


MODULAR ELECTRONICS LEARNING (MODEL) PROJECT



INDUCTORS AND INDUCTIVE CIRCUITS

© 2018-2025 BY TONY R. KUPHALDT – UNDER THE TERMS AND CONDITIONS OF THE
CREATIVE COMMONS ATTRIBUTION 4.0 INTERNATIONAL PUBLIC LICENSE

LAST UPDATE = 27 MARCH 2025

This is a copyrighted work, but licensed under the Creative Commons Attribution 4.0 International Public License. A copy of this license is found in the last Appendix of this document. Alternatively, you may 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.

Contents

1	Introduction	3
1.1	Recommendations for students	3
1.2	Challenging concepts related to inductance and inductors	5
1.3	Recommendations for instructors	7
2	Case Tutorial	9
2.1	Example: superconducting energy-storage inductor	10
2.2	Example: inductor storing and releasing energy	11
2.3	Example: series versus parallel inductors	14
2.4	Example: current source energizing an inductor	16
2.5	Example: voltage source energizing an inductor	19
2.6	Example: LR time-delay network values	22
2.7	Example: simplifying complex LR networks using Norton's Theorem	25
2.8	Example: interesting relay behavior	28
2.9	Example: custom inductor winding and testing	32
2.10	Example: measuring signal rates of change	38
3	Tutorial	43
3.1	Magnetic fields and energy storage	43
3.2	Voltage, current, energy, and inductance	47
3.3	Applications of inductance	49
3.4	Inverse exponential growth/decay	52
3.5	Inverse-exponential calculation examples	55
3.6	Inductor types and limitations	57
3.7	Series and parallel inductance	59
3.8	Making custom inductors	60
4	Historical References	63
4.1	Early examples of inductors	64
5	Derivations and Technical References	71
5.1	IEC standard component values	72
5.2	Inductor labeling	74
5.3	Inductor parameters and core types	76

5.4	Inductor parasitics	80
5.4.1	Model of a real inductor	80
5.4.2	Parasitic resistance in inductors	81
5.4.3	Parasitic capacitance in inductors	82
5.4.4	Other parasitic effects in inductors	82
5.5	Magnetic field quantities	83
5.6	Derivation of electron drift velocity	93
5.7	Time delay circuits and Thévenin's/Norton's Theorems	95
6	Programming References	99
6.1	Programming in C++	100
6.2	Programming in Python	104
6.3	Modeling inverse exponential growth and decay using C++	109
6.4	Modeling inverse exponential growth and decay using Python	116
6.5	Modeling an energizing inductor using C++	118
6.6	Solving RC and LR circuit differential equations using C++	121
6.6.1	Differential equation solver for an RC circuit and voltage source	122
6.6.2	Differential equation solver for an RC circuit and current source	123
6.6.3	Differential equation solver for an LR circuit and voltage source	124
6.6.4	Differential equation solver for an LR circuit and current source	125
7	Questions	127
7.1	Conceptual reasoning	131
7.1.1	Reading outline and reflections	132
7.1.2	Foundational concepts	133
7.1.3	Magnetic flux lines	135
7.1.4	Wirewound resistor	135
7.1.5	Impossibility of storing an energized inductor	136
7.1.6	Inductor as source versus load	137
7.1.7	Inductor charging circuit	138
7.1.8	A shocking experience	141
7.1.9	Delayed relay action	142
7.1.10	Explaining the meaning of calculations	144
7.2	Quantitative reasoning	146
7.2.1	Miscellaneous physical constants	147
7.2.2	Introduction to spreadsheets	148
7.2.3	Number of time constants	151
7.2.4	Inverse exponential functions	152
7.2.5	Inductor color codes	153
7.2.6	Voltage versus rate-of-current-change	153
7.2.7	Inductance of three-inductor network	154
7.2.8	Unknown inductor value	155
7.2.9	Resistance for specified time constant	155
7.2.10	Cylindrical inductor wire turns	156
7.2.11	Toroidal inductor wire turns	156
7.2.12	SPDT inductor switching circuit	157

7.2.13	Tabulating inductor charging	158
7.2.14	Time required to reach specified currents	159
7.2.15	Measuring inductance by time delay	160
7.2.16	Power line reactors	161
7.2.17	SPICE analysis of an energizing inductor	162
7.2.18	SPICE analysis of a de-energizing inductor	164
7.3	Diagnostic reasoning	166
7.3.1	Insulation breakdown	166
7.3.2	Incorrect voltage/current curves	167
7.3.3	Inductor proximity	167
7.3.4	Parasitic inductance on a PCB	168
7.3.5	Inductor saturation	168
A	Problem-Solving Strategies	169
B	Instructional philosophy	171
C	Tools used	177
D	Creative Commons License	181
E	References	189
F	Version history	191
	Index	196

Chapter 1

Introduction

1.1 Recommendations for students

Inductance is the ability to store energy in a magnetic field, caused by a flow of electric charge carriers through a conductor, typically a coil of wire wrapped around a ferromagnetic material such as iron. Any device constructed to exhibit this property is called an *inductor*. The amount of inductance offered by any inductor is directly proportional to the area of the wire coil, directly proportional to the permeability of the core material within the wire coil, directly proportional to the square of the number of wire turns, and inversely proportional to the axial space between turns of the coil (i.e. total length for any given number of turns). The amount of energy stored in any inductor is directly proportional to the inductance (measured in Henrys) and directly proportional to the square of the current as described by the formula $E_L = \frac{1}{2}LI^2$.

Important concepts related to inductance include **Conservation of Energy**, **magnetic fields**, **permeability**, **sources versus loads**, **“Ohm’s Law”** for inductors, **rates of change** (e.g. “speed”), natural **decay** functions, **time constant** (τ), inductor **ratings**, properties of **series** networks, properties of **parallel networks**, and **energy storage**.

Here are some good questions to ask of yourself while studying this subject:

- How might an experiment be designed and conducted to explore the phenomenon of inductance? What hypothesis (i.e. prediction) might you pose for that experiment, and what result(s) would either support or disprove that hypothesis?
- What is a *field*, and how might we measure its presence?
- How may we construct an inductor?
- What factors influence the amount of inductance offered by an inductor?
- Is the amount of energy stored within an inductor a function of its voltage, its current, or both?
- Which parameter is able to vary most rapidly for an inductor, voltage or current?

- What is the mathematical relationship between voltage and current for an inductor?
- What are some practical applications for inductors?
- Why do voltage and current tend to rise or “decay” in inverse-exponential fashion in a resistor-inductor circuit?
- What does *time constant* mean?
- What factors influence the maximum voltage or current ratings for an inductor?
- Why does inductance add in series and diminish in parallel? Explain in terms of the energy stored by each component versus by the larger network!
- What “parasitic” properties are found in real inductors other than inductance?
- Is there such a thing as a *polarized* inductor?

A useful resource for understanding inverse-exponential growth and decay as it relates to a resistor-inductor network is found in the *Programming References* chapter. There you will find example code showing how these changes in voltage and current occur over time.

The calculation of voltages and currents in inductor-resistor circuits is universally challenging for students new to the subject. Part of this challenge is a general unfamiliarity with exponential equations, but perhaps the biggest obstacle is the unrealistic expectation of students that *concepts should be completely understandable through explanation*. The truth is that deep concepts almost never clarify in one’s mind without *focused effort* and *application*. A very direct and effective way to learn new concepts with focused effort and application is to try solving the exact same problems shown as examples in the text, including *performing the author’s calculations on your own until you can see the patterns yourself*, explaining to yourself *the first principles underlying those calculations*, and *annotating the schematic diagrams* with your own notations showing current directions, voltage polarities, etc.

This advice – to actively articulate and solve (or at least attempt to solve) the example problems used by the author – is both simple and effective, but surprisingly difficult to inculcate as a habit. Most of my students seeing subjects like this for the first time find themselves staring at the text, trying to understand the author’s words, and shortly giving up when the desired understanding does not spring to mind. This is true even when they have practiced active approaches studying earlier concepts such as series-parallel resistor network analysis. They’ve been advised what to do, and indeed have done it before, but when presented with a new and daunting concept their behavior reverts to a familiar passivity.


Becoming a proficient technologist therefore consists in much more than just learning new concepts, most importantly learning *new habits of thought and of action*. When you find yourself challenged by a complicated concept, *actively apply it* and you will be surprised by how much better you will comprehend it! Always link each of your analytical steps to trusted principles so that you may explain *why* you’re doing the necessary steps, never satisfying yourself with mere procedures. Always bear in mind that your goal is not just to learn how to analyze circuits, but to learn how to tackle *any* new problem and learn *any* new concept more effectively.

1.2 Challenging concepts related to inductance and inductors

The following list cites concepts related to this module's topic that are easily misunderstood, along with suggestions for properly understanding them:

- **Determining initial and final values** – the fact that inductor current cannot instantaneously change but inductor voltage can is something many students try to rote-memorize. A far better approach is to reason from first principles. In this case, always begin by analyzing the circuit in terms of energy storage and energy conservation: inductors store energy in magnetic fields which are associated with current; this means an inductor's current cannot vary without work being done on or done by the inductor, which always requires some amount of time. Voltage, on the other hand, is what constitutes power (i.e. energy exchange per unit time) when combined with current. Therefore, voltage can jump to new values instantaneously across an inductor, such as when a de-energized inductor is suddenly connected to a source, or a pre-energized inductor is suddenly connected to a load.
- **Rates of change** – When learning the relationships between voltage and current for inductors and capacitors, one must think in terms of how fast a variable is changing. The amount of voltage induced across an inductor is proportional to how *quickly* the current through it changes, not how strong the current is. Likewise, the amount of current “through” a capacitor is proportional to how *quickly* the voltage across it changes. This is the first hurdle in calculus: to comprehend what a rate of change is, and it is far from obvious. A helpful strategy here is to relate electrical rates of change to other, more common, rates of change such as *speed* (i.e. the rate-of-change of physical position over time).
- **The time-constant equation** – Many students find the time-constant equation difficult because it involves exponents, particularly exponents of Euler's constant e . This exponent is often expressed as a negative quantity, making it even more difficult to understand. The single most popular mathematical mistake I see students make with this equation is failing to properly follow algebraic order of operations. Some students try to overcome this weakness by using calculators which allow parenthetical entries, nesting parentheses in such a way that the calculator performs the proper order of operations. However, if you don't understand order of operations yourself, you will not know where to properly place the parentheses. If you have trouble with algebraic order of operations, there is no solution but to invest the necessary time and learn it!

- Determining voltage and current trajectories** – an important step in analyzing voltages and currents in DC circuits containing energy-storing components is to qualitatively graph the trajectories of those voltages and currents. This begins with determining initial and final values, but continues to show either decay or growth that is inverse-exponential, the values in question always asymptotically approaching some final value over time. A very common misconception is for students to plot these graphs with the wrong shape, showing voltages and/or currents changing slowly at first and then more rapidly as time progresses. This is a sign of attempting to memorize shapes (and recalling them incorrectly) rather than thinking through how inverse-exponential growth and decay actually works from a mathematical perspective. A gentle way to approach this is to reflect on half-life decay which is based on negative powers of two (i.e. after one half-life there is $\frac{1}{2}$ the original quantity left, after two half-lives there is $\frac{1}{4}$ remaining, etc.) and then extend that concept to $e^{-\frac{t}{\tau}}$.



- Parasitic effects** – phenomena such as resistance, capacitance, and inductance don't just exist where we want them to, but in fact always exist everywhere they can. All conductors (except for superconductors) have some amount of resistance, all conductors separated by an insulating gap exhibit capacitance, and all conductors possess inductance by virtue of their ability to produce a magnetic field with the passage of current. Often these "parasitic" effects are small enough to ignore, but in some applications they may be severe.

1.3 Recommendations for instructors

This section lists realistic student learning outcomes supported by the content of the module as well as suggested means of assessing (measuring) student learning. The outcomes state what learners should be able to do, and the assessments are specific challenges to prove students have learned.

- **Outcome** – Demonstrate effective technical reading and writing
 - Assessment – Students present their outlines of this module’s instructional chapters (e.g. Case Tutorial, Tutorial, Historical References, etc.) ideally as an entry to a larger Journal document chronicling their learning. These outlines should exhibit good-faith effort at summarizing major concepts explained in the text.
 - Assessment – Students show how inverse-exponential quantitative results were obtained by the author in the Tutorial chapter’s examples.
- **Outcome** – Apply the concept of inverse exponential functions to inductor charging and discharging
 - Assessment – Calculate voltages and currents at specified times for a given resistor-inductor network.
 - Assessment – Calculate times when specified voltages and currents will reach specified values for a given resistor-inductor network.
 - Note – The importance of having students work through all their own calculations in class cannot be overstated!
- **Outcome** – Demonstrate how to create inductance
 - Assessment – Conduct an experiment showing how to create different amounts of inductance with materials such as enameled “magnet” wire, steel bolts, etc. using an LRC meter or other instrument capable of inductance measurement.
- **Outcome** – Prove the concept of inductive energy storage by experiment
 - Assessment – Energize an inductor and then transfer that energy to a load such as an LED, judging the amount of energy by the brightness and duration of the LED’s illumination.
 - Assessment – Re-run these experiments with multiple inductors connected in series versus parallel to compare their relative energy-storage abilities.
- **Outcome** – Independent research
 - Assessment – Locate inductor datasheets and properly interpret some of the information contained in those documents including inductance tolerance, current ratings, saturation limits, etc.

Chapter 2


Case Tutorial

The idea behind a *Case Tutorial* is to explore new concepts by way of example. In this chapter you will read less presentation of theory compared to other Tutorial chapters, but by close observation and comparison of the given examples be able to discern patterns and principles much the same way as a scientific experimenter. Hopefully you will find these cases illuminating, and a good supplement to text-based tutorials.

These examples also serve well as challenges following your reading of the other Tutorial(s) in this module – can you explain *why* the circuits behave as they do?

