), 1829
 Filtering, positive (spectroscopy), 1829
 Final Control Element, 499
 Fire signals, 987
 Fire triangle, 2601
 Firewall, computer, 2735
 First Law of Motion, 74
 First-order differential equation, 2431
 First-order lag, 2431
 Fiscal measurement, 1638
 Fisher “Level-Trol” displacer instrument, 1451, 1459, 1466
 Fisher “Whisper” low-noise trim, 2223
 Fisher AC² analog electronic controller, 2349
 Fisher E-body control valve, 3167
 Fisher E-plug control valve, 2094
 Fisher Micro-Flat Cavitation trim, 2214
 Fisher model 1098EGR pilot-operated pressure regulator, 2129
 Fisher model 2625 volume boosting relay, 2141
 Fisher model 3582 valve positioner, 2143
 Fisher model 546 I/P transducer, 961, 2110
 Fisher model 760 pressure relief valve, 2668
 Fisher model 846 I/P transducer, 2110
 Fisher model DVC6000 valve positioner, 2144, 2154

- Fisher model DVC6200 valve positioner, [2144](#)
- Fisher MultiTrol pneumatic controller, [2329](#)
- Fisher Provox distributed control system (DCS), [2376](#)
- Fisher ROC digital controllers, [2370](#)
- Fisher ValveLink software, [2156](#)
- Fisher-Rosemount model 846 I/P transducer, [966](#)
- Fission chamber, [1501](#)
- Fission, nuclear, [2421](#), [2698](#), [2707](#)
- FIT, [2628](#)
- Five-point calibration, [1278](#)
- Five-valve manifold, [1362](#)
- FIX/Intellution HMI software, [854](#)
- Fixed Programming Language (FPL), [787](#)
- Fixed-point variable, [981](#), [986](#)
- Flame ionization detector, GC, [1790](#)
- Flame photometric detector, GC, [1789](#)
- Flame safety system, [840](#), [2687](#)
- Flange taps (orifice plate), [1618](#)
- Flange, pipe, [564](#)
- Flapper, [921](#)
- Flare gas flow measurement, [1698](#)
- Flashing, [2209](#)
- Flashover, power line insulator, [1942](#)
- Flexure, [1258](#)
- Float level measurement, [1410](#)
- Floating control action, [2291](#), [2298](#)
- Floating voltage signal, [1010](#)
- Floating-point number, [982](#)
- Floating-point variable, [2389](#)
- Flooded displacer, [1462](#)
- Flow conditioner, [1634](#)
- Flow control (serial data communication), [1038](#), [1051](#)
- Flow integrator, [42](#)
- Flow programming, chromatograph, [1800](#)
- Flow prover, [1301](#), [1666](#)
- Flow switch, [686](#)
- Flow totalizer, [42](#)
- Flow tube, [1629](#)
- Flow-straightening vanes, [1634](#)
- Fluid, [170](#)
- Fluidized sand, [1293](#)
- Fluke brand multimeters, [2803](#)
- Fluke Documenting Process Calibrator (DPC) instruments, [1264](#)
- Fluke DPCTrack2 calibration management software, [1264](#)
- Fluke model 525A temperature calibrator, [1289](#)
- Fluke model 744 calibrator, [1559](#)
- Fluke model 744 Documenting Process Calibrator (DPC), [1263](#)
- Fluke model 771 clamp-on milliammeter, [898](#)
- Fluke model 801 differential voltmeter, [339](#)
- Fluke SV225 stray voltage adapter, [2806](#)
- Flume, [1657](#), [2063](#)
- Fluorescence, [634](#), [1838](#)
- Fluorescence quenching, [634](#)
- Flywheel, [98](#)
- Foam, [229](#)
- Follower, [2150](#)
- Force balance system, [928](#), [947](#), [949](#), [953](#), [963](#), [1339](#), [2319](#)
- Force-balance valve positioner, [2147](#)
- Fork terminal, [600](#)
- Form-A switch contact, [654](#), [750](#)
- Form-B switch contact, [654](#), [750](#)
- Form-C switch contact, [660](#), [661](#), [750](#)
- Formula weight, [263](#), [2516](#)
- Fortescue, Charles Legeyt, [458](#)
- FOUNDATION Fieldbus, [1096](#), [1133](#), [2537](#)
- FOUNDATION Fieldbus (FF), [2379](#)
- FOUNDATION Fieldbus H1, [1028](#), [1140](#)
- FOUNDATION Fieldbus H2, [1140](#)
- FOUNDATION Fieldbus HSE, [1140](#)
- Four-wire RTD circuit, [1521](#)
- Fourier series, [458](#), [1877](#), [2244](#), [2250](#)
- Fourier transform, [1216](#)
- Fourier, Jean Baptiste Joseph, [458](#), [1877](#), [2244](#), [2250](#), [2258](#)
- Foxboro (Invensys) I/A distributed control system (DCS), [2376](#)
- Foxboro FOXNET process data network, [2376](#)
- Foxboro INTERSPEC process data network, [2376](#)
- Foxboro magnetic flowtube, [1685](#)
- Foxboro model 13 differential pressure transmitter, [953](#)
- Foxboro model 130 pneumatic controller, [925](#), [2334](#)
- Foxboro model 14 flow totalizer, [43](#)
- Foxboro model 43AP pneumatic controller, [2332](#)

- Foxboro model 557 pneumatic square root extractor, 1602
- Foxboro model 62H analog electronic controller, 2349
- Foxboro model E69 I/P transducer, 956
- Foxboro model E69F I/P transducer, 2110
- Foxboro model IDP10 differential pressure transmitter, 1325, 1343, 1434
- Foxboro SPEC 200 analog electronic control system, 2352, 2376, 2387, 2641
- Foxboro SPECTRUM distributed control system (DCS), 2376
- Fractionation, 133, 1794
- Fragment, IP packet, 1072
- Frame check sequence, 1038
- Fraunhofer lines, 1814
- Fraunhofer, Joseph von, 1814
- Freezing point of water, 1291
- Frequency domain, 1880
- Frequency Shift Keying, 1029
- Frequency shift keying, 1097
- Fresnel zone, 1230
- Fribance, Austin E., 2
- Friction, static versus dynamic, 2204
- FSK, 1029, 1097
- Fuel cell oxygen sensor, 1860
- Fugitive emissions, 2098, 2663
- Fulcrum, 99
- Fulcrum, torque tube, 1457
- Full Variability Language (FVL), 787
- Full-active bridge circuit, 335
- Full-duplex, 1040
- Full-flow taps (orifice plates), 1619
- Fully developed turbulent flow, 203
- FUN, Fieldbus addressing, 1157
- Function block programming, 2357, 2380
- Function, inverse, 25, 2055, 2763
- Function, piecewise, 2068
- Functional diagram, 540
- Functions, transfer, 411
- Fundamental frequency, 487, 491, 1209, 1877
- Fundamental Theorem of Calculus, 26
- Fusion splicing, optical fiber, 642

- G DIN rail, 603
- G, unit of acceleration, 1875
- Gain and bias function block, 2535, 2548
- Gain margin, 2440
- Gain, amplifier, 1210
- Gain, controller, 2279
- Galilei, Galileo, 74
- Galvanic corrosion, valve stem packing, 2104
- Galvanic isolation, 736
- Galvanometer, 328, 338
- Gamma ray radiation, 1497
- Gas, 170
- Gas centrifuge, 2708
- Gas centrifuge cascade, 2710
- Gas centrifuge stage, 2710
- Gas expansion factor, 1639
- Gas Filter Correlation spectroscopy, 1829
- Gas Laws, 198
- Gas phase effect, radar level instrument, 1476, 1481
- Gas, calibration, 264, 1262, 1304, 1835, 1845, 1859
- Gas, span, 1304, 1835, 1845, 1859
- Gas, zero, 1835, 1845
- Gate valve, 691, 2080
- Gate, logic, 747
- Gauge line, 579, 1349
- Gauge pressure, 186, 1312
- Gauge tube, 579, 1349
- Gauge, tube fitting, 581
- Gay-Lussac's Law, 198
- GE Multilin model 369 protective relay, 737
- GE Multilin model 745 protective relay, 2008
- GE Series One PLC, 774
- Gear ratio, 96
- Gear set, 106
- Gear, bevel, 111
- Gear, double-helical, 107
- Gear, herringbone, 107
- Gear, hypoid, 112
- Gear, miter, 111
- Gear, planetary, 109
- Gear, single-helical, 107
- Gear, spur, 106, 2124
- Gear, worm, 113, 2124
- Geiger counter, 1500
- Geiger-Muller tube, 1500

- General Electric “Magneblast” circuit breaker, 1928
- General Electric model 121AC time-overcurrent protective relay, 1989
- Generator, 300
- Generator synchronization, 1914
- Gentile flow tube, 1629
- Gerlach scale, 181
- GFC spectroscopy, 1829
- Gibbs free energy, 226
- Gibbs’ phase rule, 159
- Gilbert, 340
- Global Positioning System (GPS), 385
- Globe valve, 2080
- Google, Internet search engine, 1086
- GPS satellite system, 385
- Graded index optical fiber, 640
- Graphic User Interface, 1076
- Graphite rupture disk, 2658
- Gravitational potential energy, 80
- Groth model 1208 pressure/vacuum safety valve, 2662
- Ground, 322
- Ground fault, 2005
- Ground loop, 622, 1015, 1200, 1964
- Ground-referenced voltage signal, 1009
- Grounding, magnetic flowmeters, 1682
- Group, hazardous area, 2600
- GUI, 1076
- Guide ring, pressure safety relief valve, 2666
- Guided wave radar, 483, 1473

- Hagen-Poiseuille equation, 207, 1648, 3164
- Half-duplex, 1040
- Half-life, of a radioactive substance, 1498
- Half-wave dipole antenna, 491, 1208
- Hall Effect sensor, 1337, 2144
- Hand controller, 2540
- Hand switch, 659
- Hand valve actuator, 2126
- Handshaking (serial data communication), 1038, 1051
- Handwheel, control valve, 2127
- Handwheel, valve, 2111
- Hard alarm, 2386
- Hard constraint, 2584
- Hard determinism, network, 1156
- Hard override, 2582
- Hardware compensation, thermocouple, 1547
- Harmonic frequency, 487, 491, 1209, 1877, 2258, 2808
- Harmonic restraint, 87 relay, 2019
- Harrier jet airplane, 2420
- HART, 1029
- HART analog-digital hybrid, 627, 974, 1091, 1269
- HART communicator, 1095
- HART multidrop mode, 1104
- Hazardous location, 2598
- Head (fluid), 208
- Heat, 1510
- Heat exchanger, 130, 2266
- Heat of formation, 278
- Heat of reaction, 274
- Heat pump, 126
- Heat tape, 1389
- Heat tracing, 1387
- Heat transfer by conduction, 128
- Heat transfer by convection, 130
- Heat transfer by radiation, 127
- Heater, overload, 731
- Heisenberg, Werner, 252
- Helical bourdon tube, 1300, 1315
- Hello World
 - program, 857
- Henry, 350
- Herringbone gear, 107
- Herschel, Clemens, 1629
- Hertz, 274, 1291
- Hess’s Law, 279
- HEU, 2707
- Heuristic PID tuning example, 2480
- High-limit function, 2565
- High-performance butterfly valve, 2092
- High-select function, 2563
- High-side transistor switch, 665
- High-speed pulse test, 2651
- High-voltage circuit breaker, 1934
- HMI panel, 849, 984
- Hold-off distance, radar, 1483
- Home position, CNC machine, 661
- Honed meter run, 1641

- Honeywell Experion PKS distributed control system (DCS), 2376
- Honeywell model UDC3000 controller, 2358
- Honeywell Radiamatic, 1564
- Honeywell TDC2000 distributed control system (DCS), 2375
- Hooke's Law, 91, 2133, 2142, 2157
- Hot standby, 2645
- Hot terminal, 454
- Hot-swappable PLC I/O, 775
- Hot-tapping, 1738
- Hot-wire anemometer, 1726
- HTTP, 1090
- HTTPS, 1090
- Huddling chamber, pressure safety valve, 2665
- Human-Machine Interface panel, 849, 984
- HVAC, 1520, 2095
- Hydration, pH electrode, 1770
- Hydraulic, 173
- Hydraulic lift, 172
- Hydraulic load cell, 1494
- Hydraulic valve actuator, 2117
- Hydrogen ion, 282, 284, 1752
- Hydrometer, 193
- Hydronium ion, 282, 284, 1752
- Hydrostatic pressure, 178
- Hydroxyl ion, 282, 284, 1752
- Hyperterminal, 1033
- Hypoid gear, 112
- Hysteresis, 1278, 2447
- Hysteresis error, 1258

- I.S. system, 2606
- I/O, 766, 775
- I/O, analog (PLC), 784
- I/O, discrete (PLC), 777
- I/O, hot-swappable, 775
- I/O, network (PLC), 786
- I/O, remote, 776
- I/P transducer, 902, 920, 2110
- IANA, 1076
- ICANN, 1076, 1084
- Ice cube relay, 751
- Ice point, thermocouple, 1541, 1547
- ICMP, 1074
- ICS-Triplex TMR safety control system, 2682
- Ideal Gas Law, 197, 1514, 1705, 2771
- Ideal PID equation, 2316, 2347, 2394, 2427
- Identifier, Fieldbus device, 1158
- IEC standard 61131-3 (PLC programming languages), 787
- IEC standard 62591 (WirelessHART field instrument communications protocol), 1207, 1232
- IEEE codes for power systems, 1992
- IEEE standard C57.12.00-2010 for transformers, 2022
- IFC spectroscopy, 1829
- Ifconfig, utility program, 1085
- Imaginary number, 371, 373
- Immunity, noise, 1022
- Impedance, 360, 392, 406, 407, 411
- Impedance, characteristic, 476
- Impedance, percentage, 370
- Impedance, surge, 476
- Impedance, transformer, 370
- Impeller-turbine mass flowmeter, 1710
- Impulse line, 579, 1349
- Impulse tube, 579, 1349, 1363
- In situ measurement, 1851
- Inches of mercury, 178
- Inches of water column, 178
- Inclined manometer, 183, 1310
- Inclined plane, 104
- Increment (counter), 822
- Index of refraction, 636
- Indicator, 511
- Inductance, 350
- Induction cup protective relay mechanism, 2041
- Induction disk protective relay mechanism, 1989, 1998
- Induction motor, 722, 2245, 2813
- Induction, mutual, 364
- Induction, self, 362
- Inductor, 350
- Inferred variable, 1176, 1307, 1441, 1640, 2056, 2058, 2531, 2548
- Information Technology (IT), 2705
- Infrared thermocouple, 1565
- Inherent valve characteristic, 2190, 2196
- Inrush current, 723
- Installed valve characteristic, 2190, 2196, 2507

- Instantaneous overcurrent protection, 1996
- Instrument air systems, 2643
- Instrument Protective Function (IPF), 2669
- Instrument tray cable (ITC), 1146
- Instrument tube bundle, 1387
- Instrumentation amplifier, 1006
- Insulation, R-value, 129
- Insulator, electrical, 298
- Integer number, 978
- Integer variable, 2389
- Integral control action, 2291, 2298
- Integral deadband, 2294, 2450
- Integral orifice plate, 1622, 1644
- Integral windup, 2294, 2385, 2411
- Integral windup, limit controls, 2568
- Integral windup, override controls, 2584
- Integral, calculus, 8, 20
- Integrating process, 2412
- Integration, applied to RMS waveform value, 357
- Integrator, 42
- Interacting PID equation, 2317, 2348
- Interactive zero and span adjustments, 1254, 1300
- INTERBUS-S, 1058
- Interface level measurement, 1403
- Interference Filter Correlation spectroscopy, 1829
- Interlock, reversing motor starter, 742
- Internal combustion engine, 53
- Internal fault, power system protection zone, 2010
- International Morse code, 988
- International Practical Temperature Scale (ITS-90), 1293, 1538, 1556, 1559
- Internet Control Message Protocol, 1074
- Internet Protocol, 1071
- Interoperability, Fieldbus devices, 1181
- Interposing, 761
- Intrinsic safety, 1341
- Intrinsic safety barrier circuit, 2607
- Intrinsic standard, 1282, 1285
- Intrinsically safe system, 2606
- Inverse function, 25, 2055, 2763
- Inverse square law of radiation, 1227, 1565, 1567
- Inverted classroom, 3178, 3182
- Inviscid flow, 201
- Ion, 228
- Ion-selective membrane, 2074
- Ionic substance, 282
- Ionization, 1752
- Ionization chamber for detecting nuclear radiation, 1500
- IP, 1071
- Ipconfig, utility program, 1084
- IPF, 2669
- IPv4, 1073
- IPv6, 1082
- ISA 84, 787
- ISA PID equation, 2316, 2394, 2427
- ISA100.11a, 1043
- ISEL Fieldbus function block, 2573
- ISM 2.4 GHz radio band, 1236
- Isolating diaphragm, 576, 1324, 1326, 1329, 1371
- Isolation, galvanic, 736
- Isomer, 228
- Isopotential point, pH, 1783
- Isothermal terminal block, 598, 1535
- Isotope, 228, 2707
- Isotopes, 237
- Isotropic antenna, 1220
- ITC, 1146
- ITS-90, 1293, 1538, 1556, 1559
- ITT safety valve, 2689
- Jabber, 1045
- Jabber latch, 1045
- Jacket, reactor vessel, 507
- Jam packing, valve, 2100
- Join Key, WirelessHART, 1246
- Josephson junction array voltage standard, 1283
- Joule, 293
- Joules' Law, 312
- Jumper wire, 2811
- K-Tek brand magnetic float level indicator, 1415
- Kallen, Howard P., 2
- KCL, 322, 2004
- Kelvin, 123
- Kelvin resistance measurement, 1521, 1756
- Kermit, 1033
- Key, cryptographic, 2736
- Keyphasor, 1887
- Kinematic viscosity, 199

- Kinetic energy, 88
- Kirchhoff's Current Law, 322, 2004
- Kirchhoff's Voltage Law, 320, 414
- Knife-edge bearing, 1457
- Knockout drum, 1451
- Koyo CLICK PLC, 772, 803, 855
- Koyo DL06 PLC, 773
- KVL, 320, 414