2.1 Example: superconducting energy-storage inductor

In a US patent by inventor Gary Van Stephenson, a toroidal inductor wound with a coil made of superconducting wire (i.e. having zero electrical resistance) is used to store energy. Figures 3A, 3B, and 3C in this patent show the inductor's magnetic field and current over time as it is receiving energy from a source (3A), idly storing energy previously received (3B), and delivering energy to a load (3C), shown here from left to right:




Note the graphs of current (I) and magnetic field (B) plotted over time for each of these conditions, as well as the existence of an external source or external load required to alter these quantities over time.


2.2 Example: inductor storing and releasing energy

It is easy to use a battery, inductor¹, resistor, and LED to demonstrate the principle of inductance storing and releasing energy:

Inductor storing energy



Inductor releasing energy

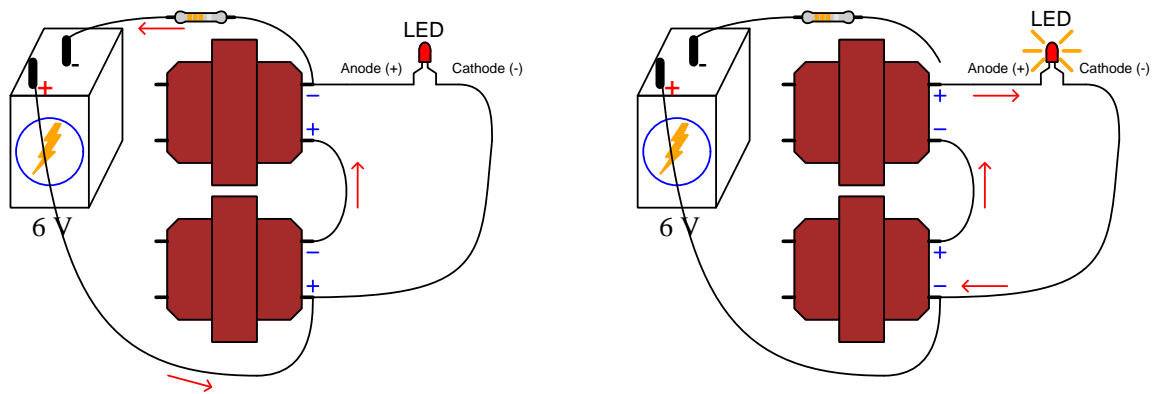


In the upper configuration the inductor acts as a load. The LED does not illuminate here because the polarity is backwards. In the lower configuration (lifting one wire) the inductor acts as a source, which energizes the LED because the inductor's voltage polarity reverses. The purpose of the resistor is to limit current through the inductor to a value that is safe for the LED, since the LED's initial current will be equal to the inductor's final current. This resistor value will need to be customized for your experiment, so that the series combination of this resistor and the inductor's winding resistance together limit the energization current to an LED-safe value.


¹The high-voltage winding of a step-down power transformer generally works well for this inductance. If the transformer has dual-voltage capability (e.g. 120/240 Volts), wire it for the highest-voltage configuration in order to have the greatest amount of inductance.

A recommended variation of this experiment is to try using multiple inductors of the same rating as the original, energized by the same source current as well. The point of these experiments is to compare how long the LED remains energized when powered by the series inductor network, versus when powered by the parallel inductor network, versus when powered by single inductor, total current being equal in all cases:

Series inductors



Parallel inductors



Since inductance is defined as the capacity to store energy in the form of a magnetic field, the comparative results of these experiments will prove whether inductance is enhanced in parallel or in series. The configuration resulting in the LED remaining energized longest is the one having the most inductance.

An important principle in scientific experimentation is to *control* for all variables in an experiment except the one variable we wish to explore. This is why it is important here to use identically-valued inductors in all cases, and use the same amount of source current too. Fixing the source current while using a *voltage* source (e.g. battery) and experiencing different amounts of inductor winding

resistance in each case requires the resistor value be customized for each iteration of the experiment, using an ammeter to measure source and adjusting the resistor value in each version of the circuit to achieve the same total current for each. This way, the only variable is the total amount of inductance resulting from the number of inductors and the manner in which they interconnect with each other.

2.3 Example: series versus parallel inductors

Calculating energy stored within a series network of 20 mH inductors based on the number of inductors (n) connected in series, in each case energized by a 5 Ampere source. The amount of energy stored within each inductor is given by the following formula, and the total energy stored is the simple sum of all inductors' stored energy:


$$E_L = \frac{1}{2}LI^2$$

Where,

E_L = inductor energy in Joules (J)

L = inductance in Henrys (H)

I = current in Amperes (A)



Number of inductors (n)	Current through each inductor	Energy stored in each inductor	Energy stored (total)
1	5.000 A	250.00 mJ	250.00 mJ
2	5.000 A	250.00 mJ	500.00 mJ
3	5.000 A	250.00 mJ	750.00 mJ
4	5.000 A	250.00 mJ	1.0000 J

Connecting more inductors in series results in a direct increase in stored energy, because each inductor receives the same source current (5 Amperes) and therefore stores the same amount of energy.

Since inductance is defined as the ability to store energy within a magnetic field, we must conclude that connecting inductors in series results in *more* over-all inductance.

Calculating energy stored within a parallel network of 20 mH inductors based on the number of inductors (n) connected in parallel, in each case energized by a 5 Ampere source. The amount of energy stored within each inductor is given by the following formula, and the total energy stored is the simple sum of all inductors' stored energy:


$$E_L = \frac{1}{2}LI^2$$

Where,

E_L = inductor energy in Joules (J)

L = inductance in Henrys (H)

I = current in Amperes (A)




Number of inductors (n)	Current through each inductor	Energy stored in each inductor	Energy stored (total)
1	5.000 A	250.00 mJ	250.00 mJ
2	2.500 A	62.500 mJ	125.00 mJ
3	1.667 A	27.778 mJ	83.333 mJ
4	1.250 A	15.625 mJ	62.500 mJ

Connecting more inductors in parallel results in the source current (5 Amperes) being divided into smaller portions. With less current through each inductor, there is dramatically less energy stored within each owing to the squared current term in the $E_L = \frac{1}{2}LI^2$ equation, so that with more inductors in the parallel network there is actually less total energy stored for the entire network.

Since inductance is defined as the ability to store energy within a magnetic field, we must conclude that connecting inductors in parallel results in *less* over-all inductance.

2.4 Example: current source energizing an inductor

Assume the inductor begins in a completely de-energized state. The switch begins in a closed state and is opened at time $t = 0$.



Switch opening time	I_L	I_R	V
-1 seconds	0 mA	0 mA	0 V
0 seconds	0 mA	6 A	21 V
1 second	1.7719 A	4.2281 A	14.798 V
2 seconds	3.0205 A	2.9795 A	10.428 V
3 seconds	3.9004 A	2.0996 A	7.349 V
4 seconds	4.5204 A	1.4796 A	5.179 V
5 seconds	4.9574 A	1.0426 A	3.649 V
6 seconds	5.2653 A	737.74 mA	2.572 V
7 seconds	5.4822 A	517.76 mA	1.812 V
8 seconds	5.6351 A	364.86 mA	1.277 V
∞	6 A	0 mA	0 V

After the switch opens the inductor is free to accept energy from the current source and so its magnetic field increases in strength and so does its current. Thus, I_L follows the inverse-exponential growth equation from 0 Amperes at $t = 0$ building up to a maximum of 6 Amperes at infinite time:

$$I_L = (6 \text{ A}) \left(1 - e^{-\frac{t}{\tau}}\right)$$

Once inductor current is calculated for each time value, all other circuit parameters may be calculated from those I_L values following Kirchhoff's Current Law and Ohm's Law.

Here is a step-by-step conceptual analysis of this circuit:


1. **Zero inductor energy** before the switch is opened, stated as a given condition.
2. **Zero circuit voltage** before the switch is opened, since shorts guarantee zero voltage and this is a parallel circuit which means voltage must be equal for all components.
3. **Zero inductor energy** at the precise moment the switch opens, since energy transfer takes time and so far no time has yet elapsed with the switch in the open state. Again, this means zero inductor current at the precise moment of switch opening.
4. **Full resistor current** (6 Amperes) at the precise moment of switch opening, as per Kirchhoff's Current Law, flowing downward through the resistor.
5. **Full resistor voltage** (21 Volts) at the precise moment of switch opening, as per Ohm's Law, positive on top and negative on bottom across the resistor.
6. **Full voltage across all components** (21 Volts positive top and negative bottom) at the precise moment of switch opening, since this is a parallel circuit and voltage must be equal across all parallel-connected components.
7. **Increasing inductor energy** over time as the source imposes a voltage across it and transfers energy into it. Since inductors store energy in magnetic fields, and magnetic fields directly relate to current, we may conclude that the inductor's current must increase as well over time, flowing downward through it because the inductor is acting as a load now.
8. **Decreasing resistor current** over time as the inductor's current increases, as per Kirchhoff's Current Law.
9. **Decreasing resistor voltage** over time as the resistor's current decreases, as per Ohm's Law.
10. **Decreasing circuit voltage** over time, being a parallel circuit where all components share the same voltage value at each point in time.
11. **Full inductor current** after the switch has been open a long while and the entirety of the source current passes through it (6 Amperes), following the increasing trend established after the switch opened. This means the inductor now possesses the maximum amount of energy it can for this circuit, since inductor energy is directly related to magnetic field strength and inductor current.
12. **Zero resistor current** after the switch has been open a long while, as per Kirchhoff's Current Law with the inductor passing the entire source current.
13. **Zero resistor voltage** after the switch has been open a long while, as per Ohm's Law.
14. **Zero circuit voltage** after the switch has been open a long while, all components being connected in parallel and therefore sharing the same voltage.

15. *SUMMARY:* inductor current grows from zero to 6 Amperes when the shorting switch opens; resistor voltage spikes to 21 Volts when the switch opens and then decays to zero over time; circuit voltage spikes to 21 Volts when the shorting switch opens and then decays to zero over time. *We are now ready to commence inverse-exponential calculations of voltages and currents over time!*

Many students have a bad tendency to skip these conceptual steps in lieu of jumping immediately to the inverse-exponential formulae which they know will ultimately give them the answers they seek for voltages and currents at certain times. Remember, it is important to always let the physical concepts guide the application of mathematics, lest you fall into the common error of blindly plugging numbers into equations in the hope of getting the right answers! Your schematic diagram should be fully annotated with current arrows and voltage polarities before any calculations are performed.

2.5 Example: voltage source energizing an inductor

Note: these calculations all assume an ideal inductor having no internal resistance, as well as perfect wire connections having no resistance either. The switch begins in an open state and is closed at time $t = 0$.



Switch closure time	V_L	V_R	I
-1 seconds	0 V	0 V	0 mA
0 seconds	12 V	0 V	0 mA
1 second	10.2344 V	1.7656 V	882.8 mA
2 seconds	8.7285 V	3.2715 V	1.636 A
3 seconds	7.4442 V	4.5558 V	2.278 A
4 seconds	6.3489 V	5.6511 V	2.826 A
5 seconds	5.4148 V	6.5852 V	3.293 A
6 seconds	4.6181 V	7.3819 V	3.691 A
7 seconds	3.9386 V	8.0614 V	4.031 A
8 seconds	3.3591 V	8.6409 V	4.320 A
∞	0 V	12 V	6 A

The switch's open contact status prior to and at time $t = 0$ ensures zero current through the inductor, which means it will have no magnetic field within and will store no energy. After the switch closes the inductor is free to accept energy from the voltage source and so its magnetic field increases in strength and so does its current. Thus, I_L follows the inverse-exponential growth equation from 0 Amperes at $t = 0$ building up to a maximum of 6 Amperes at infinite time:

$$I_L = (6 \text{ A}) \left(1 - e^{-\frac{t}{\tau}}\right)$$

Once inductor current is calculated for each time value, all other circuit parameters may be calculated from those I_L values following Kirchhoff's Voltage Law and Ohm's Law.

Here is a step-by-step conceptual analysis of this circuit:


1. **Zero circuit current** before the switch is opened, since opens guarantee zero current and this is a series circuit which means current must be equal for all components.
2. **Zero inductor energy** before the switch is closed, since inductors store energy in magnetic fields, and magnetic fields directly relate to current which we know at this time is zero.
3. **Zero inductor energy** at the precise moment of switch closure, since energy transfer takes time and so far no time has yet elapsed with the switch in the closed state. Again, this means zero inductor current at the precise moment of switch closure.
4. **Zero current through all components** at the precise moment of switch closure, since this is a series circuit and current must be equal through all series-connected components.
5. **Zero resistor voltage** at the precise moment of switch closure, as per Ohm's Law.
6. **Full inductor voltage** (12 Volts) at the precise moment of switch closure, as per Kirchhoff's Voltage Law, positive on top and negative on bottom across the inductor.
7. **Increasing inductor energy** over time as the source imposes a voltage across it and transfers energy into it. Since inductors store energy in magnetic fields, and magnetic fields directly relate to current, we may conclude that the inductor's current must increase as well over time, flowing downward through it because the inductor is acting as a load now.
8. **Increasing resistor current** over time as the inductor's current increases, since they are connected in series with each other and series-connected components must share the same current.
9. **Increasing resistor voltage** over time as the resistor's current increases, as per Ohm's Law.
10. **Decreasing inductor voltage** over time, as per Kirchhoff's Voltage Law.
11. **Zero inductor voltage** after the switch has been closed a long while, following the decreasing trend established after the switch closed.
12. **Full resistor voltage** after the switch has been closed a long while, as per Kirchhoff's Voltage Law with the inductor dropping zero voltage.
13. **Full resistor current** after the switch has been closed a long while, as per Ohm's Law (6 Amperes, given 12 Volts dropped across the 2 Ohm resistor).
14. **Full inductor energy** after the switch has been closed a long while, since inductor energy is directly related to magnetic field strength and inductor current.
15. *SUMMARY: inductor current grows from zero to 6 Amperes when the switch closes; resistor voltage grows from zero to 12 Volts over time; inductor voltage spikes to 12 Volts when the switch closes and then decays to zero over time. We are now ready to commence inverse-exponential calculations of voltages and currents over time!*

Many students have a bad tendency to skip these conceptual steps in lieu of jumping immediately to the inverse-exponential formulae which they know will ultimately give them the answers they seek for voltages and currents at certain times. Remember, it is important to always let the physical concepts guide the application of mathematics, lest you fall into the common error of blindly plugging numbers into equations in the hope of getting the right answers! Your schematic diagram should be fully annotated with current arrows and voltage polarities before any calculations are performed.

2.6 Example: LR time-delay network values

Example #1

Determine the trajectories of inductor voltage and current in this circuit once the switch toggles position ($t = 0$), assuming the switch has been in the position shown for a very long time prior to toggling, as well as values for V_L and I_L at $t = 20$ microseconds, $t = 30$ microseconds, and $t = 40$ microseconds.



Before the switch toggles ($t < 0$) the inductor is connected in series with the 10 Volt source through the 2.7 kΩ resistor. After the inductor has absorbed all the energy it will from the voltage source, the inductor will neither behave as a source nor as a load but will simply pass current with no voltage drop. This current is determined by Ohm's Law, being 3.7037 mA.


At the moment of switch toggling ($t = 0$) the inductor still has the same amount of current passing through it in the same direction as it did before because inductors store energy in magnetic fields (manifest as current), and there has been no time for any energy to transfer out of it into the 6.8 kΩ load resistor, and energy cannot simply disappear into nothing. By Ohm's Law this means the 6.8 kΩ resistor will initially develop 25.185 Volts across it. Note that this voltage, which is greater than the DC source's, is possible because the conserved energy within an inductor at the moment of switch toggling manifests as *current* which was defined by the DC source voltage and a different (lesser) resistance of 2.7 kΩ.

As the switch remains in its new position ($t > 0$) the inductor loses energy as it powers the 6.8 kΩ resistor, eventually depleting its stored energy as the resistor dissipates it in the form of heat. This means inductor current must decay to zero over time, and with that the resistor's voltage will also decay to zero. Once voltage reaches its final value of zero, the inductor will no longer be a load (nor a source) and its energy value and current value will remain constant at zero.

- At $t = 20$ microseconds: $V_L = 10.171$ Volts and $I_L = 1.4958$ milliAmperes
- At $t = 30$ microseconds: $V_L = 6.4640$ Volts and $I_L = 0.95060$ milliAmperes
- At $t = 40$ microseconds: $V_L = 4.1079$ Volts and $I_L = 0.60411$ milliAmperes

Example #2

Determine the trajectories of inductor voltage and current in this circuit once the switch closes ($t = 0$), assuming the switch has been open for a very long time prior, as well as values for V_L and I_L at $t = 100$ microseconds, $t = 150$ microseconds, and $t = 200$ microseconds.



Before the switch closes ($t < 0$) the inductor completely de-energized. We know this because inductors store energy magnetically, manifesting as current, and the open switch prohibits current everywhere in this series circuit.


At the moment of switch closure ($t = 0$) the inductor will still have no current passing through it because there has been no time for any energy to transfer into it from the 5 Volt DC source, and energy cannot simply appear from nothing. Zero current everywhere in this series circuit means neither resistor drops any voltage (by Ohm's Law), and using Kirchhoff's Voltage Law we must conclude the entirety of the DC source's voltage must drop across the inductor.

As the switch remains closed ($t > 0$) the inductor gains energy from the source, gradually increasing current. When this growing current develops a combined total of 5 Volts' worth of drop across the two resistors, Kirchhoff's Voltage Law predicts the inductor's voltage will reach zero. At this point the inductor behaves as neither a source nor a load, and energy transfer ceases. Current will stabilize at this maximum value defined by the 5 Volts DC source and the two series resistors which is 3.7594 mA.

- At $t = 100$ microseconds: $V_L = 2.7316$ Volts and $I_L = 1.7056$ milliAmperes
- At $t = 150$ microseconds: $V_L = 2.0190$ Volts and $I_L = 2.2413$ milliAmperes
- At $t = 200$ microseconds: $V_L = 1.4923$ Volts and $I_L = 2.6373$ milliAmperes

Example #3

Determine the trajectories of inductor voltage and current in this circuit once the switch opens ($t = 0$), assuming the switch has been closed for a very long time prior, as well as values for V_L and I_L at $t = 300$ microseconds, $t = 500$ microseconds, and $t = 1$ millisecond.



Before the switch opens ($t < 0$) the inductors should collectively have no current through them, having had ample time to exhaust their stored energy in circuit wiring resistance. Voltage will also be zero in this de-energized state given the closed status of switch preventing any voltage drop.

At the moment of switch opening ($t = 0$) the inductors will still have no current through them because there has been no time for any energy to transfer into them from the 50 milliAmpere source, and we know stored energy in any inductor manifests as current through its terminals. With the switch in its new position, though, the inductors will be able to absorb energy from the source over time. Kirchhoff's Current Law tells us the two resistors (taken together as a 57Ω series equivalent resistance) must carry the full current of the 50 mA source, and by Ohm's Law this must yield a voltage drop of 2.85 Volts.

As the switch remains open ($t > 0$) the inductors will act as loads absorbing energy from the 50 mA source, which means both inductors' currents will grow over time until they collectively (in parallel) reach 50 milliAmperes. At that point the resistors will also have no current through them (by Kirchhoff's Current Law) and this means voltage will ultimately decay down to zero.

It is helpful to simplify this circuit by considering it as one with a single inductor (42.857 mH) and a single resistor (57Ω), having a time constant of $t = 751.88$ microseconds.

- At $t = 300$ microseconds: $V_L = 1.9123$ Volts and $I_L = 16.450$ milliAmperes
- At $t = 500$ microseconds: $V_L = 1.4657$ Volts and $I_L = 24.286$ milliAmperes
- At $t = 1$ millisecond: $V_L = 0.75376$ Volts and $I_L = 36.776$ milliAmperes


2.7 Example: simplifying complex LR networks using Norton's Theorem

Determining the time constant (τ) and initial/final voltages for an elementary LR network is usually a straightforward process, especially time constant which is simply the quotient $\tau = \frac{L}{R}$. However, for networks containing multiple resistors it may not be clear upon first inspection exactly how those resistors interact with the inductor. In some multi-resistor networks *every* resistor impacts the time constant, while in others only some do.

An extremely useful tool for this application is Norton's Theorem, applying it to the original LR network in such a way that the inductor is identified as the load, reducing the rest of the original network to a single current source (I_N) and a single equivalent resistance (R_N). That Norton equivalent current limits the inductor's maximum current at any time (assuming the inductor was not pre-charged to some greater current value prior to connection to the network), and that Norton resistance combined with the "load" inductance directly defines time constant ($\tau = \frac{L}{R_N}$).


The following LR network examples showcase the utility of Norton's Theorem to help make comprehensible what might otherwise defy intuition. You, as the reader, are strongly encouraged to Nortonize each of these example networks on your own in order to see for yourself how the equations for time constant and maximum current emerge. Simply regard the inductor in each case as the load, remove that load from the network and determine R_N by replacing the source with its ideal internal resistance value (shorts for voltage sources and opens for current sources) and then determining total resistance as measured from the load terminals, and then determine load terminal current with the source activated once again and the load terminals shorted together.

Example #1




$$R_N = R_2 \quad \tau = \frac{L}{R_2} \quad I_N = I$$

In this network the value of R_1 is irrelevant to time constant, and neither resistance value affects maximum inductor current.

Example #2

$$R_N = R_1 + R_2 \quad \tau = \frac{L}{R_1 + R_2} \quad I_N = \frac{I}{R_1 \left(\frac{1}{R_1} + \frac{1}{R_2} \right)}$$

In this network both the values of R_1 and R_2 affect time constant and maximum inductor current. That maximum current (I_N) is derived using the current divider equation, since R_1 and R_2 effectively form a current-division network with current source I from the perspective of the (short-circuited) load terminals.


Example #3

$$R_N = R_1 \quad \tau = \frac{L}{R_1} \quad I_N = \frac{V}{R_1}$$

In this network the value of R_2 is irrelevant both to time constant and to maximum inductor current, while R_1 affects both.

2.7. EXAMPLE: SIMPLIFYING COMPLEX LR NETWORKS USING NORTON'S THEOREM 27

Example #4




$$R_N = \frac{1}{\frac{1}{R_1} + \frac{1}{R_2}} \quad \tau = L \left(\frac{1}{R_1} + \frac{1}{R_2} \right) \quad I_N = \frac{V}{R_1}$$


In this network both the values of R_1 and R_2 affect time constant but only R_1 affects maximum inductor current.

2.8 Example: interesting relay behavior

An electromechanical *relay* is a device using the magnetic field from an electromagnet coil to attract a movable iron piece (called an “armature”) which in turn actuates switch contacts. The following schematic symbols show a relay with a normally-open contact (left) and one with a normally-closed contact (right):

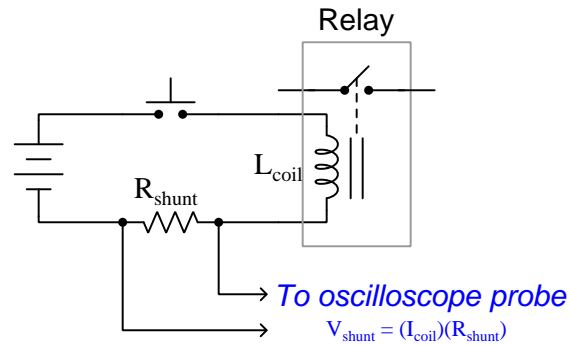


In the following photograph we see a relay with its plastic cover removed:




A coil spring holds the switch contacts in their “normal” (resting) positions, and is overcome by the electromagnet’s attraction when sufficient current passes through the relay’s coil.

If we perform an experiment by connecting a low-resistance *shunt resistor* in series with a relay's coil and using an oscilloscope to measure voltage across that resistor, the oscilloscope's indication will be proportional to the *current* through the relay coil:

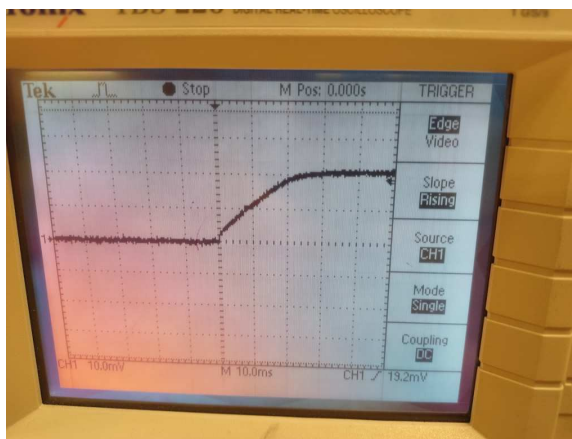


What we might expect upon pressing the pushbutton switch is a classic inverse-exponential rise from zero current to full current due to the inductance (L) of the relay's electromagnet coil, with a time constant (τ) equal to $\frac{L}{R}$, but what we get instead is this:



A normal inverse-exponential growth curve for an ordinary inductor would smoothly rise over time, but this curve appears to have a sharp “dip” mid-way up. What might be the cause of this?

Continuing our experiment, we will now manually hold the relay's armature in the attracted position while pressing the pushbutton switch. Current still passes through the relay coil upon energization, but now the coil will not do any useful work because we will manually do that work ourselves. What we see during this second run of the experiment is a standard inverse-exponential growth curve with no unusual "dip" anywhere to be found:



Clearly, moving the relay's armature by hand affects the rise of current through the coil over time – but why? How can a *mechanical* stimulus impact an *electrical* behavior?

Reviewing foundational principles of magnetic fields and of inductance, we know that inductors store energy in their magnetic fields, and that magnetic field strength is a direct function of current strength ($E_L = \frac{1}{2}LI^2$). From this we may conclude the higher the oscilloscope's trace climbs on the display, the more energy is stored in that magnetic field. Energy, as we know, is the ability to set matter into motion or to perform *work*, and we also know that energy must always be conserved. The act of pulling the relay's armature over a short distance against the restraint of the coil spring represents a definite quantity of work. Thus, during the first experiment when the relay's armature moved in attraction to the coil's magnetic field, the work exerted in that motion and force required energy came from somewhere (because energy is conserved), and the only source of energy for doing that work was the coil's magnetic field. This work entailed a transfer of energy away from the magnetic field to the mechanical spring, and so when the armature moved the magnetic field weakened, which meant the current had to weaken as well, thus explaining the dip.

In the second oscilloscope screenshot where we manually actuated the relay armature, we preemptively performed the work that would normally be done by the relay's coil, and therefore the energy invested in the coil had no outlet. Thus, the relay coil acted as a plain and simple inductor, storing all the energy it received from the battery rather than doing any useful work with that energy. This is why no "dip" appeared during the second experiment.

Another variation on this experiment is to reduce the DC power source voltage to a value below the relay's pick-up voltage. This means the power source will not output enough voltage to produce a strong enough relay coil current to actually prompt armature motion inside the relay. With this reduced voltage, the magnetic field produced by the energized relay coil will not generate an attractive force strong enough to overcome the relay's armature spring, and so the armature will

remain in its “resting” position even with the coil energized.

Doing this will cause the current’s profile over time to increase in a smooth, inverse-exponential fashion much like it did when we manually forced the armature into the energized position. The reason for this is much the same: by configuring the experiment such that the armature will not actually move, there will be no mechanical work done by the energized relay coil, which means all the energy of the magnetic field resides with the coil and does not leave it.

This experiment serves as a fantastic illustration of the power of mastering foundational concepts. If all one ever learned about inductors is that voltage and current trace inverse-exponential trajectories over time, the “dip” phenomenon would be a complete mystery. However, understanding the relationship of an inductor’s current to its magnetic field strength, and in turn the relationship to energy (and its conservation), allows us to understand these “strange” effects. It also goes to show just how universal the Law of Energy Conservation, which of course is a benefit because we may apply this law to all physical phenomena and thereby gain understanding we might not obtain otherwise.