- L2F optical flowmeter, 1695
- L-Type, FOUNDATION Fieldbus AI block parameter, 1196
- Ladder Diagram, 765
- Ladder Diagram programming, 757, 802
- Lag time, 2434, 2546
- Lag time function, 2550
- Lambert-Beer Law, 1810
- Laminar flow, 204, 207, 3164
- Laminar flowmeter, 1648
- LAN, computer, 2735
- Lantern ring, 2100
- Laplace transform, 417, 1216
- Lapping valve plugs and seats, 2085
- LAS, 1155
- Laser, 1834
- Laser-two-focus optical flowmeter, 1695
- Latent heat, 144
- Latent heat of fusion, 145
- Latent heat of vaporization, 145
- Law of Continuity (fluids), 205, 1589, 1661, 2209
- Law of Intermediate Metals, thermocouple circuits, 1544
- LD, 757, 802
- Le Châtelier, Henry Louis, 2746
- Lead function used for dynamic compensation, 2550
- Lead time function, 2550
- Lead/lag function, analog circuit, 2553
- Lead/lag function, digital implementation, 2561
- LED, 1834
- LEL, 1861, 2602, 2696
- Lenz's Law, 624, 1998, 2251
- LEU, 2707
- Level gauge, 1404
- Level switch, 673
- Lever, 99
- Lift pressure, safety valve, 2658, 2665
- Limit switch, 661
- Limited Variability Language (LVL), 787
- Limiting case, 198, 206
- Line current, 448
- Line of action, 93
- Line of Sight (RF), 1230
- Line pressure effect, pressure transmitter, 1345
- Line reactor, 2244, 2260
- Line voltage, 448
- Linear time-invariant (LTI) system, 437
- Linear valve characterization, 2198
- Linearity error, 1257
- Linearization, 2062
- Link Active Scheduler, 1155
- Link Active Scheduler, FOUNDATION Fieldbus, 2380
- Lipták, Béla, 2, 1428, 1661, 2188, 3162
- Liquid, 170
- Liquid interface detection with radar, 1478
- Liquid valve sizing equation, 2179
- Liquid-tight flexible conduit, 607
- Live List, Fieldbus, 1159
- Live zero, 862, 865, 882, 1253
- Live-load packing, valve, 2101
- Lo-Loss flow tube, 1629
- Load, 300, 2271, 2284
- Load cell, 332, 1490
- Load cell, hydraulic, 1494
- Load encroachment blocking, protective relay, 2048
- Load line, 2193
- Load versus source, 306
- Load, process, 2522
- Loading pressure, 2120
- Lock-out, tag-out, 85, 1499, 1925
- Lockout auxiliary relay, 2049
- Logarithm, common, 1210
- Logic gate, 747
- Logic solver, 2680
- Logix5000 PLC software, Rockwell, 837
- LonWorks, 2365
- Loop calibrator, 906
- Loop diagram, 537
- Loop sheet, 537
- Loop, computer programming, 2388

- Loop-powered transmitter, [311](#), [891](#)
- Loopback address, [1076](#)
- Louvre, [2095](#)
- Low flow cutoff, vortex flow transmitter, [1673](#)
- Low-limit function, [2565](#)
- Low-select function, [2563](#)
- Low-side transistor switch, [665](#)
- Lower explosive limit (LEL), [1861](#), [2602](#), [2696](#)
- Lower range value, [499](#), [1253](#), [1270](#), [1300](#)
- LRV, [499](#), [1253](#), [1270](#), [1300](#)
- LTI system, [437](#)
- Lubricator, valve packing, [2103](#), [2207](#)
- Luft detector, [1822](#)
- Lug, electrical connection, [593](#), [599](#)
- Luminiferous ether, [1061](#)
- Lusser's Law, [2618](#), [2686](#)
- LVDT, [1337](#), [2125](#), [2146](#)

- MAC address, Ethernet, [1062](#), [1073](#), [1158](#)
- Machine, simple, [99](#)
- MacNeil, Blair, [2495](#)
- Macrocycle, [1157](#)
- Madelung rule, [247](#)
- Magnetic flowmeter, [1676](#)
- Magnetic remanence, [1977](#)
- Magnetic shielding, [623](#)
- Magnetostriction, [1416](#), [1486](#)
- Magnetrol liquid level switch, [674](#)
- MagTech level gauge, [1406](#)
- Make-before-break switch contacts, [1967](#)
- Man-in-the-middle attack, [2723](#)
- Manchester encoding, [1028](#)
- Manifold, pressure transmitter, [1359](#), [1362](#)
- Manipulated variable, [499](#), [2268](#)
- Manipulating algebraic equations, [2761](#)
- Manometer, [182](#), [1299](#), [1308](#)
- Manometer, cistern, [1310](#)
- Manometer, inclined, [183](#), [1310](#)
- Manometer, nonlinear, [2060](#)
- Manometer, raised well, [1310](#)
- Manometer, slack tube, [1300](#)
- Manometer, U-tube, [1310](#)
- Manometer, well, [1310](#)
- Mantissa, floating-point number, [983](#)
- Manual loading station, [2540](#), [2567](#)
- Manual mode, [500](#), [2383](#)

- Manual valve actuator, [2126](#)
- Mapping, Modbus, [1114](#)
- Mark, [1027](#), [1029](#)
- Mask, Allen-Bradley SQO sequencer output instruction, [845](#)
- Mask, subnet, [1077](#)
- Masoneilan model 21000 control valve, [2081](#)
- Mass balance, [73](#), [2415](#), [2530](#)
- Mass density, [57](#)
- Master Terminal Unit (MTU), [976](#), [2368](#)
- Master-slave channel arbitration, [1041](#)
- Maximum experimental safe gap (MESG), [2600](#), [2605](#)
- Maximum working pressure, [1348](#)
- Maxon safety valve, [2688](#)
- Maxwell's electromagnetic equations, [489](#)
- Maxwell, James Clerk, [489](#)
- MCC, [733](#)
- MCC room, [611](#), [1924](#)
- Mean life (of a component or system), [2634](#)
- Mean Time Between Failures (MTBF), [2631](#)
- Mean Time To Failure (MTTF), [2631](#)
- Measurement electrode, [1765](#)
- Measurement junction, thermocouple, [1531](#)
- Mechanical advantage, [100](#)
- Mechanical interlock, reversing motor starter, [742](#)
- Median signal select, [2571](#)
- Medium-voltage circuit breaker, [1927](#), [1985](#)
- Memory map, [797](#)
- MEMS, [1335](#)
- Meniscus, [1309](#)
- Mercoid pressure switch, [669](#)
- Mercury, [1399](#)
- Mercury barometer, [1312](#)
- Mercury cell, [1557](#)
- Mercury tilt switch, [669](#), [675](#)
- Merrick weighfeeder, [1732](#)
- Metal fatigue, [1322](#)
- Metal rupture disk, [2657](#)
- Metallic liquid-tight conduit, [608](#)
- Meter run, orifice (honed), [1641](#)
- Metering pump, [2262](#)
- Metrology, [1282](#)
- Mho element, protective relay, [2042](#)
- Micro fuel cell oxygen sensor, [1860](#)

- Micro Motion Coriolis mass flowmeter, [1719](#)
 Micro Motion ELITE model Coriolis mass flowmeter, [1281](#)
 Micromanometer, [184](#)
 Micron, [637](#), [639](#)
 MIE, [2602](#)
 Mil, [1875](#)
 Miller, Richard W., [1638](#)
 Milton-Roy metering pump, [2262](#)
 Minicom, [1033](#)
 Minimum bend radius for optical fiber, [642](#)
 Minimum ignition current ratio (MICR), [2600](#)
 Minimum Ignition Energy, [2602](#)
 Minimum linear flow rate, turbine flowmeter, [1670](#)
 Miter gear, [111](#)
 Mixing valve, [2088](#), [2161](#)
 Mixture, [229](#)
 Mobile phase, [1784](#)
 Modal dispersion, [639](#)
 Modbus, [786](#), [1106](#)
 Modbus 984 addressing, [1113](#)
 Modbus ASCII, [1111](#)
 Modbus mapping, [1114](#)
 Modbus Plus, [1106](#)
 Modbus RTU, [1111](#), [1241](#)
 Modbus TCP, [1090](#), [1112](#), [1241](#)
 Modulus, bulk, [1467](#), [1469](#), [1487](#), [1644](#), [1689](#), [1690](#)
 Molarity, [263](#), [1302](#)
 Mole, [262](#)
 Molecular chemical formula, [233](#)
 Molecular substance, [282](#)
 Molecular weight, [263](#)
 Molecule, [228](#)
 Moment arm, [93](#)
 Moment balance system, [947](#)
 Monochromatic light, [1834](#)
 MooN redundancy notation, [2671](#)
 Moore Industries model IPT I/P transducer, [2110](#)
 Moore Industries model SPA alarm module, [523](#), [685](#)
 Moore Products “Nullmatic” temperature transmitter, [1516](#)
 Moore Products model 353 digital controller, [2356](#), [2359](#)
 Moore Products model 65 pneumatic square root extractor, [1602](#)
 Moore Syncro analog electronic controller, [2350](#)
 Morse code, [988](#), [1024](#)
 Motion balance system, [948](#), [949](#), [957](#), [2321](#)
 Motion control system, [2508](#)
 Motion-balance valve positioner, [2151](#)
 Motional EMF, [1676](#)
 Motor Control Center, [733](#)
 Motor control center (MCC), [611](#), [1924](#)
 Motor drive, [84](#), [2236](#)
 Motor overload protection, [730](#)
 Motor valve actuator, [2122](#)
 MOV, [2122](#)
 MTBF, [2631](#)
 MTS M-Series magnetostrictive float level transmitter, [1487](#)
 MTTF, [2631](#)
 MTU, [976](#), [2368](#)
 Mu metal, [623](#)
 Multi-column chromatograph, [1801](#)
 Multi-mode optical fiber, [639](#)
 Multi-segment characterizer, [2068](#)
 Multi-variable transmitter, [1105](#), [1478](#), [1643](#), [1724](#), [1794](#)
 Multidrop, HART, [1104](#)
 Multilin model 369 protective relay, [737](#)
 Multipath ultrasonic flowmeter, [1692](#)
 Multiplexer, [1007](#)
 Multiplication factor, nuclear fission, [2421](#)
 Multiplying relay, [2514](#)
 Mutual induction, [364](#)
 Mux, [1007](#)
 MV, [499](#)
 MWP, [1348](#)
 NAMUR recommendation NE-43, [2653](#)
 NaN, [982](#)
 Nassau model 8060 loop calibrator, [911](#)
 National Bureau of Standards, [1282](#)
 National Electrical Code (NEC), [1146](#), [2599](#)
 National Fire Protection Association (NFPA), [2599](#)