2.9 Example: custom inductor winding and testing

In this example we see a custom inductor made by winding 19 turns of 23-gauge enamel-coated copper wire around a Micrometals T50-66C powdered iron toroidal core. This core’s datasheet specifies an inductance factor (A_L) of 51 nanoHenrys per turn-squared, so the inductor’s value should be:

$$L = A_L N^2$$

$$L = (51 \times 10^{-9})(19^2) = 18.411 \mu\text{H}$$

Micrometal’s general catalog of core materials lists the type 66 as being “well-suited from 100 kHz to 500 kHz” and so a test was conducted using a NanoVNA vector network analyzer over that same range of signal frequencies:




Analysis from the VNA shows the measured inductance to be quite close to the predicted value of 18.411 microHenrys over the “swept” signal range of 100 kHz to 500 kHz:

Frequency	Impedance	Series L	Deviation from predicted
100 kHz	$0.148 + j11.3 \Omega$	$17.974 \mu\text{H}$	-2.37%
300 kHz	$1.11 + j33.5 \Omega$	$17.782 \mu\text{H}$	-3.42%
500 kHz	$2.76 + j56.1 \Omega$	$17.849 \mu\text{H}$	-3.05%

The “real” portion of the complex impedance value represents series equivalent resistance in Ohms, equal in this case to the resistance of the wire in the inductor’s coil as well as any connection resistances between the inductor and the VNA. This value rises slightly with signal frequency due to the *skin effect*, where electric current at very high frequency becomes excluded from traveling in the wire’s center and therefore decreases the effective cross-sectional area of the conductor. The “imaginary” portion of the complex impedance value is the inductive reactance (X_L) in Ohms, relating to inductance by the formula $X_L = 2\pi fL$ where f is the signal frequency in Hertz.


Like most vector network analyzers, the NanoVNA produces a Smith chart plot of complex impedance (normalized to 50 Ohms). As expected for a nearly-pure inductance, the path traced during the 100-500 kHz “sweep” closely follows the 0 Ohm resistive circle circumference while intersecting various inductive reactance circles as frequency rises. Here the red marker signifies 100 kHz (just above the +j0.2 circle denoting +j10 Ohms), the green marker 300 kHz (between the +j0.5 and +j1.0 circles), and the blue marker 500 kHz (just above the +j1.0 circle denoting +j50 Ohms):



Next we see the results of winding 30 turns of 23-gauge wire around another T50-66C powdered iron core with the same inductance factor of 51 nanoHenrys per turn-squared:

$$L = A_L N^2$$

$$L = (51 \times 10^{-9})(30^2) = 45.900 \mu\text{H}$$



Analysis from the VNA shows the measured inductance to deviate more significantly from the predicted value of 45.9 microHenrys over the “swept” signal range of 100 kHz to 500 kHz than it did with the previous (19-turn) inductor:


Frequency	Impedance	Series L	Deviation from predicted
100 kHz	$0.827 + j26.8 \Omega$	$42.665 \mu\text{H}$	-7.048%
300 kHz	$3.86 + j80.6 \Omega$	$42.771 \mu\text{H}$	-6.817%
500 kHz	$6.29 + j135 \Omega$	$42.941 \mu\text{H}$	-6.447%

A greater number of wire turns means an increase in two different parasitic properties: first, the wire will have greater total resistance because there is simply more of it for current to pass through; second, the greater number of turns packed onto the same size toroid means individual turns being closer together to one another than before which will increase the amount of inter-turn capacitance².

²Recall that capacitance between any two conductive surfaces is a direct function of the overlapping area of those surfaces, and an inverse function of separation distance. Here, with 30 turns of wire being packed closer together on the same toroid than 19 turns of wire, the inter-turn capacitance must be greater.

Here again we see the NanoVNA's Smith chart showing complex impedance normalized to 50 Ohms, this time for the 30-turn inductor:

S11 Smith Chart




As one would expect, the curve lies farther to the right due to the increased inductive reactance of the larger inductance over the same frequency range. Increased resistance from the longer wire drives the curve closer to the middle of the Smith chart rather than remain on the periphery as before.

Using an entirely different toroidal core, I performed an experiment where I tried different numbers of turns of the same 23-gauge enameled wire and measured inductance using a different VNA (a PicoVNA model 106) sweeping from 300 kHz to 400 kHz:

Turns	300 kHz	354 kHz	381 kHz	400 kHz
2	6.06 μH	6.31 μH	6.43 μH	6.52 μH
4	22.79 μH	23.70 μH	24.13 μH	24.45 μH
5	36.36 μH	37.82 μH	38.51 μH	38.93 μH
10	149.63 μH	156.85 μH	159.24 μH	160.16 μH

The following photograph shows this iron core with four turns of wire wrapped around it, the bare iron core itself being wrapped in white electrical tape to prevent the enameled wire from chafing against the bare metal and possibly creating a coil-to-core shorted fault:



If you take any two inductance values at the same test frequency but for a different number of turns, the ratio of those two inductance values approximately equals the square of the ratio of the turns. For example, 160.16 microHenrys and 38.93 microHenrys at 400 kHz for 10 turns and 5 turns, respectively, makes an inductance ratio of nearly four to one (4.11405:1) for a turns ratio that is two to one (2:1). This confirms the general principle that inductance is proportional to the square of the number of wire turns wrapped around the core.

Computing the average A_L value at each test frequency based on the four measured inductance values and turns values ($\frac{L}{N^2}$) yields the following:

Frequency	Average A_L
300 kHz	1.473 μH per turn ²
354 kHz	1.535 μH per turn ²
381 kHz	1.562 μH per turn ²
400 kHz	1.579 μH per turn ²

Averaging these four A_L values yields an aggregate A_L of 1.537 μH per turn-squared for this particular toroidal iron core. This figure, of course, will be useful for creating custom inductors based on this toroid as the core material.

2.10 Example: measuring signal rates of change

Capacitors and inductors relate voltage to current by rates of change. For a capacitor, the amount of current is proportional to how quickly voltage across that capacitor either rises or falls over time ($I = C \frac{dV}{dt}$). For an inductor, the amount of voltage is proportional to how quickly current through that inductor rises or falls over time ($V = L \frac{dI}{dt}$).

For example, a 330 microFarad capacitor experiencing a voltage increasing at a rate of 45 Volts per second will pass 14.85 milliAmperes. If the voltage happens to decrease at an equivalent rate (i.e. $\frac{dV}{dt} = -45$ Volts per second) then the 14.85 milliAmpere current will reverse direction through the capacitor compared to how it flowed with the increasing voltage.

Similarly, a 100 milliHenry inductor experiencing a current increasing at a rate of 5 Amperes per second will induce a voltage of 500 milliVolts. If the current happens to decrease at an equivalent rate (i.e. $\frac{dI}{dt} = -5$ Amperes per second) then the 500 milliVolt voltage induced across the inductor will reverse polarity from what it was during the period of increasing current.


Not only are rates-of-change important in determining how energy-storing devices such as capacitors and inductors will respond in circuits, but rates-of-change are also important for determining how *parasitic* capacitances and inductances will affect intended circuit behavior. Parasitic capacitance exists between any two conducting surfaces separated by an electrically insulating medium, and parasitic inductance exists along any length of conductor. This means any rate-of-change of voltage over time between two separated conductors will cause some amount of current to “pass” between them, and that any rate-of-change of current over time through any single conductor will cause some amount of voltage to drop across its length. In many circuits these parasitic effects are negligible, but in circuits experiencing extremely fast rates of change for voltage and/or current the effects can be significant or even severe.

Oscilloscopes are ideal for performing empirical measurements of voltage rates-of-change, and of current rates-of-change given the proper accessories³. Some skill is required to do this, though, and here we will explore practical examples to show how it is done.


For any signal plotted in the time domain, where the horizontal axis of the plot is expressed in units of seconds, milliseconds, microseconds, etc., the signal’s rate of change at any given point will be the *slope* or *pitch* of the waveform, mathematically defined as its *rise over run*. A great aid to discerning slope at any location on a waveform is to sketch a straight line visually matching the wave’s slope at that point, then use locations along that straight line to more easily discern how far it rises (or falls) over some “run” of time. We call this straight line a *tangent line*.

³For example, a *current probe* converting a sensed current into a voltage the oscilloscope may directly sense, or a *shunt resistor* placed in the circuit developing an oscilloscope-measurable voltage drop for any current passing through.

Here we see an example of a waveform with sloping sections. In the first image we see a specific location on the waveform where we wish to measure voltage rate-of-change ($\frac{dV}{dt}$):




Next we see a tangent line drawn to match the slope of the waveform at the specified location, with convenient points for fall/run measurements taken on that line against the oscilloscope grid's major divisions. In this example, the oscilloscope's vertical sensitivity has been set for 0.5 Volts per division, and the horizontal timebase for 0.2 milliseconds per division:



As we can see, the tangent line falls 2 vertical divisions (-1 Volt) over a timespan of 8 horizontal divisions (1.6 milliseconds), yielding a $\frac{dV}{dt}$ quotient of -625 Volts per second, which may also be expressed as -0.625 Volts per millisecond. The negative sign is important, as it distinguishes this particular rate-of-change as *falling* rather than rising over time.

Such rate-of-change measurements are necessarily approximate, as they require us to visually gauge where a tangent line may be overlaid on the waveform's oscillograph, and also to visually assess the slope of that tangent line using the grid provided on the instrument's display screen. However, in most applications extremely precise rate-of-changes are not necessary, and such techniques suffice quite well.

Below is another example of a waveform with sloping sections, the oscilloscope configured for 500 milliVolts per division on the vertical axis and 5 milliseconds per division on the horizontal:




Approximating the slope for each rising section of this wave, we count one division of rise over 2.8 divisions of run, or 500 milliVolts rise over 14 milliseconds of run. This is a rate-of-rise of +35.7 Volts per second.

Approximating the slope for each falling section of this wave, we count one division of fall over 1.4 divisions of run, or -500 milliVolts fall over 7 milliseconds of run. This is a rate-of-fall of -71.4 Volts per second.

If these rates-of-change appear suspiciously large compared to the actual amplitude of the waveform, which barely crests over +1 Volt on the oscillograph, bear in mind that we are calculating *rates of change* for voltage and not absolute values of voltage itself. This is analogous to the distinction between *speed* and *distance*: traveling at a rate of 30 kilometers per hour does not necessarily mean you will travel 30 kilometers, as the actual distance traveled depends on how long that speed is sustained. A voltage rising at a rate of 35.7 Volts per second would indeed rise 35.7 Volts *if given a full second to do so*, but since each rising/falling portion of this waveform is so short in duration the actual amount of rise or fall in each case is only one-half of one Volt. It is therefore perfectly appropriate to consider any $\frac{dV}{dt}$ value as being the *speed* at which a voltage increases or decreases over time, distinct from the actual value of that voltage at any particular moment in time.

Here we see another oscillograph, this one zoomed into the rising edge of a square wave. For this measurement the oscilloscope was configured for 20 milliVolts per division of vertical sensitivity and a timebase of 250 nanoseconds per division on the horizontal:



The tangent line overlaid on this screenshot for the purpose of measuring the pulse edge's rate-of-change rises approximately 4 vertical divisions over a run of 1 division, which is 80 milliVolts of rise over 250 nanoseconds of run. The pulse edge's rate-of-change, therefore, is approximately +320,000 Volts per second, or +320 Volts per millisecond, or +0.32 Volts per microsecond (all equivalent expressions of $\frac{dV}{dt}$).


Chapter 3

Tutorial

3.1 Magnetic fields and energy storage


A *magnetic field* is an invisible web of interaction between moving electrical charges, able to exert forces against those charges when separated from each other over space. Magnetic fields differ from electric fields in that they only influence *moving* charges and not stationary charges. Electric fields influence electric charges regardless of velocity.

Any conductor passing an electric current exhibits an encircling magnetic field perpendicular to the direction of that current. If the conductor is bent into the shape of a coil, these encircling fields will link together to form a larger magnetic parallel to the coil's centerline:



Just as any moving mass possesses kinetic energy, moving electric charge carriers creating their own magnetic field also possess electrical energy related to their motion. That is to say, the magnetic field surrounding a current-carrying conductor causes those moving charge carriers to exhibit a kind of *inertia* such that an investment of energy is necessary to speed up their drift and a divestment of energy is necessary to make them slow down. This is the basis of *inductance*: the storing of energy within a magnetic field. The amount of energy stored by any inductance is a function of current: all other factors being equal, more current passing through a conductor means more energy stored in the magnetic field.

Electrical components called *inductors*¹ exploit this phenomenon, constructed of a single conductor coiled around some core, the core usually being made of alloy steel or some other ferromagnetic substance. The amount of inductance exhibited by any inductor is directly proportional to the cross-sectional area of the coil (in square meters), directly proportional to the square of the number of turns of wire in the coil, directly proportional to the magnetic permeability² (μ) of the core material, and inversely proportional to the axial length of the coil (in meters):



Energy stored in any inductance is a function of the inductance and the amount of current passing through it as described below:

$$E_L = \frac{1}{2}LI^2$$

Where,

E_L = energy in Joules (J)

L = inductance in Henrys (H)

I = current in Amperes (A)

Note the similarities between this equation and those describing potential energy stored within a mechanical spring ($E_p = \frac{1}{2}kx^2$) and kinetic energy stored within a moving mass ($E_k = \frac{1}{2}mv^2$). Just as kinetic energy varies with the square of the mass's velocity (v), energy stored within an inductor exhibits a kinetic character³ as it varies with the square of the current through it.

¹Also referred to as *reactors* in high-power electrical circuits, and sometimes as *chokes* in electronic circuits.

²Permeability, simply defined, is a measure of how effective a substance is at magnetically storing energy for any given amount of magnetomotive force (MMF). It may be measured in absolute terms, expressed in units of Henrys per meter, or it may be expressed as a ratio relative to the permeability of a perfect vacuum.



³In most applications the actual kinetic energy of the moving charge carriers is inconsequentially small due to their extremely slow drift velocity, but since the magnetic field and its stored energy only occurs with charge carrier *motion* it is not inappropriate to regard the energy stored in the magnetic field of an inductance as being kinetic rather than potential. The inductance's charge carriers, once in motion, tend to remain in motion and will not speed up or slow down without work being done.

One way to make a *variable* inductor is to make the ferromagnetic core movable in relation to the wire coil. In the example below we see a variable inductor with a screwdriver-adjustable core:




Another way to make a variable inductance is to alter the physical orientations of two or more adjacent inductors wired in series and/or parallel with each other. Those individual coils' magnetic fields interfere either constructively or destructively based on their relative orientations, resulting in a different total inductance value for each position.

A clever implementation of this concept is seen in the following photographs, where an AC generator having both stationary and rotating coils is wired to form a single inductance, that inductance value depending on the interconnections of the coils as well as the shaft's angle. On the left we see an explanatory schematic showing how the AC generator's coils were re-wired to serve as a variable inductor, and on the right we see a label affixed to the generator's side showing different series/parallel configurations with measured inductance values:



Energy is always conserved, which means it must always be accounted for in any system and cannot simply appear or vanish. Another way of stating this same fact is to say that the only way for the amount of stored energy in any system to vary is if energy *transfers* to or from some other place⁴. Since energy is a function of current for any inductance, the amount of current through a fixed inductor's coil cannot vary unless energy is added to or removed from that inductance. That energy transfer occurs when voltage appears between the terminals of an inductor. **In other words, given that energy is always conserved and that stored energy within an inductor manifests as current, it must be current that is conserved while voltage changes polarity based on the accumulation or release of energy.**



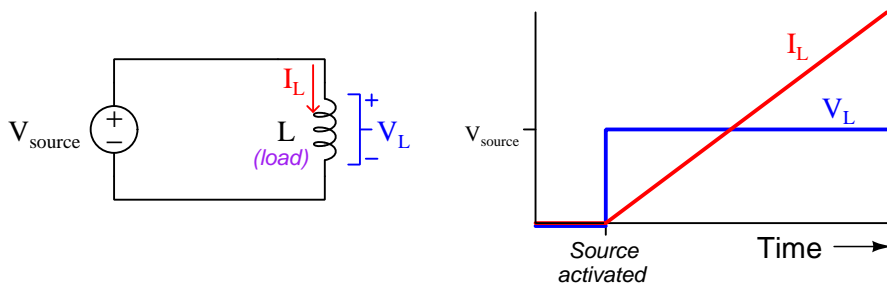
Thus, an inductance maintains a steady store of energy when its current is constant and its voltage⁵ is zero. In this condition of steady current and zero voltage the inductance simply behaves as a *short*. However, if current increases through an inductance it means its stored energy must also increase, which in turn means the inductance must act as a load as it absorbs energy from some external source. Conversely, if current decreases through an inductor it means its stored energy must also decrease, which in turn means the inductance must act as a *source* as it releases energy to some external load.

⁴A close analogy based on the Conservation of Mass is that of a container holding water. Since water has mass, the only way for the amount of water in that container to vary is if water leaves it or enters it. A container's water mass cannot simply vary without any cause, because mass is a conserved quantity.

⁵A true zero-voltage condition with finite current is possible only when the inductor's coil is made of superconducting wire. Such inductors do indeed exist and have been used for practical energy-storage applications. Non-superconducting inductors possess some non-zero resistance in their wire coils, and so will drop voltage in accordance with Ohm's Law ($V = IR$) in addition to voltage occurring as a result of the inductor absorbing or releasing energy.

3.2 Voltage, current, energy, and inductance

If we connect a DC voltage source to a pure⁶ inductor, we find that the current through the inductor rises at a steady rate determined by the amount of inductance and the amount of voltage. Inductor current must begin at zero because inductors store energy by means of magnetic fields, and since current is a function of magnetic field strength this means a de-energized inductor is guaranteed to have zero current. As energy gradually transfers from the voltage source to the inductor, its magnetic field strength increases and along with that so must its current: :



Since energy transfer in or out of an inductor requires voltage across the inductor's terminals, the amount of voltage impressed across the inductor dictates energy transfer rate and therefore how rapidly inductor current varies. This may be expressed mathematically as follows:

$$V = L \frac{dI}{dt}$$

Where,

V = voltage in Volts (V)

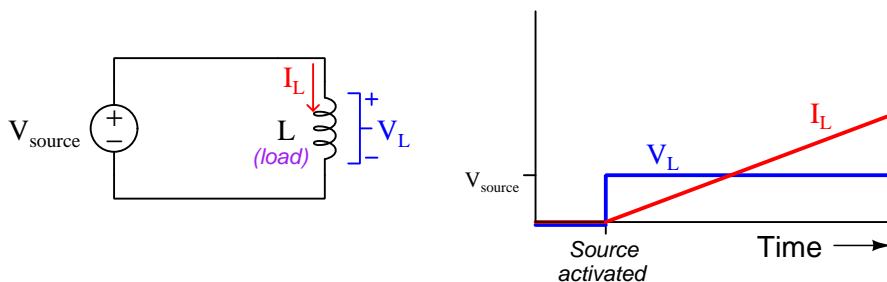
L = inductance in Henrys (H)

$\frac{dI}{dt}$ = rate-of-change of current in Amperes per second (A/s)

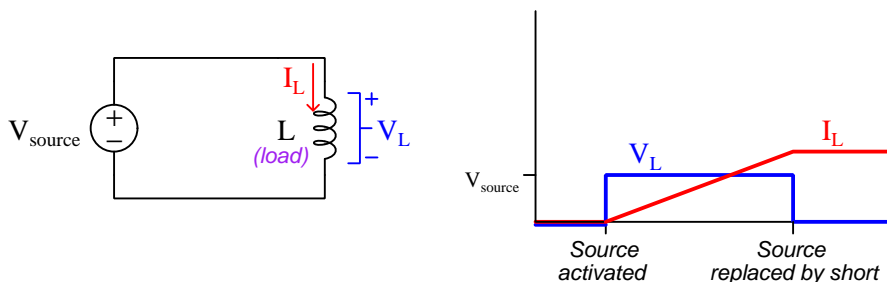
This formula may be thought of as “Ohm’s Law” for inductance. Unlike a resistor, where voltage and current are always in direct and immediate proportion to each other, the *speed* of an inductor’s current is directly proportional to its voltage. The presence of time in the characteristic equation makes inductance behave quite different from resistance, and any circuit containing a substantial amount of inductance must be analyzed with *time* as one of its variables.

⁶This would be an inductor possessing only inductance, with no wire resistance whatsoever. This is actually possible to achieve, by using superconducting wire!

If we were to repeat this experiment on the same size inductor, but using a voltage source with only half as much voltage, we would find the current still rises linearly over time but at half the rate. The decreased voltage means the rate of energy transfer to the inductor from the source is less, and with that the rate of current rise must also be less because (once again) inductor energy and inductor current are directly related to each other:



Replacing the voltage source with a short (wire) at some point in time prior to the end of the experiment results in the inductor holding its store of energy within its magnetic field and therefore (ideally⁷) maintaining a constant current:



Thinking in terms of *energy storage* is essential for proper reasoning regarding voltage and current for any inductance. Inductance is defined as the storage of energy within a *magnetic field*, and magnetic fields directly relate to *current*. Current cannot instantaneously jump from one value to another for any inductance because in order for inductive current to vary at all there must be a transfer of energy either to or from somewhere else, and this always requires time to occur. Voltage, on the other hand, may very well immediately jump from one value to another for an inductance because voltage merely dictates the *rate* at which energy transfers in or out rather than the absolute amount of energy stored. Joule's Law is helpful here, *power* being the product of voltage and current ($P = IV$): for any given amount of current, the rate of energy transfer (i.e. power) depends on how much voltage there is.

⁷This will be true only if the circuit contains absolutely no resistance to dissipate energy as current continues to flow. If resistance is present within the inductor, it will be unable to indefinitely store energy, and as a result its energy store will dissipate and likewise its current will gradually decrease.

3.3 Applications of inductance

Inductors find many applications in DC circuits. As we see from the previous experiments, we may use an inductor to store energy in a manner similar to that of a secondary-cell (“rechargeable”) battery. Connecting an inductor across the terminals of a DC current source will cause that inductor to energize and absorb energy from that source. Subsequently short-circuiting that inductor results in the inductor holding that stored energy (assuming no resistance exists either in the inductor or in the shorting wire), and then when the inductor is connected to a load it delivers that stored energy to the load.

Storing any substantial amount of energy in an inductor necessitates very large amounts of inductance, requiring coils with *many* turns of wire. The problem with using high-turns coils is that wire resistance will dissipate any stored energy rather rapidly. Therefore, the only practical way to make an inductive energy storage device capable of holding large amounts of energy over usable periods of time is to use *superconducting* wire coils having absolutely zero resistance. Such inductors comprise the heart of *Superconducting Magnetic Energy Storage systems*, or *SMES*⁸.

⁸Even using superconducting wire, current state-of-the-art SMES are relatively limited in applications. Some have been applied as backup power sources for certain industrial applications which are intolerant of even momentary lapses in utility power, where the SMES need only supply power for very brief outages.

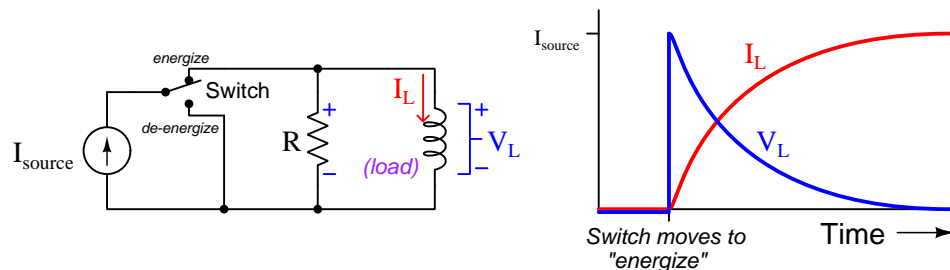
Another use of inductors is *current stabilization*, exploiting the principle that an inductance's current cannot vary unless energy transfer occurs. An inductor connected in series with some other component will force that component's current to be more stable than it might be otherwise, the inductor acting as a load when current rises and acting as a source when current falls. This is analogous to a *flywheel*⁹ acting to stabilize the rotating speed of a machine by serving as a mechanical load as speed increases and serving as a mechanical source as speed decreases. The following photograph shows a set of three *line reactors* used to stabilize current in high-voltage power line conductors during fault events:



Short-circuit faults in electric power systems result in extremely high levels of current which can be damaging to life, property, and the components of the power system. The purpose of the line reactors (inductors) is to retard the rise of current at the start of a fault, so that protective devices such as circuit breakers have more time to react to the fault and de-energize the faulted line(s) before the current magnitude reaches unreasonably high levels.

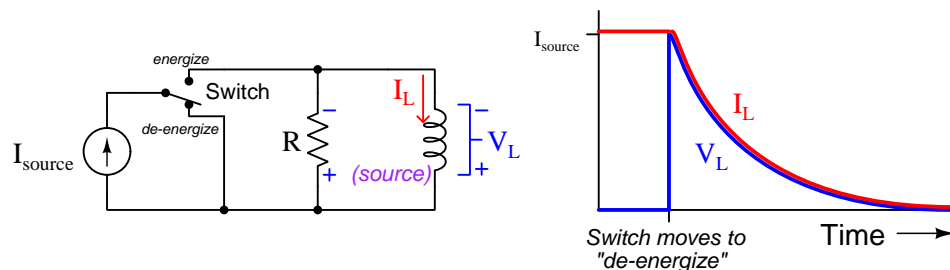
⁹A “flywheel” is typically a heavy wheel coupled to a rotating shaft, the mass of that wheel serving to store kinetic energy as it spins.

Yet another use of inductors is creating time-delayed voltages and currents. Consider this example circuit showing an inductor and resistor connected to a DC current source through a single-pole double-throw (SPDT) switch:



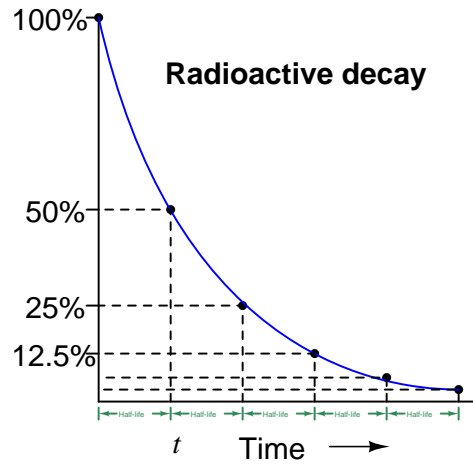
We know inductors store energy by means of magnetic fields, and that the strength of an inductor's magnetic field is directly related to current through it. If we assume the inductor begins in a fully de-energized state, its current therefore must be zero when the switch initially moves to the "energize" position. With zero I_L , Kirchhoff's Current Law demands the resistor conduct the full source current, and from Ohm's Law ($V = IR$) we can tell voltage must start at its maximum value. A high voltage value results in a high rate-of-change for inductor current ($V = L \frac{dI}{dt}$), and so I_L increases rapidly. As I_L increases, less of the source's current will go through the resistor in accordance with Kirchhoff's Current Law, and by Ohm's Law this means less voltage drop. Less voltage results in a slower rate-of-rise for current ($V = L \frac{dI}{dt}$), and so I_L increases not linearly but rather it rises at a slower and slower pace until it finally levels off at the source current value. As I_L rises, I_R falls, and so voltage must fall as well in accordance with Ohm's Law, eventually settling at zero but at a slower and slower speed over time.