- National Institute of Standards and Technology, 1282
- Natural convection, 135, 165
- NBS, 1282
- NDE, 717
- NDIR spectroscopy, 1815
- NDUV spectroscopy, 1815
- NDVIS spectroscopy, 1815
- NE, 716
- NEC, 1146, 2599
- Needle valve, 1368, 2082
- Negative feedback, 37, 926, 940, 1387, 2132, 2343, 2417, 2421, 2425, 2439, 2441, 2443, 2490, 2502
- Negative filtering (spectroscopy), 1829
- Negative lag, 2420
- Negative pressure, 189
- Negative self-regulation, 2420
- NEMA 7 enclosure, 2605
- NEMA 8 enclosure, 2605
- Nernst equation, 1763, 1775, 2074
- Netstat, utility program, 1090
- Network ID, WirelessHART, 1246
- Network Manager, WirelessHART, 1235, 1237
- Neutral pH, pure water, 284
- Neutral terminal, 454
- Neutralization, pH, 288
- Neutron, 228
- Neutron backscatter, 1497
- Neutron radiation, 1497
- Newton's Law of Cooling, 2431
- Newton, Isaac, 74
- Nichols, N.B., 2457
- NIST, 1282
- NIST traceability, 1284
- Nitrogen-phosphorus detector, GC, 1789
- Noise floor, 649, 1224–1226, 1229
- Noise immunity, 1022
- Noise, control valve, 2179, 2186, 2188, 2222
- Non-bleeding pneumatic relay, 937
- Non-contact radar, 1473
- Non-dispersive chemical analyzer, 1815
- Non-inertial reference frame, 1712
- Non-metallic liquid-tight conduit, 608
- Non-Newtonian fluid, 200
- Non-retentive instruction, PLC program, 818, 828
- Non-Return-to-Zero, 1027
- Nonincendive circuit, 2606
- Nonlinear manometer, 2060
- NooM redundancy notation, 2671
- Normal energization state of a solenoid, 716
- Normal state of a circuit breaker auxiliary contact, 1987, 1995
- Normal state of a PLC program contact, 759, 788, 804, 808, 814
- Normal state of a relay contact, 749, 755
- Normal state of a switch, 654, 656, 659, 661, 663, 668, 673, 682, 686
- Normal state of a thermal overload switch, 730
- Normal state of a valve, 703, 706, 712
- Normal state of a vibration switch, 1891
- Normal state of an 86 lockout relay, 2050
- Normally de-energized (NDE), 717
- Normally energized (NE), 716
- Not a Number (NaN), 982
- NOx emissions, 2074
- Nozzle, 921
- Nozzle, process vessel, 1407, 1462, 2648
- NPN output switch, 665
- NPT pipe threads, 572
- NRC, Nuclear Regulatory Commission, 1501
- NRZ, 1027
- NSS Labs, 2729
- Nuclear fission reactor, 2421, 2698
- Nuclear radiation, 1497
- Nuclear Regulatory Commission, 1501
- Nuclear versus chemical reaction, 265
- Null modem, 1053, 1066
- Null zone, radar, 1483
- Numeration systems, 978
- NUN, Fieldbus addressing, 1157
- Nupro pressure relief valve, miniature, 2660
- Nyquist Sampling Theorem, 996
- Octal base relay, 751
- OFC, 644
- Off-delay timer, PLC programming, 827
- Ohm, 313
- Ohm's Law, 313
- Ohm's Law analogy for turbulent fluids, 2178

- Ohm, Georg Simon, [313](#)
- Oil bath temperature calibrator, [1293](#)
- Oleo, degrees, [181](#)
- Omega model CL-351A dry-block temperature calibrator, [1295](#)
- Omega model OS-36 infrared thermocouples, [1565](#)
- On-delay timer, PLC programming, [827](#)
- On-off control, [2275](#)
- OOS mode, Fieldbus, [1196](#)
- Open Fiber Control safety system, [644](#)
- Open, electrical fault, [324](#)
- Open-loop test, [2408](#), [2409](#), [2412](#), [2419](#), [2424](#), [2442](#)
- Operating coil, protective relay, [2006](#), [2038](#)
- Optical beam splitter, [1833](#), [1836](#)
- Optical fiber cable, [627](#)
- Opto 22 Optomux network, [1058](#)
- Orbital, electron, [241](#), [242](#)
- Order of magnitude, [942](#)
- Order of operations, [2764](#)
- Orifice meter run, [1641](#)
- Orifice plate, [1609](#), [2057](#)
- Orifice plate, concentric, [1610](#)
- Orifice plate, conical entrance, [1617](#)
- Orifice plate, eccentric, [1612](#)
- Orifice plate, integral, [1622](#), [1644](#)
- Orifice plate, quadrant edge, [1616](#)
- Orifice plate, segmental, [1613](#)
- Orifice plate, square-edged, [1610](#)
- Oscilloscope probe, $\times 10$, [2557](#)
- Oscilloscope, differential measurement mode, [1203](#)
- Out coil, PLC programming, [818](#)
- Out Of Service (OOS) mode, Fieldbus, [1196](#)
- OUT.Scale, FOUNDATION Fieldbus AI block parameter, [1196](#)
- Output limit, PID controller, [2386](#)
- Output tracking, [2384](#)
- Over-damping, [430](#)
- Over-lean condition, [2601](#)
- Over-rich condition, [2601](#)
- Overload “heater”, [731](#)
- Overload protective device, [730](#)
- Overreaching, protective relay, [2027](#)
- Override, hard, [2582](#)
- Override, hard versus soft, [2584](#)
- Override, soft, [2583](#)
- Overtone frequency, [1877](#)
- Oxygen control, burner, [2074](#)
- P&ID, [535](#)
- Packet, IP data, [1071](#)
- Packing lubricator, [2103](#), [2207](#)
- Packing, bellows seal, [2105](#)
- Packing, jam, [2100](#)
- Packing, live-loaded, [2101](#)
- Packing, valve, [2098](#)
- Paperless chart recorder, [517](#)
- Parallel damper, [2095](#)
- Parallel digital data, [992](#)
- Parallel PID equation, [2315](#)
- Parallel pipe threads, [574](#)
- Parallel versus serial digital data, [1023](#)
- Parity bit, [1034](#)
- Parshall flume, [2063](#)
- Partial stroke valve testing, [2651](#), [2684](#)
- Particle, [228](#)
- Parts per billion (ppb), [1852](#)
- Parts per million (ppm), [283](#), [1304](#), [1845](#), [1852](#), [1858](#), [1862](#), [1864](#)
- Pascal’s principle, [175](#), [2220](#)
- Pascal, pressure unit, [171](#)
- Pass Token (PT), Fieldbus, [1159](#)
- Passivation layer, metals, [268](#), [2215](#), [2229](#)
- Password lockout, [2733](#)
- Password timeout, [2733](#)
- Passwords, [2723](#), [2725](#), [2727](#), [2729](#), [2732](#)
- Path loss, RF, [1227](#)
- Pauli Exclusion Principle, [241](#)
- Payload, computer virus, [2716](#)
- PC-ControLab software, [2475](#)
- PCOF, [638](#)
- Per unit, [869](#), [883](#)
- Percent, [869](#)
- Periodic Table of the Elements, [236](#)
- Periodic waveform, [1877](#)
- Permanent pressure drop, [1742](#), [2211](#)
- Permanent pressure loss, [219](#)
- Permittivity, [346](#)
- Permittivity, relative, [1475](#)
- PFD, [533](#), [2615](#), [2639](#), [2685](#)