Moving the SPDT switch to the "de-energize" position causes the inductor to behave as a source now rather than as a load, delivering its stored energy to the resistor. As the inductor releases energy its magnetic field weakens and its current decreases, causing voltage to decrease as well (Ohm's Law, $V = IR$), resulting in a slower and slower descent for both voltage and current over time:



3.4 Inverse exponential growth/decay

The voltage and current curves plotted in the previous section for circuits containing inductors and resistors are actually common to a wide range of phenomena. For example, if we plot the temperature of a hot object as it cools down to equal ambient air temperature, we will obtain a graph having the same shape as the one showing an inductor de-energizing through a resistor. Another example of similar mathematical behavior is found in radioactive materials, where atoms randomly disintegrate to become other types of elements. An easily-comprehended way to quantify the decay of a radioactive substance is by a measure called *half-life*, which is the amount of time required for one-half of the remaining atoms to disintegrate:



The proportion of original radioactive atoms remaining after each half-life period may be calculated using reciprocated powers of two: after the first half-life period only $\frac{1}{2}$ of the original atoms are intact; after the second half-life period only $\frac{1}{4}$ (or $\frac{1}{2^2}$) of the original atoms are intact; after the third half-life period only $\frac{1}{8}$ (or $\frac{1}{2^3}$) of the original atoms remain, etc. An alternative way to express this mathematically is to use *negative powers of two* such as 2^{-2} to represent $\frac{1}{4}$ and 2^{-3} to represent $\frac{1}{8}$.

Simple powers of two work well to describe the proportion of original atoms remaining after whole intervals of the half-life period, but if we wish to calculate the proportion at any arbitrary time we will need a more sophisticated formula. Here we let t represent time and h represent the half-life of the substance:

$$\text{Proportion remaining} = 2^{-\frac{t}{h}}$$

The fraction $\frac{t}{h}$ simply declares *how many half-lives have elapsed*, and the negative sign ensures the power-of-two will have a fractional value between one and zero. For example, we know that the radioactive gas Tritium has a half-life of 12.3 years. If we take a pure sample of Tritium and keep it in a sealed vessel for exactly 4 years, we would expect 79.82 percent of the original Tritium atoms to remain in that vessel after those 4 years, the decayed atoms having converted into Helium atoms:

$$\text{Tritium proportion remaining} = 2^{-\frac{4}{12.3}}$$

$$\text{Tritium proportion remaining} = 2^{-0.3252}$$

$$\text{Tritium proportion remaining} = 0.7982 = 79.82\%$$


We may modify this equation to tell us the proportion of Tritium atoms that have decayed into Helium over time t as well, which is simply the mathematical complement of what remains:

$$\text{Helium proportion} = 1 - 2^{-\frac{4}{12.3}}$$


$$\text{Helium proportion} = 1 - 0.7982$$

$$\text{Helium proportion} = 0.2018 = 20.18\%$$

If we could continue this test through the end of time, eventually the Tritium proportion would reach zero and the Helium proportion would reach 100%.



When analyzing resistor-inductor networks we find that negative powers of e work better than negative powers of two to model growth and decay, e being a mathematical constant approximately equal to 2.718. Instead of working with *half-life* quantities we instead measure the growth/decay in terms of *time constant* which is symbolized by the Greek letter “tau” (τ) and is simply equal to inductance divided by resistance ($\tau = \frac{L}{R}$). Any decaying value in one of these networks follows the proportion of $e^{-\frac{t}{\tau}}$ while any growing value follows $1 - e^{-\frac{t}{\tau}}$ as shown in the following graphs:



To illustrate how each of these formulae work, we will calculate the values of $e^{-\frac{t}{\tau}}$ and $1 - e^{-\frac{t}{\tau}}$ for several whole-numbered time constant periods beginning with zero representing the system at its start time and progressing to increasingly *negative* values:


$-\frac{t}{\tau}$	$e^{-\frac{t}{\tau}}$	$1 - e^{-\frac{t}{\tau}}$
0	1	0
-1	0.3679	0.6321
-2	0.1353	0.8647
-3	0.04979	0.9502
-4	0.01832	0.9817
-5	0.006738	0.9933

After one time constant ($t = \tau$) the decaying variable will be at 36.79% of its original value while the growing variable reaches 63.21% of its final value. After five time constants' worth of time, all values have settled to well within 1% of their final values.

It is important to realize that the formulae $e^{-\frac{t}{\tau}}$ and $1 - e^{-\frac{t}{\tau}}$ merely describe *fractional proportions* of some initial or final quantity and not necessarily the quantity itself. For example, the current through a resistor-inductor network losing energy over time will *follow* the $e^{-\frac{t}{\tau}}$ “decay” formula as it goes from a value of 1 at the starting time ($t = 0$) and ends up at 0 after an infinite amount of time has passed, but to calculate the actual amount of current in units of Amperes we must incorporate a multiplier describing that initial current. For example, a resistor-inductor network initially energized to a current of 5 Amperes with a time constant of 100 milliseconds would be most fully described by the decay formula $I_L = 5e^{-\frac{t}{0.1}}$ where I_L is the inductor’s current and t is the discharge time in seconds. Similarly, a resistor-inductor-source network starting at an inductor current of zero at $t = 0$ and building up to a maximum of 270 milliAmperes with a time constant of 2 seconds would be most fully described by the growth formula $I_L = 0.27(1 - e^{-\frac{t}{2}})$.

3.5 Inverse-exponential calculation examples

Suppose the inductor in the following circuit begins in a completely de-energized state, with the two SPST¹⁰ switches in the positions shown:



What will happen if we flip the on/off switch to the “On” position while leaving the charge/discharge switch open? To answer this question, we should first consider the inductor from the perspective of energy storage: we know inductors store energy using magnetic fields, and that magnetic field strength is proportional to current, therefore an inductor with no energy must exhibit no current, and as it gains energy from the source its current must rise over time. Therefore, we know we must use the inverse-exponential *growth* equation $(1 - e^{-\frac{t}{\tau}})$ to predict how the inductor’s current will change over time.

The final value for current in this circuit is limited by source voltage and total resistance, which in this case is 2 Volts and 600 Ohms, respectively, giving a final current value of 3.333 milliAmperes. Incorporating this as a multiplier for the inverse-exponential growth equation gives us this form:

$$I_L = (3.333 \text{ mA})(1 - e^{-\frac{t}{\tau}})$$

We also know the circuit’s time constant (symbolized by the Greek letter “tau” τ) will be equal to the quotient of 5 Henrys and 600 Ohms, which is 8.333 milliseconds. If we evaluate this equation at 10 milliseconds, 20 milliseconds, and 30 milliseconds we get inductor current values of 2.329 milliAmperes, 3.031 milliAmperes, and 3.242 milliAmperes, respectively.

¹⁰“SPST” refers to *single-pole, single-throw*.

Multiple ways exist to calculate voltages in this circuit. One way is to take each of the calculated inductor current values at specific times and use Ohm's Law to compute the resistor voltage drops at each of those times, then use those voltage values and Kirchhoff's Voltage Law to calculate inductor voltage. Being a series network, of course, the inductor's current value at any given time will be the same as current through all other component in the network.

Alternatively, we could select the appropriate inverse-exponential equation to directly calculate inductor voltage at any time. Knowing that the inductor begins in a completely de-energized state and gradually absorbs energy over time, we know its current must begin at zero and rise over time. If inductor current begins at zero and rises over time, all resistors (sharing that same current in series) will experience voltages beginning at zero and also rising over time. Kirchhoff's Voltage Law demands the inductor's voltage must begin at the full source value of 2 Volts and decay to zero over time, which we may predict using the inverse-exponential decay equation:

$$V_L = (2 \text{ V})(e^{-\frac{t}{\tau}})$$

If we evaluate this equation at 10 milliseconds, 20 milliseconds, and 30 milliseconds we get inductor voltage values of 0.6024 Volts, 0.1814 Volts, and 54.65 milliVolts, respectively. These voltage values are identical to those calculated the other way, using inductor current values at the same times along with Kirchhoff's Voltage Law and Ohm's Law.

What happens if we now close the "Charge/Discharge" switch after having let the inductor settle to a maximum current value of 3.333 milliAmperes? Doing so will connect the energized inductor directly to the 330 Ohm resistor where it will deliver its stored energy to that load over time. As the inductor loses its store of energy, its magnetic field strength must decrease because that is the form its stored energy takes. The decreasing magnetic field means inductor current must also decay to zero, and from this we may know the inverse-exponential decay equation will describe the inductor's current at different points in time following the closure of the "Charge/Discharge" switch.

The 330 Ohm resistor is still in series with the inductor, which means those two component currents must be the same at any given time. With the "Charge/Discharge" switch in the closed position, that resistor is also in parallel with the inductor which means the two components will also share the exact same voltage at any given time. The time constant (τ) will be different than before with only 330 Ohms in series with the 5 Henry inductor, though:

$$\tau = \frac{5}{330} = 15.15 \text{ milliseconds}$$

$$I_L = (3.33 \text{ mA})(e^{-\frac{t}{\tau}})$$

$$V_R = V_L = (3.33 \text{ mA})(330 \Omega)(e^{-\frac{t}{\tau}})$$

Evaluating these equations at 10 milliseconds, 20 milliseconds, and 30 milliseconds yields results of 0.5685 Volts at 1.723 milliAmperes, 0.2938 Volts at 0.8905 milliAmperes, and 151.9 milliVolts at 0.4602 milliAmperes, respectively.

3.6 Inductor types and limitations

Inductors are broadly categorized by the type of core material used. Air, iron, steel, and ferrite are among the different materials employed as core media within the wire coil. These core materials vary in *permeability* (i.e. the degree to which they enhance magnetic fields) and *magnetic hysteresis* (i.e. the degree to which they remain magnetized after exposure to the coil's magnetic field). As shown in the following equation, permeability directly relates to the inductance of a coil for any given coil area, length, and number of wire turns:

$$L = \frac{\mu AN^2}{l}$$

Where,

L = Inductance in Henrys (H)

μ = Absolute¹¹ permeability of the core material in Henrys per meter (H/m)

A = Cross-sectional area of the solenoid coil in square meters (m²)

N = Number of turns of wire in the solenoid coil (unitless)

l = Length of the solenoid coil in meters (m)

Air-core inductors exhibit much less inductance than inductors with ferromagnetic core materials such as iron or steel or ferrite. Ferrous-core inductors may be symbolized in schematic diagrams by a pair of lines drawn parallel to the coil's axis:



A major advantage of air-core inductors is that they are immune to the effects of *magnetic saturation*¹², because empty space has no limit for magnetic flux. Ferromagnetic materials, by contrast, reach a saturation point where all their magnetic domains are aligned, after which the material is as magnetized as it can be and cannot further enhance the magnetic field. Beyond the saturation point for a ferrous core any further energy storage for the inductor is on the basis of an air-core, which is to say the inductance of an inductor dramatically decreases after saturation.

¹¹Material permeability is often specified as *relative* to a vacuum, in which case μ_r is equal to 1 for a vacuum and greater than 1 for all other substances. The permeability of free space (μ_0 is approximately $1.25663706212 \times 10^{-6}$ Henrys per meter) must be multiplied by any given relative permeability value in order to arrive at the absolute permeability of the substance in question.

¹²The alignment of magnetic “domains” within a ferromagnetic material as it experiences an external magnetic field has the effect of multiplying the effect of that magnetic field. However, this “multiplying” effect only occurs when there are domains within the core material that are *not yet* aligned, *but will* align with further magnetization. As coil current increases, more and more of the material's domains align with the growing external field, leaving fewer and fewer un-aligned domains ready to pivot, and this means diminishing returns for any further increases in magnetic field. At some high value of coil current, *all* domains within the core material will be aligned with the coil's field, at which point the core can contribute no further “multiplication” effect. The field may be strengthened beyond this saturation point by further increases in coil current, but these increases will occur at the same rate as for vacuum or air rather than at the accelerated rate when the ferromagnetic material was unsaturated. In other words, once we reach saturation the core ceases to be of further benefit.

The most elementary ratings for an inductor are *inductance* and *maximum current*. Inductance is expressed in Henrys, typically by means of a printed text label on the inductor body. The maximum current rating for any inductor is limited by either the wire's ampacity¹³ or the magnetic core's saturation, whichever is lower.

Other inductor ratings include *winding resistance* which is the total resistance of the wire coil in Ohms (ideally zero) and *parasitic capacitance*¹⁴ between the insulated turns of wire.


For any given type of core material, the physical size of an inductor is proportional to its inductance as well as to its current rating. An inductor having more inductance but the same current rating must be made either with more turns of wire in the coil or more cross-sectional coil area, and this makes it larger. An inductor having a higher current rating but the same inductance must be made with thicker wire for greater ampacity and/or a core with greater cross-sectional area to accommodate more lines of magnetic flux, and this again makes it larger.

¹³“Ampacity” is the term used to denote the maximum amount of current any conductor may continuously sustain based on the amount of heat dissipated by the conductor's resistance.

¹⁴Capacitance is the ability to store energy in an electric field between two insulated conductors. A “parasitic” property is one that exists as an artifact of a component's construction, and is not necessarily desired. For an inductor, any stray capacitance is considered parasitic, as is any resistance.


3.7 Series and parallel inductance

Like resistors, inductors may be connected together to form networks with different total inductance than any single inductor. The relationship between series versus parallel connections and total inductance is the same as it is for resistors: connecting inductors in series accumulates total inductance, while connecting inductors in parallel diminishes total inductance:



A few simple thought experiments prove how inductance adds in series and diminishes in parallel. Imagine a 3 Henry inductor carrying 2 Amperes of current, storing 6 Joules of energy as predicted by $E_L = \frac{1}{2}LI^2$. Now imagine three of these same inductors connected in series so all of them share the same 2 Ampere current. Total stored energy must be the sum of the three individual inductors' energies and therefore is 18 Joules, *three times* that of the single inductor. Now imagine three of these same inductors connected in parallel, splitting the 2 Ampere total current into thirds. Each 3-Henry inductor carrying $\frac{2}{3}$ Ampere of current will store only $\frac{2}{3}$ Joule of energy with the three together storing only 2 Joules of energy in total, just *one third* that of the single inductor.