- pH, 284, 1302
- pH neutralization, 288
- Phase change, 1291
- Phase current, 448
- Phase margin, 2440
- Phase reference signal, vibration monitoring, 1887
- Phase rotation, 441, 458, 467
- Phase sequence, 441, 458, 467, 723
- Phase shift, process dynamic, 2439, 2443
- Phase voltage, 448
- Phase-shift oscillator circuit, 2440
- Phasometer, 382
- Phasor, 375
- Phasor diagram, 442
- Photoelectric effect, 1843
- Photoionization detector, GC, 1789
- Photomultiplier tube, 1842, 1848
- Photon, 250
- Pickoff coil, 1664
- Pickup coil, 1664
- Pickup current, protective relay, 1997
- Pickup, protective relay, 1990, 1996
- Piecewise function, 2068
- Pieruschka, Erich, 2618
- Piezometer, 1588
- Pigtail siphon, 1393
- Pilot burner, 2690
- Pilot valve, 930
- Pilot, protective relay, 2012
- Pilot-loaded pressure regulator, 2120
- Pilot-operated control valve, 2118
- Pilot-operated valve, for overpressure protection, 2668
- Ping, utility program, 1074, 1078
- Ping6, utility program, 1082
- Pipe elbow flow element, 1631
- Pipe flange, 564
- Pipe hanger, 1493
- Pipe taps (orifice plate), 1619
- Piping and Instrument Diagram (P&ID), 535
- PIS, 2669
- Pitch, thread, 572
- Pitot tube, 1625
- Planck's constant, 250, 1805
- Planck, Max, 250, 1805
- Planetary gear set, 109
- Plato, degrees, 181
- PLC, 748, 758, 761, 765, 878, 1107
- Plug, tube, 587
- Plugging, AC motor, 2250, 2256
- PMV model 1500 valve positioner, 2149
- Pneumatic, 173
- Pneumatic "resistor", 1649
- Pneumatic control system, 503
- Pneumatic deadweight tester, 1298
- Pneumatic diaphragm valve actuator, 2109
- Pneumatic piston valve actuator, 2114
- Pneumatic relay, 932
- Pneumatic valve actuator, 2109
- PNP output switch, 665
- Poise, 199
- Polar form, complex number, 375
- Polarity, 296
- Pole, transfer function, 411, 413, 419
- Pole-zero plot, 415
- Polling, 1041
- Polybius, Greek historian, 987
- Polynomial expression, 429
- Porpoising of process variable, 2474
- Port number (OSI layer 4), 1090
- Port, control valve trim, 2184
- Port-guided globe valve, 2082
- Position algorithm, defined, 2395
- Positive displacement pump, 695
- Positive feedback, 2420, 2441, 2443
- Positive filtering (spectroscopy), 1829
- Postel, Jon, 1076
- Potential energy, 80, 278, 292
- Potential transformer (PT), 736, 1899, 1945
- Potential transformer (PT) safety, 1964
- Pound per square inch, pressure unit, 171
- Poundal, 1709
- Power, 75
- Power Line Carrier (PLC) communications, 1902
- Power line carrier telemetry, 976
- Power line charging current, 2012
- Power reflection factor, 1478
- Powers and roots, 2065
- ppb, 1852
- ppm, 283, 1304, 1845, 1852, 1858, 1862, 1864
- Pre-act control action, 2295, 2299

- Preamplifier, pH probe, 1778
- Precipitate, 229
- Precision potentiometer, 1289, 1557
- Predictive maintenance, 1260
- Pressure, 170, 171, 1398
- Pressure gauge mechanism, typical, 1316
- Pressure recovery, 219
- Pressure recovery factor, 2210
- Pressure Relief Valve (PRV), 2121, 2658
- Pressure Safety Valve (PSV), 2121, 2658
- Pressure snubber, 1368
- Pressure switch, 668
- Pressure, absolute, 186
- Pressure, differential, 186
- Pressure, gauge, 186
- Pressure, hydrostatic, 178
- Pressure, negative, 189
- Pressure, within solids, 189
- Pressure-based flowmeters, 1583
- Pressurized Water Reactor (PWR), 163, 2698
- Preventive maintenance, 2642
- Primary Coated Optical Fiber, 638
- Primary injection test, current transformer, 1978
- Primary sensing element, 499
- Prism, 1812
- Probability, 2612
- Probability and Boolean values, 2614
- Probability of Failure on Demand (PFD), 2615, 2639, 2685
- Probe Node (PN) token, Fieldbus, 1159
- Problem-solving technique: applying numerical values, 2206
- Problem-solving technique: divide and conquer, 1075
- Problem-solving technique: limiting cases, 412, 960, 2626, 2786, 2788
- Problem-solving technique: thought experiment, 508, 949, 1420, 1445, 1446, 1448, 1449, 1463, 1476, 1531, 1634, 1735, 1819, 1823, 2204, 2206, 2285, 2327, 2398, 2413, 2472, 2542, 2572, 2576, 2586, 2678, 2755, 2787, 2788, 2790
- Procedure, energy control, 86
- Process, 498, 2266
- Process alarm, 2386
- Process and Instrument Diagram (P&ID), 535
- Process Flow Diagram (PFD), 533
- Process load, 2522
- Process switch, 521, 654
- Process variable, 498, 2267
- Product, chemical reaction, 265
- Profibus, 786, 1096
- Profibus PA, 1028, 1135
- Profile factor, multipath ultrasonic flowmeter, 1693
- Programmable Logic Controller, 748, 758, 761, 765, 878, 1107
- Programming, chromatograph, 1800
- Projectile physics, 89
- Proof of closure switch (safety valve), 2690
- Proof testing, 2650
- Proportional band, 2282, 2297
- Proportional control, 2276
- Proportional control action, 2297
- Proportional weir, 1655
- Proportional-only offset, 2286, 2292, 2410
- Protection zone, power system, 2010, 2027
- Protective Instrument System (PIS), 2669
- Protective relay, 464, 735, 1903, 1927, 1985, 2670
- Protective relay codes, 1992
- Proton, 228
- Prover, flow, 1301, 1666
- Proximitors, Bently-Nevada, 1884
- Proximity switch, 663
- Prussian blue, 2085
- PRV, 2121, 2658
- Pseudocode, 36, 46, 2388
- PSI, 171
- PSV, 2121, 2658
- PT, 736
- PT, potential transformer, 1899, 1945
- PTFE (Teflon) valve stem packing, 2104
- Puget Sound Energy, 1910
- Pull string, 606
- Pulley, 101
- Pulse test, 2651
- Pulse width modulation, 2240
- Pump curve, 2196
- Purge cycle, 2690
- Purge flow rate, 1384, 1424
- Purge flow regulator, 1386
- Purged impulse line, 1384

- PWM, 2240
- PWR, 163, 2698

- QCL, 1837
- Quadrant-edge orifice plate, 1616
- Quadrature pulse, 825
- Quality factor, RLC filter circuit, 436
- Quantization error, 993
- Quantum cascade laser (QCL), 1837
- Quantum theory of light, 251
- Quarter-active bridge circuit, 334
- Quarter-wave antenna, 492, 1208
- Quarter-wave damping, 2459
- QUB, Fieldbus, 1161
- Quick-opening valve characterization, 2199
- QUU, Fieldbus, 1161

- R-value, 129
- R-X diagram, 2032
- Racking out, circuit breaker, 1927
- Radar detection of liquid interfaces, 1478
- Radar level instrument, 483, 1473
- Radial damper, 2096
- Radiation, heat, 1510
- Radiation, heat transfer, 127
- Radiation, nuclear, 1497
- Radio Frequency (RF), 1146, 1222
- Radio frequency interference from motor drive circuits, 2244, 2250
- Radio wave, 488
- Radioactivity, 237
- Radiotelegraph, 1024
- Raised well manometer, 1310
- Raised-Face (RF) flange, 567
- Ramp-and-soak setpoints, 2500
- Range wheel, 953
- Rangeability, 1660, 1776, 2165
- Rangeability, optical flowmeter, 1698
- Rangeability, pressure-based flowmeter, 1642
- Rangedown, 1281
- Ranging, 1251
- Rankine, 123
- Ransomware, 2718
- Rate control, 2295
- Rate control action, 2295, 2299
- Rate limit function, 2565

- Ratio control strategy, 2510
- Ratio station, 2514
- Ratio, gear, 106
- RC phase-shift oscillator circuit, 2440
- Reactance, 360
- Reactance versus resistance, 386
- Reactant, chemical reaction, 265
- Reaction rate, 2463
- Reaction, enthalpy of, 274
- Reaction, heat of, 274
- Reactor, power line, 2260
- Reading, active, 2744
- Real Gas Law, 198
- Real number, 371
- Real number (floating-point), 982
- Receiver gauge, 916, 922, 927, 1605, 1607, 2059
- Reclosing relay, 1986
- Recorder, 514, 2456
- Rectangular form, complex number, 375
- Rectangular weir, 1653
- Rectifier, SCR controlled, 2241
- Recursion, computer programming, 2393
- Red Lion Controls panel-mounted indicator, 513
- Red-line editing, 754
- Reduced-port control valve trim, 2184
- Reducing union, tube, 585
- Redundancy, 2373, 2645
- Redundant sensors, 2676
- Redundant transmitters, 2571
- Reference electrode, 1767
- Reference junction compensation, 1540
- Reference junction, thermocouple, 1531
- Reference probe, radar level instrument, 1481
- Reference pulse, radar, 1482
- Reflection factor, 1478
- Reflection grating, 1813
- Refractive index, 636
- Regenerative braking, AC motor, 2250, 2254
- Regulator, purge gas flow, 1386
- Relation control strategy, 2518
- Relative flow capacity, 2187
- Relative gas density, 180
- Relative permittivity, 1475
- Relative permittivity, influence on radar level measurement accuracy, 1480
- Relay, 499

- Relay Ladder Logic programming, 757, 802
- Relay, ice cube, 751
- Reliability, 2610, 2623, 2636
- Relief valve, 2121, 2658
- Remanence, CT, 1977
- Remote backup, protective relay, 2046
- Remote PLC I/O, 776
- Remote seal, 1372
- Remote setpoint, 2383, 2500, 2502
- Remote telemetry system, 976, 2372
- Remote Terminal Unit (RTU), 976, 2368
- Repose, angle of, 1472, 2072
- Request timed out, error message, 1079
- Required Safety Availability (RSA), 2685
- Reseat pressure, safety valve, 2658, 2665
- Reset coil, PLC programming, 818
- Reset control action, 2291, 2298
- Reset deadband, 2294, 2450
- Reset windup, 2294, 2385, 2411
- Reset windup, limit controls, 2568
- Reset windup, override controls, 2584
- Residence time, see Retention time, 2527
- Resistance, 313, 360
- Resistance versus reactance, 386
- Resonance, 423, 485
- Resonance, mechanical, 2713
- Resonant wire pressure sensor, 1335
- Resource block, Fieldbus, 1176
- Response factor, chromatograph detector, 1793
- Restrained differential relay, 2007
- Restraint coil, protective relay, 2007
- Restricted-capacity control valve trim, 2184
- Retention time, 1784, 1787, 1803, 2526, 2527, 2694
- Retentive instruction, PLC programming, 818, 828
- Reverse indication, controller output, 888, 2133
- Reverse valve actuator, 2128
- Reverse valve body, 2129
- Reverse-acting controller, 508, 2278, 2585
- Reverse-acting pneumatic relay, 933
- Reverse-acting transmitter, 539
- Reverse-acting valve body, 2080
- Reynolds number, 201
- Reynolds number, for laminar versus turbulent flow regimes, 204
- RF, 1146, 1222
- RF flange, 567
- RFI, 2244, 2250
- Richter scale, 1776
- Riemann sum, 19
- Riemann, Bernhard, 19
- Right-hand rule, 93, 304, 343
- Ring terminal, 600
- Ring-Type Joint (RTJ) flange, 567
- Rising stem valve actuator, 2126
- RLL, 757, 802
- RMS quantities, 356
- Robertshaw Vibraswitch, 1891
- Rockwell ControlLogix 5000 PLC, 771
- Rockwell PLC-5, 768
- Rockwell SLC 500 PLC, 769
- Root privilege, computer, 2728
- Root, polynomial, 429
- Root-mean-square (RMS) quantities, 356
- Roots and powers, 2065
- Rosemount Analytical X-STREAM X2 gas analyzer, 1828
- Rosemount field-mounted indicator, 513
- Rosemount Micro-Motion Coriolis mass flowmeter, 1719
- Rosemount model 1151 differential pressure transmitter, 966, 1328, 1342, 1419, 1621
- Rosemount model 1151 gauge pressure transmitter, 1377
- Rosemount model 268 HART communicator, 1095
- Rosemount model 3051 differential pressure transmitter, 904, 1331, 1342, 1426, 1642
- Rosemount model 3051S differential pressure transmitter, 1333
- Rosemount model 3095MV multi-variable transmitter, 1176, 1643
- Rosemount model 3301 guided-wave radar transmitter, 2071
- Rosemount model 3301 level transmitter, 1482
- Rosemount model 5708 3D Solids Scanner, 2072
- Rosemount model 648 WirelessHART temperature transmitter, 1245
- Rosemount model 8700 magnetic flowmeter, 1683
- Rosemount model 8800C vortex flow transmitter, 1674

- Rosemount model CT5400 stack gas analyzer, 1836
- Rotabolt self-indicating bolt, 568
- Rotameter, 1384, 1424, 1650, 2108
- Rotating magnetic field, 719, 2245, 2813
- Rotating paddle level switch, 677
- Rotork electric valve actuator, 2124
- Router, 1073
- RS-232 serial communication, 1050, 1066, 1074, 1089, 1106, 2737
- RS-422 serial communication, 1054, 2737
- RS-485 serial communication, 1054, 1066, 1106, 2737
- RSA, 2685
- RSView HMI software, Rockwell, 854
- RTD, 1289, 1517
- RTD table, 1519
- RTJ flange, 567
- RTU, 976, 2368
- RTU Modbus frames, 1111
- Run tee fitting, 586
- Runaway process, 2419
- Rung, PLC programming, 804
- Rupture disk, 2657
- Rupture disk, graphite, 2658
- Rupture disk, metal, 2657
- RVDT, 2125, 2146
- SAE straight thread pipe fittings, 574
- Safety barrier circuit, intrinsic, 2607
- Safety Instrumented Function (SIF), 2669
- Safety Instrumented System (SIS), 2669
- Safety PLC, 2680
- Safety relief valve, 2665
- Safety valve, 2121, 2658
- Safety, instrument transformer, 1964
- Sage “Prime” model thermal mass flowmeter, 1727
- Salt, 288
- SAMA diagram, 540
- Sample rate, 996
- Sample time, 996
- Sample-and-hold PID algorithm, 2446
- Sand bath temperature calibrator, 1293
- Saturated steam, 165
- SCADA, 976, 1902, 1921, 2368
- SCFM, 1704
- Scheduled communication, Fieldbus, 1155, 1159
- Schrödinger, Erwin, 252
- Schweitzer “Best-Choice Ground Directional Element” protective relay logic, 2025
- Schweitzer Engineering Laboratories model SEL-387L differential current relay, 2012
- Schweitzer Engineering Laboratories model SEL-551 overcurrent relay, 1991, 2003
- Scintillation optical flowmeter, 1697
- SCOF, 638
- Scram, 2700
- Screwless terminal block, 596
- Seal pot, used on wet leg of level measurement system, 1435
- Seal-in contact, 739, 794, 1999
- Seat load, control valve, 2138, 2145
- Seating profile, electronic valve positioner diagnostic, 2157
- Seattle City Light, 1910
- Secant line, 34
- Second derivative, calculus, 12
- Second Law of Motion, 74, 97, 192
- Second-order lag, 2436
- Secondary Coated Optical Fiber, 638
- Secondary emission, electrons, 1843
- Security, 2610, 2611
- Security, versus dependability, 2670
- Segmental orifice plate, 1613
- Segmental wedge, 1631
- Segmented ball valve, 2093
- Segway personal transport, 2420
- Selective permeability, 1764
- Selectivity, protective relay, 2010
- Self-balancing bridge, 329
- Self-balancing system, 926, 1338
- Self-diagnostics, 2653
- Self-induction, 362
- Self-powered transmitter, 889
- Self-purged impulse line, 1390
- Self-regulating process, 2409
- Sense wires, 4-wire RTD circuit, 1521
- Sensible heat, 144
- Sensing line, 579, 1349
- Sensing tube, 579, 1349
- Sensitivity, radio receiver, 1226

- Sequenced control valves, 2159
- Sequencer Compare instruction, Allen-Bradley PLC programming, 847
- Sequencer Load instruction, Allen-Bradley PLC programming, 847
- Sequencer Output instruction, Allen-Bradley PLC programming, 843
- Serial digital data, 992
- Serial versus parallel digital data, 1023
- Series PID equation, 2317, 2348
- Servo motor control, 2508
- Set coil, PLC programming, 818
- Setpoint, 499, 2269
- Setpoint tracking, 543, 2358, 2385
- Setpoint, remote, 2383, 2500, 2502
- Shaded-pole AC induction motor, 726
- Shading coil, 726
- Sheave, 117
- Shelf life, pH electrode, 1770
- Shelf life, span gas, 1859
- Shell, electron, 242
- Shell-and-tube heat exchanger, 132
- Shielded cables, 620
- Shielding, magnetic, 623
- Shift register, FIFO, 1267
- Shift register, used to implement dead time, 2552
- Shinsky, Francis Greg, 3
- Short, electrical fault, 324
- Shunt resistor, 901
- Shutter, radioactive source lock-out, 1499
- Siemens 505 PLC, 767
- Siemens APACS control system, 2681
- Siemens model 353 digital controller, 2356, 2359
- Siemens Procidia controller GUI software, 2360
- Siemens Quadlog safety PLC, 2680
- Siemens S7-300 PLC, 771, 2729
- SIF, 2669
- Sightfeed bubbler, 1424
- Sightglass, 1404
- Signal, elevated, 1011
- Signal, floating, 1010
- Signal, ground-referenced, 1009
- Signature, control valve diagnostics, 2156
- Signed integer, 979
- Silicon resonator pressure sensor, 1335
- Silicon strain gauge element, 1322
- Simple apparatus (intrinsically safe), 2608
- Simple machine, 99
- Simplex, 1040
- Simultaneous systems of linear equations, 268
- Single-ended analog input, 1013
- Single-ended signaling, 1050
- Single-helical gear, 107
- Single-mode optical fiber, 639
- Single-phase AC induction motor, 723
- Single-phasing a three-phase AC induction motor, 727
- Single-precision floating-point number, 983, 1126
- Sinking current, 304, 778, 781
- Sinking output switch, 664
- SIP, 575, 1371
- SIS, 2669
- Skydrol hydraulic fluid, 698
- Slack diaphragm, 1315
- Slack-tube manometer, 1300
- Slide rule, 1624, 2180
- Slip speed, 723, 2245, 2246
- Slip-stick cycle, 2449
- Slope, pH instrument, 1780
- Slurry, 2224
- SMA style connector, 641
- Smart instrument, 1269
- Smart transmitter, 905
- SMTP, 1090
- Snell's Law, 636
- Snubber, pressure, 1368
- Society of Automotive Engineers (SAE), 574
- Socrates, 3185
- Sodium error, pH measurement, 1766
- Soft alarm, 2386
- Soft constraint, 2584
- Soft override, 2583
- Software compensation, thermocouple, 1548
- Sol, 229
- Solenoid, 341
- Solenoid valve, 342, 702
- Solid, 170
- Solute, 229
- Solution, 229
- Solvent, 229
- Solver, logic, 2680
- Solving for variables in an equation, 2761