The series and parallel relationships shown above assume the inductors are entirely independent of each other. If the magnetic fields of two or more inductors happen to link, the relationships become much more complicated. If current through one inductor is able to induce voltage in an adjacent inductor, the effect is called *mutual inductance* (M). We may perform a "thought experiment" whereby two identical coils of wire are wrapped around a common ferromagnetic core so that they completely share each others' magnetic fields:



In the series-aiding case, the two coils' magnetic fields always point in the same directions, and so enhance each other to make a stronger magnetic field than either coil would on its own. The result will be a greater amount of total inductance, even greater than two independent (non-linked) inductors wired in series. In the series-opposing case, the two coils' magnetic fields are always opposite each other and therefore completely cancel. The result here will be zero magnetic field, and therefore zero energy storage and zero inductance.

3.8 Making custom inductors

Inductors are the easiest of the basic electrical components to *build* yourself, since they simply consist of a wire coil. The formula predicting the theoretical inductance of a wire coil where the length (l) is much greater than its diameter is, or for a toroidal coil, is as follows:



$$L = \frac{\mu AN^2}{l}$$

Where,

L = Inductance in Henrys (H)

μ = Absolute permeability of core material in Tesla-meters per Ampere ($4\pi \times 10^{-7}$ for empty space T·m/A)

A = Cross-sectional area of the core in square meters (m^2)

N = Number of turns of wire in the coil

l = Length of coil in meters (m)

Of these two inductor forms, toroidal is generally preferred because it is more compact and also has the beneficial characteristic of better containment of its magnetic field. A small inductor using an epoxy-coated powdered-iron toroidal core is shown below, with the wire coil wound properly so that it covers most of the toroid's diameter:



It is not difficult to build your own inductor by wrapping enamel-insulated copper wire around a toroidal powdered-iron or ferrite core. To simplify the inductance calculations, manufacturers of toroidal magnetic cores typically advertise an *inductance index* or *inductance factor* for each core based on physical dimensions and core material composition. This value is symbolized as A_L and is often expressed in units of microHenrys or nanoHenrys per turn-squared. What this factor does is consolidate¹⁵ permeability (μ), core cross-sectional area (A), and magnetic path length (l) into a single coefficient describing the magnetic characteristics of a particular toroidal core.

$$L = \frac{\mu AN^2}{l} \quad (\text{original inductance formula})$$

$$L = A_L N^2 \quad (\text{simplified inductance formula})$$

Where,

L = Inductance in Henrys (H)

A_L = Inductance factor in Henrys per turn-squared (H/turn²)

N = Number of turns of wire in the coil

For example, If we were to wrap 25 turns of wire around a toroidal core having an A_L value of 16,000 nanoHenrys per turn-squared, the inductance would be:

$$L = (16000 \times 10^{-9})(25^2) = 10 \text{ mH}$$

A very important caveat for all inductance calculations using ferromagnetic core materials is that the permeability of these materials is never constant. The more strongly magnetized a ferromagnetic material is, the less permeability it exhibits. This, in fact, is why these materials exhibit *B-H curves* rather than *B-H lines*, and why they also are susceptible to *saturation*. Magnetic core manufacturers typically specify the *initial permeability* (μ_i) of the core material, which is its permeability at zero magnetic flux and thus represents a maximum (ideal). Since the inductance of any inductor is directly proportional to core permeability, if you happen to know what the permeability of the material will be under typical operating conditions (μ_{op}), you may calculate the ratio of this operating permeability to initial permeability and arrive at a “de-rating” factor which you may then apply to your inductance calculation:

$$L = \left(\frac{\mu_{op}}{\mu_i} \right) A_L N^2$$

Air-core inductors (e.g. with wire wrapped around non-magnetic forms made of plastic) do not suffer from variations in permeability or from magnetic saturation, but of course without the presence of a ferromagnetic core the amount of inductance possible for any given coil size is very limited. We typically encounter air-core inductors in applications where mere microHenrys or nanoHenrys of inductance is sufficient. For the vast majority of inductor applications ferromagnetic cores are a necessity.

¹⁵In other words, $A_L = \frac{\mu A}{l}$

Chapter 4

Historical References

This chapter is where you will find references to historical texts and technologies related to the module's topic.

Readers may wonder why historical references might be included in any modern lesson on a subject. Why dwell on old ideas and obsolete technologies? One answer to this question is that the initial discoveries and early applications of scientific principles typically present those principles in forms that are unusually easy to grasp. Anyone who first discovers a new principle must necessarily do so from a perspective of ignorance (i.e. if you truly *discover* something yourself, it means you must have come to that discovery with no prior knowledge of it and no hints from others knowledgeable in it), and in so doing the discoverer lacks any hindsight or advantage that might have otherwise come from a more advanced perspective. Thus, discoverers are forced to think and express themselves in less-advanced terms, and this often makes their explanations more readily accessible to others who, like the discoverer, comes to this idea with no prior knowledge. Furthermore, early discoverers often faced the daunting challenge of explaining their new and complex ideas to a naturally skeptical scientific community, and this pressure incentivized clear and compelling communication. As James Clerk Maxwell eloquently stated in the Preface to his book *A Treatise on Electricity and Magnetism* written in 1873,

It is of great advantage to the student of any subject to read the original memoirs on that subject, for science is always most completely assimilated when it is in its nascent state . . . [page xi]

Furthermore, grasping the historical context of technological discoveries is important for understanding how science intersects with culture and civilization, which is ever important because new discoveries and new applications of existing discoveries will always continue to impact our lives. One will often find themselves impressed by the ingenuity of previous generations, and by the high degree of refinement to which now-obsolete technologies were once raised. There is much to learn and much inspiration to be drawn from the technological past, and to the inquisitive mind these historical references are treasures waiting to be (re)-discovered.

4.1 Early examples of inductors

Inductance is the ability to store energy within a magnetic field, and an *inductor* is a device built expressly for that purpose. Any length of wire exhibits inductance simply due to the magnetic field that necessarily forms around it when electrical charges pass through as a current, but most inductors maximize this energy-storage effect by wrapping the conductor into a coil shape, and often wrapped around a ferromagnetic *core material* to enhance the magnetic field.

The following photograph taken from the book *Wireless Telegraphy and Telephony* by Alfred P. Morgan and published in 1920 shows a coil of wire wrapped around a wooden frame to form what that author refers to as a *helix*, suitable for use in radio communication circuits of the day. In modern parlance we would refer to this as an *air-core inductor*:




FIG. 48.—Helix.

Coils form stronger magnetic fields than straight conductors because the close proximity of each “turn” of the coil results in the circular magnetic fields around each turn reinforcing the fields around adjacent turns, all the turns together acting to produce one large magnetic field with lines running through the center and looping around the outside of the coil assembly.


This next photograph comes from page 84 of *Proceedings of The Institute of Radio Engineers* (Volume 2, Number 1) dated March 1914, showing a *spiral* inductor made from heavy metal strip suitable for use in a high-power radio transmitter circuit:



FIGURE 15—Twenty-five Microhenry Inductance Coil Unit, 200 Amperes.

As the caption indicates, this air-cored inductor is rated at $25 \mu\text{H}$ and may carry upwards of 200 Amperes of radio-frequency AC current.

A set of spiral-wound air-core inductors for radio system use are seen in the following photograph:



The close proximity of these inductors means there will be some *magnetic coupling* between them, such that they will function as a *transformer* linking different inductive circuits together.

Early inductor technology also included variable inductors that worked on the principle of a moving “slide” contact that would connect to the coil at different locations when moved along its length. An example of a “dual-slide” inductor is illustrated in the 1920 *Wireless Telegraphy and Telephony* book by Alfred P. Morgan:




FIG. 78.—Tuning coil of the double slide type.

An ingenious design of variable inductor called the *Variometer* used two coils of wire, one stationary and the other mounted on a pivot such that the angular relation of one coil to the other could be varied by turning a knob. With the two coils connected in series with each other, changing the relative angle between them would allow their respective magnetic fields to reinforce or cancel one another to varying degrees, thereby adjusting the amount of total inductance offered by the assembly.

The book *The Wireless Experiments Manual* written by Elmer E. Bucher and published in 1920 gives instructions on the construction and usage of a variometer. Below we see an illustration of a Variometer using cardboard as forms to hold the wire coils:




FIG. 145. View of the assembled variometer of Fig. 144.

A commercial Variometer of that era is shown in a photograph from the same book:





FIG. 150a. Clapp-Eastham-type S.S. Variometer for receiving circuits.

Next we see a schematic diagram from Bucher's manual showing a Variometer (labeled as *VM*) connected in series with the antenna (*A*) of a radio receiver circuit, the two series-connected coils clearly shown in the diagram:



In this receiver circuit the *VC* is a variable capacitor (or variable *condenser* as it was more commonly known in 1920), *FC* is a fixed capacitor, *D* is a diode (or *detector*), and *P* was a pair of audio headphones for listening to the demodulated radio signal. Together, the parallel-connected variable inductor and variable capacitor formed a *resonant tuning circuit* that could be used to select one radio frequency signal from other signals present at the antenna.

Chapter 5

Derivations and Technical References

This chapter is where you will find mathematical derivations too detailed to include in the tutorial, and/or tables and other technical reference material.

5.1 IEC standard component values

Components such as resistors, inductors, and capacitors are manufactured in several *standard values*, described by IEC standard 60063. Rather than having a single series of standard values, the IEC publishes lists called *E series* based on the number of unique values spanning a single *decade* (i.e. a 10:1 range).

The shortest of these series, called *E3* contains just three values: 10, 22, and 47. The next series is called *E6* with six unique values: 10, 15, 22, 33, 47, and 68. These values represent *significant values* for components, meaning the decimal point may be freely moved to create values spanning multiple decades. For example, “33” simply means one can expect to find components manufactured in values of 33, 3.3, 0.33, and 0.033 as well as 330, 3.3 k, 33 k, etc.

Although this may seem like a strange standard for component manufacturers to follow, there is a compelling logic to it. The terms of each series are closer-spaced at the low end than at the high end, and this allows for *series* and/or *parallel* combinations of components to achieve most any desired value. For example, in the E6 series we only have values with the significant figures 10, 15, 22, 33, 47, and 68, but this doesn’t mean we are limited to *total* values with these significant figures. For example, if we needed 80 Ohms of resistance we could connect a 33 Ohm and 47 Ohm resistor together in series. 50 Ohms could be made from two 68 Ohm resistors in parallel (making 34 Ohms) plus a 15 Ohm and 1 Ohm resistor in series.

On the next page is a table showing the four most common E-series specified by IEC standard 60063.

E3	E6	E12	E24
10	10	10	10
			11
		12	12
			13
	15	15	15
			16
		18	18
			20
22	22	22	22
			24
		27	27
			30
	33	33	33
			36
		39	39
			43
47	47	47	47
			51
		56	56
			62
	68	68	68
			75
		82	82
			91

E48, *E96*, and *E192* series are also found in the IEC 60063 standard, used for components with tighter tolerance ratings than typical.

5.2 Inductor labeling

Inductors store electrical energy using magnetic fields, the amount of energy stored being proportional to the square of the applied current. The ratio of stored energy to current (squared) is called *inductance* (L), measured in the unit of the *Henry* (H), and it is the primary characteristic of any inductor. Color codes used to be popular for denoting inductance value, following a very similar pattern to resistors, but numerical labels are now the norm.

The color code used to mark inductors follows the same chromatic and numerical sequence used for resistors¹, with the first and second bands representing significant digits, the third band representing a power-of-ten multiplier, and the fourth band representing tolerance. However, rather than directly indicating inductance in Henrys, the implied metric prefix is *micro*.

Band color	First digit value	Second digit value	Multiplier	Tolerance
Black	0	0	10^0	$\pm 20\%$
Brown	1	1	10^1	$\pm 1\%$
Red	2	2	10^2	$\pm 2\%$
Orange	3	3	10^3	$\pm 3\%$
Yellow	4	4	10^4	$\pm 4\%$
Green	5	5	10^5	
Blue	6	6	10^6	
Violet	7	7	10^7	
Grey	8	8	10^8	
White	9	9	10^9	
Gold			10^{-1}	$\pm 5\%$
Silver			10^{-2}	$\pm 10\%$
None				$\pm 20\%$

For example, an inductor marked with colored bands Yellow, Violet, Orange, and Gold would be 47×10^3 microHenrys, or 47 mH.

Inductors labeled with numerical markings follow the same digit-digit-multiplier pattern as the color code: two significant digits followed by a power-of-ten multiplier, with an assumed base prefix of *micro*. For example, 331 would represent 33×10^1 microHenrys, which is 330 μH or 0.33 mH. If a decimal point is required for a numerically-labeled inductor, a capital letter “R” is used. For example, instead of labeling a 0.01 μH inductor as 0.01, it would be shown as R01.

¹A useful mnemonic for associating these colors with decimal digits 0 through 9 and the percentages 5-10-20% is as follows: “*Better Be Right Or Your Great Big Venture Goes Wrong. Get Started Now.*”.

Very small inductors are sometimes numerically labeled on the basis of *nano*Henry rather than microHenry. In such cases a capital letter “N” represents the decimal point for the nanoHenry value: for example, 3N9 represents 3.9 nH and 27N represents 27 nH.

Inductor *tolerance* is another important parameter which is often printed on the body of the inductor. For physically small inductors where not enough surface area exists to print percentage figures, we must resort to other means for expressing tolerance. To this end, a system of letter-codes has been developed, shown here in the following table:

Letter code	Tolerance
B	± 0.15 nH
C	± 0.2 nH
S	± 0.3 nH
D	± 0.5 nH
F	$\pm 1\%$
G	$\pm 2\%$
H	$\pm 3\%$
J	$\pm 5\%$
K	$\pm 10\%$
L	$\pm 15\%$
M	$\pm 20\%$
V	$\pm 25\%$
N	$\pm 30\%$

Note how the first four of these codes refer to absolute tolerances in nanoHenry, while the rest represent tolerances in percent. When printed on a inductor’s body, the tolerance code typically follows the numerical inductance code. For example, an inductor with “821J” on its body would be $820 \mu\text{H} \pm 5\%$.