- Sonic flow, 2220
- Sonic level instrument, 1468
- Source versus load, 306
- Sourcing current, 304, 778, 781
- Sourcing output switch, 664
- Soxhlet, degrees, 181
- Space, 1027, 1029
- Span, 499
- Span adjustment, 1254
- Span gas, 1304, 1835, 1845, 1855, 1858
- Span shift, 1256
- SPDT switch contact, 750
- SPEC 200 analog electronic control system, 2352, 2376, 2387, 2641
- Species, chemical, 1784
- Species, chemical composition, 284, 1766, 1788, 1816, 1818
- Specific gravity, 180, 192, 1380, 1418
- Specific heat, 1729
- Specific volume, 180
- Spectacle blind, pipe, 569
- Spectroscope, 252
- Spectroscopic notation, 244
- Spectroscopy, absorption, 256
- Spectroscopy, emission, 253
- Speed of light, 636
- Speed of sound measurement, transit-time ultrasonic flowmeter, 1691
- Speed of sound through a fluid, 1689
- Spiral bourdon tube, 1315
- Split-range control valves, 2159
- SPLT Fieldbus function block, 2171
- Spool valve illustration, 2-way, 697
- Spool valve illustration, 4-way, 711–713
- Spool, pipe, 566
- Spread-spectrum radio, 1236, 1238, 2736
- Spring adjuster, valve actuator, 2135
- Spring-loaded pressure regulator, 2119
- Sprocket, 120
- SPST switch contacts, 749, 1966
- Spur, 1062, 1141
- Spur gear, 106, 2124
- Spurious trip, 2670
- SQC instruction, Allen-Bradley PLC programming, 847
- SQL instruction, Allen-Bradley PLC programming, 847
- SQO instruction, Allen-Bradley PLC programming, 843
- Square root characterizer, 1601, 2062, 3159, 3161, 3164
- Square root extractor, 1602
- Square root scale, 1605, 2059
- Square-D Motor Logic Plus control, 744
- Square-edged concentric orifice plate, 1610
- Squirrel-cage AC induction motor, 722
- SSH, 1090, 2729
- ST style connector, 629, 638, 641
- Stagnation pressure, 1586
- Standard cell, 1287, 1557
- Standard conditions, chemistry, 278
- Standard cubic feet per minute, 1704
- Standard enthalpy of formation, 278
- Standard heat of formation, 278
- Standard temperature and pressure (STP), 278
- Starter, motor, 731
- Static contact, PLC programming, 820
- Static friction, 2204
- Static pressure effect, pressure transmitter, 1345
- Stationary phase, 1784
- Stator, 719, 2245
- Stator winding, 439
- Status propagation, Fieldbus, 1178
- Steady-state gain, 2425
- Steam cut, valve trim, 2225
- Steam jacket, 507
- Steam table, 148, 163, 165
- Steam tracing, 1387
- Steam trap, 1387
- Steam, industrial uses of, 164
- Steam-hydrocarbon reforming process, 2514
- Steam-In-Place, 575, 1371
- Stefan-Boltzmann equation, 2073
- Stefan-Boltzmann Law, 127, 1561
- Stem connector, control valve, 2135
- Stem packing lubricator, 2103, 2207
- Stem valve, 936
- Stem-guided globe valve, 2081
- Step 7 PLC software, Siemens, 2716
- Step-index, 640
- Stiction, 2205

- Stilling well, 1502, 1659
 Stinger, pressure test accessory, 1367
 Stoichiometry, 265
 Stokes, 200
 Store and forward, 1045
 STP, Standard Temperature and Pressure, 278
 Strain gauge, 331, 1321
 Strapping table, 2070
 Stress, 171
 Strip chart recorder, 516
 Strobe light, 382, 997
 Stroboscope, 382, 997
 Strong acid, 286
 Strong base, 287
 Strong nuclear force, 230
 Strouhal number, 1671
 Strouhal, Vincenc, 1671
 Structural chemical formula, 233
 Stub, 1062, 1141
 Stuxnet computer virus, 2706
 Sublimation, phase change, 154, 2713
 Subnet mask, 1077
 Subshell, electron, 242
 Substitution, algebraic, 2774
 Subtractive transformer windings, 1959
 Superconductivity, 313
 Superfluidity, 313
 Superheat, 166, 2519
 Superheated steam, 166, 2519
 Superheated vapor, 149
 Superheater, 2519
 Supernatant, 229
 Supervisory Control And Data Acquisition, 976, 1902, 1921, 2368
 Surge impedance, 476
 Susceptance, 360
 Suspension, 229
 Sutro weir, 1655
 Swagelok instrument tube fittings, 581
 Swamping, 177, 337, 1056, 1520, 1646, 2439, 3164
 Switch, 653
 Switch, dry contact, 663
 Switch, process, 521, 654
 Switch, wet contact, 663
 Switching hub, Ethernet, 1069
 Symbol, 802
 Symbolic addressing, 802
 Symmetrical components, three-phase power system, 458, 2023
 Synchronization lamp, 1914
 Synchronization, generator, 1914
 Synchronous data transfer, 1026, 1032
 Synchronous motor, 2245
 Synchronous speed, 723, 2245
 Synchrophasor, 385
 Synergy automobile hybrid electric drivetrain (Toyota), 110
 Système International, 72, 255
 Systems of linear equations, 268
 Table, RTD values, 1519
 Table, strapping, 2070
 Table, thermocouple values, 1538
 Tag name, 802, 851
 Tag name, naming conventions for, 852
 Tangent line, 34
 Tank circuit, 423, 485
 Tank expert system, 1439
 Tap hole finish, orifice plate, 1619
 Tape-and-float level measurement, 1413
 Tapered pipe threads, 571
 Tare weight, 1490
 Target flow element, 1628
 Taylor analog electronic controller, 2350
 TCP, 1072, 1088
 TDL, 1837
 TDMA, 1043
 TDR, 627
 Tee tube fitting, branch, 586
 Tee tube fitting, run, 586
 Tee tube fitting, union, 586
 Teflon (PTFE) valve stem packing, 2104
 Telegraph, 1024
 Telemetry system, 976, 2372
 TELNET, 1090, 2729
 Temperature calibrator, dry-block, 1295
 Temperature calibrator, oil bath, 1293
 Temperature calibrator, sand bath, 1293
 Temperature coefficient of resistance, RTD wire, 1517, 1519
 Temperature programming, chromatograph, 1800

- Temperature switch, 682
- Temperature, defined for a gas, 1509
- Terminal block, 594
- Terminal strip, 594
- Terminal, fork versus ring, 600
- Termination resistor, 481, 627
- Test diode, 899
- Test switch, power instrumentation, 1965
- Test Uncertainty Ratio, 1285
- Test, computer programming, 2390
- Texas Instruments 505 PLC, 767
- Thayer ball mill controller, 45
- Thermal conductivity detector, GC, 1790
- Thermal energy, 1509
- Thermal image, circuit breaker, 1573
- Thermal image, electric motor, 1571
- Thermal image, motor starter, 1572
- Thermal imager, 1570
- Thermal mass flowmeter, 1726
- Thermal overload heater, 730
- Thermal siphon, 135, 165
- Thermistor, 1517
- Thermocouple, 1289, 1530
- Thermocouple burnout detection, 1560
- Thermocouple compensation, hardware, 1547
- Thermocouple compensation, software, 1548
- Thermocouple table, 1538
- Thermosiphon, 135, 165
- Thermowell, 1574
- Thin-layer chromatography, 1785
- Third Law of Motion, 74
- Thought experiment, 508, 949, 1420, 1445, 1446, 1448, 1449, 1463, 1476, 1531, 1634, 1735, 1819, 1823, 2204, 2206, 2285, 2327, 2398, 2413, 2472, 2542, 2572, 2576, 2586, 2678, 2755, 2787, 2788, 2790
- Thread pitch, 572
- Three-element boiler feedwater control, 2529
- Three-valve manifold, 1359
- Three-way globe valve, 2088, 2161
- Three-wire RTD circuit, 1522, 1523
- Throttling control valve, 2079
- Tieback variable, Allen-Bradley Logix5000 PLC programming, 2570
- Tilt switch, mercury, 669, 675
- Time constant, 2432, 2546
- Time constant, differentiator circuit, 2343
- Time constant, integrator circuit, 2344
- Time Distribution (TD) message, Fieldbus, 1159
- Time Division Multiple Access (TDMA), 1043
- Time domain, 1880
- Time domain reflectometry, 1466
- Time overcurrent protection, 1996
- Time overcurrent relay, 1986
- Time-Domain Reflectometer, 627
- Timer circuit, watchdog, 2653
- Timer, PLC programming, 827
- Tofino industrial network firewall, 2735
- Token Ring, 1042
- Token-passing channel arbitration, 1042
- Top Hat DIN rail, 603
- Toroidal conductivity cell, 1759
- Torque, 82
- Torque tube, 1455
- Torr, 188
- Torricelli, Evangelista, 218
- Toshiba magnetic flowmeter, 1683
- Total heat of steam, 148
- Total internal reflection, 636
- Totalizer, 42
- Toyota Prius hybrid electric automobile, 110
- Traceability, NIST, 1284
- Transducer, 499
- Transducer block, Fieldbus, 1176
- Transfer function, 437
- Transfer function, pole, 411, 413, 419
- Transfer function, zero, 411, 413, 419, 424
- Transfer functions, 411
- Transform function, 1216
- Transformer, 364
- Transformer impedance, 370
- Transformer, AC, 361
- Transformer, additive windings, 1959
- Transformer, subtractive windings, 1959
- Transient fault, power line, 1942
- Transit-time ultrasonic flowmeter, 1690
- Transition zone, radar, 1484
- Transition-sensing contact, PLC programming, 820
- Transmation model 1040 loop calibrator, 911
- Transmission Control Protocol, 1072, 1088
- Transmission line, 476, 487, 626

- Transmission power line, 1898
- Transmitter, 499
- Transport delay, 519, 1800, 2443, 2542
- Trap, 1393
- Trap, steam, 1387
- Trend recorder, 514, 2456
- Triaxial vibration probe array, 1886
- Triconex TMR safety control system, 2682
- Trim, 2079
- Trim, in a feedforward control system, 2528
- Trip curve, thermal overload, 733
- Trip solenoid, 2683
- Triple Modular Redundant (TMR) control system, 2682
- Triple point, water, 152, 1291
- Triple-cascaded loops, 2508
- True state, PLC programming, 804
- Tube bulkhead fittings, 585
- Tube cap, 587
- Tube connector, 585
- Tube fitting gauge, 581
- Tube plug, 587
- Tube union, 585, 588
- Tunable diode laser (TDL), 1837
- Tuning fork level switch, 676
- TUR, 1285
- Turbine flow element, 1662
- Turbulent flow, 204
- Turck Fieldbus coupling device, 1145
- Turndown, 1281
- Turndown ratio, flume flowmeter, 1661
- Turndown ratio, optical flowmeter, 1698
- Turndown ratio, pressure-based flowmeter, 1608, 1642
- Turndown ratio, turbine flowmeter, 1665
- Turndown ratio, weir flowmeter, 1661
- Turns ratio, transformer, 96
- Twaddell, degrees, 180
- Twin-turbine mass flowmeter, 1710
- Twisted, shielded pair cables, 625
- Two's complement, 979
- Two-wire RTD circuit, 1520
- U-tube manometer, 1310
- UDP, 1088
- UEL, 2602
- Ullage, 483, 1410, 1468, 1486
- Ultimate gain, 2458, 2459
- Ultimate period, 2458
- Ultimate sensitivity, 2458
- Ultrasonic flowmeter, 1687
- Ultrasonic level instrument, 1468
- Ultrasonic level switch, 678
- Unbalanced signaling, 1050
- Undependability, 2611, 2685
- Under-damping, 430
- Underreaching, protective relay, 2028
- Uninterruptible power supply, 1988
- Union tee fitting, 586
- Union, tube, 585, 588
- Unit conversions, 61
- Unit reaction rate, 2464
- Unity fraction, 61, 271
- Universal Serial Bus, 1050, 1055
- Universal solvent, 229
- Unreliability, 2623
- Unrestrained differential relay, 2007
- Unscheduled communication, Fieldbus, 1155, 1159
- Unsecurity, 2611
- Unshielded, twisted pair (UTP) cable, 1065
- Unsigned integer, 979
- Up-down calibration test, 1278
- Upper explosive limit, 2602
- Upper range value, 499, 1270, 1300
- UPS, 1988
- URV, 499, 1270, 1300
- USB, 1050, 1055
- Useful life of a component or system, 2634
- User Datagram Protocol, 1088
- UTP cable, 1065
- V-cone flow element, 1630
- V-notch weir, 1653
- V/F ratio, 2248
- V/Hz ratio, 2248
- V1 missile development, 2618
- Vacuum circuit breaker, 1930
- Vacuum tube voltmeter, 337
- Vacuum, produced by a venturi, 216
- Valence electrons, 246, 248, 298
- Valve actuator, 2079

- Valve body, [2079](#)
- Valve characterization, equal percentage, [2198](#)
- Valve characterization, linear, [2198](#)
- Valve characterization, quick-opening, [2199](#)
- Valve manifold, [1359](#)
- Valve packing, [2098](#)
- Valve seal overtravel interlock switch, [2690](#)
- Valve signature, electronic positioner diagnostics, [2156](#), [2208](#)
- Valve sizing equation, liquid, [2179](#)
- Valve stem packing lubricator, [2103](#)
- Valve trim, [2079](#)
- Valve, solenoid, [342](#)
- Valves, [2079](#)
- Varec pressure relief valve, [2659](#)
- Variable displacement pump, [696](#)
- Variable-area flowmeter, [1650](#)
- Variable-frequency drive, [84](#), [605](#), [887](#), [981](#), [1041](#), [1107](#), [1137](#), [1235](#), [1560](#), [2236](#), [2246](#), [2427](#), [2581](#), [2644](#), [2799](#)
- Variable-speed drive, [2240](#)
- VCR, Fieldbus, [1161](#)
- Vector cross-product, [93](#), [304](#), [1713](#)
- Vector dot-product, [94](#)
- Velocity algorithm, defined, [2395](#)
- Velocity factor, multipath ultrasonic flowmeter, [1693](#)
- Velocity factor, transmission line, [483](#), [636](#)
- Velocity of approach factor, [1646](#)
- Vena contracta, [1609](#), [2213](#), [3158](#)
- Vent hole, orifice plate, [1614](#)
- Vent valve fitting, [1365](#)
- Venturi tube, [219](#), [1588](#)
- VFD, [84](#), [605](#), [887](#), [981](#), [1041](#), [1107](#), [1137](#), [1235](#), [1560](#), [2236](#), [2246](#), [2427](#), [2581](#), [2644](#), [2799](#)
- Vibrating fork level switch, [676](#)
- Vibration loop, tubing, [589](#), [2111](#)
- Virtual Communication Relationship (VCR), Fieldbus, [1161](#)
- Virtual Private Network, [2736](#)
- Viscosity, [199](#)
- Viscosity, absolute, [199](#)
- Viscosity, kinematic, [199](#)
- Viscosity, temperature dependence, [200](#)
- Viscous flow, [201](#)
- VisSim Realtime computer software, [2495](#)
- Volt, [293](#), [295](#)
- Volta, Alessandro, [293](#)
- Voltage, [292](#)
- Voltage follower circuit, [940](#)
- Voltage signal, elevated, [1011](#)
- Voltage signal, floating, [1010](#)
- Voltage signal, ground-referenced, [1009](#)
- Voltage transformer (VT), [1899](#), [1945](#)
- Voltage-to-frequency ratio, [2248](#)
- Volume booster, [2140](#), [2170](#)
- Volume, defined, [57](#)
- von Kármán, Theodore, [1671](#)
- Vortex flowmeter, [1672](#)
- Vortex street, [1671](#)
- Voting system, [2574](#)
- VPN, [2736](#)
- VSD, [2240](#)
- VT, voltage transformer, [1899](#), [1945](#)
- VTVM, [337](#)
- Wade Associates, Inc., [2475](#)
- Wallace & Tiernan, [1300](#)
- Wally box, [1300](#)
- Wastewater disinfection, [505](#), [2272](#)
- Watchdog timer circuit, [2653](#)
- Waveform, periodic, [1877](#)
- Wavenumber, [1809](#)
- Weak acid, [286](#)
- Weak base, [287](#)
- Wear-out period, [2634](#)
- Wedge, [104](#)
- Weighfeeder, [1730](#)
- Weight density, [57](#)
- Weight-based level instrument, [1490](#)
- Weir, [1653](#), [2063](#)
- Well manometer, [1310](#)
- Weschler panel-mounted bargraph indicator, [512](#)
- Westinghouse model CO-11 overcurrent relay, [1975](#)
- Weston cell, [1287](#)
- Wet calibration, [1453](#)
- Wet contact, [663](#)
- Wet leg, [1432](#)
- Whip antenna, [492](#), [1208](#)
- Whitelisting, application, [2728](#)
- Wild flow, [2513](#), [2575](#)

- Wild variable, [2271](#), [2513](#)
- Winch, [83](#)
- Wind-up, controller, [2294](#), [2385](#), [2411](#)
- Wire duct, [612](#)
- Wire loom, [612](#)
- Wire pulling, [606](#)
- Wire, jumper, [2811](#)
- Wireless transmitter, [2446](#)
- WirelessHART, [1043](#), [1233](#), [2736](#)
- WLAN, [1045](#)
- Wonderware HMI software, [854](#)
- Work, [75](#), [293](#)
- Worm gear, [113](#), [2122](#), [2124](#)

- X-windows, [1076](#)
- XD_Scale, FOUNDATION Fieldbus AI block
parameter, [1196](#)
- XIC (examine if closed), [814](#)
- XIO (examine if open), [814](#)

- Yagi antenna, [493](#), [1208](#)
- Yokogawa CENTUM distributed control system
(DCS), [2375](#), [2376](#)
- Yokogawa DPharp pressure transmitter, [1335](#)
- Yokogawa model DYF vortex flowmeter, [1182](#)
- Yokogawa model EJA110 differential pressure
transmitter, [1335](#), [1343](#)

- Zero, [499](#)
- Zero adjustment, [1254](#)
- Zero energy state, [1365](#)
- Zero gas, [1835](#), [1845](#)
- Zero sequence, [458](#), [465](#), [737](#), [2023](#)
- Zero shift, [1255](#)
- Zero, transfer function, [411](#), [413](#), [419](#), [424](#)
- Ziegler, J.G., [2457](#)
- Ziegler-Nichols “closed-loop” (Ultimate) PID
tuning example, [2478](#)
- Ziegler-Nichols “open-loop” (Reaction rate) PID
tuning example, [2476](#)
- Zirconium oxide, [2075](#)
- Zone of protection, power system, [2010](#), [2027](#)
- Zone, hazardous area, [2599](#)