5.3 Inductor parameters and core types

Inductors come in many different types, the primary distinctions being the type of substance used for the magnetic core and the physical shape of that core. Several important performance parameters exist for inductors, listed below:


- **Insulation resistance:** *the desirable ability of the wire insulation to oppose “breakdown” resulting from high voltage applied across that insulation, either from the wire to the core material, from one wire to another (e.g. primary-to-secondary in a transformer), or from one turn of wire to an adjacent turn of wire*
- **DC resistance:** *Ohmic resistance of the wire used to make the winding(s) of the inductor, generally undesirable because this constitutes a mechanism of energy dissipation within the inductor*
- **Rated current:** *the maximum continuous DC current through the inductor based on wire heating from I^2R dissipation*
- **Saturation current:** *the amount of current at which point the core begins to magnetically saturate², and its inductance correspondingly decreases, usually based on a specified percentage of inductance decrease compared to at zero current*
- **Distributed (parasitic) capacitance:** *wire insulation separating winding turns from each other and the winding from the core material (if electrically conductive) constitute parasitic capacitance in the inductor*
- **Self-resonance:** *the frequency at which an inductor will self-resonate based on its inductance reacting with its own parasitic capacitance*
- **Leakage inductance:** *the amount of inductance inherent to magnetic flux escaping the core (i.e. “leaking” flux), which may couple with adjacent conductors*
- **Stability:** *the degree to which an inductor’s value undesirably drifts over long periods of time*
- **Temperature coefficient:** *the degree to which an inductor’s value undesirably changes with temperature*
- **Rated temperature:** *the maximum operating temperature at which an inductor will still perform within specification, limited by insulation type, thermal expansion coefficients of the wire and core materials, and the Curie temperature³ of the core material*

²Ferromagnetic materials assist in the development of magnetic fields as their internal magnetic “domains” align one-by-one with any externally-applied field such as that from an energized coil of wire. If we keep increasing the external magnetic field, we will reach a point where all the ferromagnetic core material’s domains are fully aligned with that field, at which point the core can offer no further assistance to increasing field strength. Here the core is said to be *magnetically saturated*, like a sponge fully saturated with water. Any further increases in magnetic field strength (e.g. increasing coil current beyond the saturation point) will proceed unaided by the core.

³The Curie temperature is that temperature beyond which a ferromagnetic material begins to lose its magnetic properties.

Multiple energy-dissipation mechanisms exist within inductors, and may be broadly divided into two categories: “copper” (wire) losses and “iron” (core) losses. All wires made of non-superconducting materials exhibit some electrical resistance, and this directly converts electrical energy into thermal energy ($P = I^2R$) where it will never return to the circuit. At very high frequencies where the *skin effect* causes AC current to avoid the center of the conductor and only travel near the surface (skin), the wire’s effective cross-sectional area decreases and this leads to greater I^2R losses. Ferromagnetic core materials experience molecular-level motion⁴ when the magnetic field varies in intensity, and this represents another way for electrical energy to become converted into heat and never return to the circuit. Magnetic materials dissipate energy in this way proportional to the amount of hysteresis they exhibit: magnetically “soft” materials having narrow B-H hysteresis curves are more energy-efficient as inductor and transformer cores than magnetically “hard” materials having wide B-H hysteresis curves. Furthermore, if the core material happens to be electrically conductive there will be some amount of electric current induced within the core itself as inductor current varies, as a result of electromagnetic induction. These *eddy currents* encounter electrical resistance in the core material, also resulting in I^2R losses.

A variety of core shapes exist for different sizes and purposes of inductor:




Toroidal cores are exceptionally good at confining the magnetic field to the inductor (i.e. very low “leakage” flux) but are more expensive to manufacture.

Inductor core materials exist to enhance the amount of magnetic flux resulting from electric current passing through the winding(s), as well as to “contain” the magnetic flux so that it remains inside the inductor case and does not link with any conductors outside of the inductor. Some inductors have no solid core material, and are called *air-core* inductors. These have very low inductance values and suffer from unwanted “coupling” with adjacent conductors but enjoy immunity from magnetic saturation. Most inductors use some form of *ferromagnetic* material such as iron or iron alloy.

⁴This effect when viewed on a macroscopic scale is called *magnetostriction*, and accounts for why ferromagnetic inductor and transformer cores may “hum” when their respective windings are energized by AC in the audio-frequency range.

Most inductor designs attempt to minimize energy losses in the magnetic core, one strategy for this being to minimize magnetic hysteresis losses. Hysteresis losses occur when the magnetic domains inside the material have a high degree of *remanence*, necessitating work to be done when magnetizing, unmagnetizing, and re-magnetizing the core material as inductor current varies and/or reverses direction. The amount of energy dissipated in this manner by a magnetic core is proportional to the area enclosed by the B-H curve:



Choosing a magnetic material with a naturally narrow B-H curve helps minimize hysteresis losses. So does “over-building” the core to have greater cross-sectional area than is absolutely necessary, in an effort to spread out the magnetic flux over more material and therefore decrease flux density (B) so the core remains as far from saturation as possible.

Another strategy for minimizing energy losses in the core focuses on the mechanism of *eddy currents* which are electric currents induced in the core material by way of electromagnetic induction. This is a problem when the core material happens to be electrically conductive as well as ferromagnetic (e.g. iron is a good example of this), as these circulating electric currents do no useful work yet still encounter electrical resistance in the core material causing heat dissipation which is energy extracted from the circuit never to return to it. An effective mitigation is to build the core in such a way that it electrically interrupts these currents by not providing a continuous pathway for them to circulate. If the core material is iron, constructing the core as a stack of thin iron sheets coated in (insulating) enamel accomplishes this goal, and we call this a *laminated iron core*. Another strategy is to use powdered iron rather than solid iron. Yet another is to use a material that is ferromagnetic but electrically insulating, for example *ferrite* which is a chemical oxide of iron rather than being metallic iron.

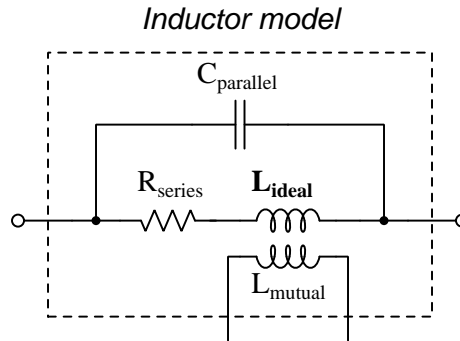
For the end-user rather than the designer of manufactured inductors, an effective strategy for minimizing energy losses is to simply over-rate the inductor in terms of current (i.e. select an inductor with a much higher current rating than the expected current for your application). Using an inductor with a much higher current rating than your expected design current means that inductor will have larger-gauge wire windings, as well as an over-sized magnetic core capable of handling much stronger magnetic fields than you will ever impress upon it. Over-sized windings minimize resistive

energy losses, while the over-sized core effectively shortens and narrows the B-H curve which in turn minimizes hysteresis losses.

5.4 Inductor parasitics

5.4.1 Model of a real inductor

An ideal inductor exhibits only inductance, with no capacitance, resistance, or other characteristics to interfere. Real inductors exhibit all these phenomena to varying degrees, and we collectively refer to these undesirable traits as *parasitic effects*. The following diagram models some of the parasitic effects observed in real inductors:



In addition to the inductance the inductor is supposed to exhibit (L_{ideal}), the inductor also has parasitic resistance (R_{series} , also known as *Equivalent Series Resistance*, or *ESR*), parasitic capacitance ($C_{parallel}$), and mutual inductance (L_{mutual}) with nearby wires and components.

Some of these parasitic effects – such as equivalent series resistance – affect the inductor’s performance in DC applications. Most of the other parasitic effects cause problems in AC and pulsed applications. For example, the effective inductor-capacitor “tank circuit” formed by L_{ideal} and $C_{parallel}$ will cause *resonance* to occur at a particular AC frequency, resulting in much more reactance at that frequency than what would be predicted by the inductive reactance formula $X_L = 2\pi fL$.

Next we will explore common mechanisms for each of these effects.

5.4.2 Parasitic resistance in inductors

Wire resistance plays a dominant role in this parasitic effect due to the typically long lengths of wire necessary to wind the coil that forms most inductors. Wire resistance is not the only dissipative mechanism at work, though. Other losses include magnetic hysteresis of the iron core material as well as *eddy currents* induced in the iron core. An “eddy” current is a circulating electric current induced within the iron core of an inductor, made possible by the fact that iron is an electrically-conductive material as well as being ferromagnetic. These circulating currents do no useful work, and dissipate energy in the form of heating the iron. They may be minimized by forming the iron core from pieces of iron that are electrically insulated from one another, e.g. forming the iron core from *laminated sheets* or *powdered particles* of iron where each sheet or particle is electrically insulated from the next by a layer of non-conductive material on its outer surface.


The series resistance of an inductor is always frequency-dependent. In DC conditions (i.e. frequency of zero Hertz) there will be the basic wire resistance of the coil at play. As frequency increases from zero, however, both the magnetic core losses from hysteresis and eddy currents also increase which add to the DC resistance to form a larger ESR. At extremely high frequencies the *skin effect*⁵ further adds to the inductor’s ESR.

⁵At high frequencies, electric current travels more toward the outer surface of a conductor rather than through the conductor’s entire cross-section, effectively decreasing the conductor’s cross-sectional area (gauge) as frequency rises.

5.4.3 Parasitic capacitance in inductors

Any time an electric field forms between two conductors, energy is stored in that electric field. We call this electric-based energy-storing capability *capacitance*, and of course all inductors must have some capacitance due to the insulating media between wire leads as well as between adjacent turns of wire within the coil (and between the wire turns and the iron core).

Parasitic capacitance is a problem for inductors in AC applications because capacitive reactance (X_C) tends to cancel out inductive reactance (X_L). If we plot the impedance of an inductor as a function of frequency, we would expect an ideal inductor to manifest a straight-line ascent on a logarithmic plot. However, what we see is that at a certain frequency the parallel parasitic capacitance resonates with the inductance to create a nearly-infinite impedance, and then past that frequency the capacitive effects overshadow the inductance:



Precious little may be done to eliminate parasitic capacitance within any inductor, whereas parasitic inductance is fairly easy to minimize within a capacitor. This explains why when faced with an equivalent choice between a circuit design using capacitors and a circuit design using inductors, capacitors nearly always win. Simply put, it is easier to make a nearly-ideal capacitor than it is to make a nearly-ideal inductor.

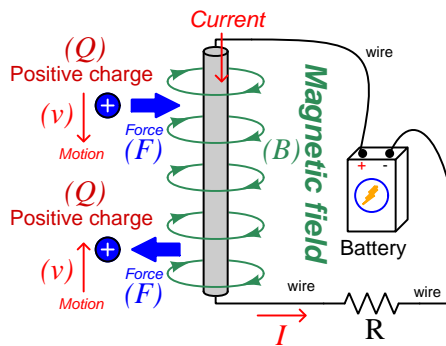
This also explains why the self-resonant frequency of most inductors is much lower than the self-resonant frequency of most capacitors: all other factors being equal, an inductor will have more parasitic capacitance in it than an equivalent capacitor will have parasitic inductance within it, making the LC product greater for the inductor than for the capacitor.

5.4.4 Other parasitic effects in inductors

Mutual inductance occurs whenever adjacent conductors' magnetic fields link with one another, which is difficult to avoid especially in physically dense circuit layouts. This parasitic effect may be minimized by proper placement of inductive components (e.g. keeping them spaced as far apart from each other as possible, orienting their axes perpendicular to each other rather than parallel) as well as by core designs with strong magnetic field containment (e.g. toroidal cores contain their magnetic fields better than rectangular cores).

5.5 Magnetic field quantities

A useful definition of magnetic field (B) is in terms of the force (F , called the *Lorentz force*) exerted on a moving electric charge (Q) influenced by that field:



$$\vec{F} = Q\vec{v} \times \vec{B}$$

Where,

\vec{F} = Force exerted on the charge (Newtons)

Q = Charge quantity (Coulombs⁶)

\vec{v} = Velocity of moving charge (meters per second)


\vec{B} = Magnetic field (Tesla, Webers per square meter, or Newtons per Ampere-meter)

The small “arrow” symbols above the variables for force and velocity and magnetic field in the equation denote those variables as *vector quantities*, having both magnitude and direction. Charge is a *scalar quantity* having only magnitude but no direction, and as a scalar quantity when multiplied by the velocity vector it simply magnifies the magnitude but does not alter the direction. The “cross-product” (\times) is a specific form of vector multiplication, and it results in a product at right angles to the vector directions of both terms. Therefore, the force and velocity and electric field vectors never all point in the same direction.


⁶One Coulomb of electric charge is equal to 6.2415×10^{18} electrons.

Vector cross-products conveniently relate to the fingers of the right hand, which is where the “right-hand rule” originates:

General "right-hand rule"
for vector cross-products




"Right-hand rule" specific
to the Lorentz force




When holding the index finger, middle finger, and thumb of your right hand perpendicular to each other, your index finger points in the direction of the velocity vector (\vec{v}), your middle finger in the direction of the magnetic field vector (\vec{B}), and your thumb in the direction of the force vector (\vec{F}). A simple mnemonic I use to remember these relationships of fingers to vectors is that the **I**ndex finger points in the direction of current⁷ (I), the **M**iddle finger points in the direction of the **m**agnetic field (B), and the **T**humb points in the direction of the **thrust** (i.e. force) acting upon the moving charge.

The Lorentz force’s effect on electrically-charged particles in motion has many applications, from redirecting the paths of charged-particle beams in particle accelerator machines, to bending the trajectory of electron beams in a cathode-ray tube (CRT), to forcing electrons to travel in spiral or circular trajectories inside of magnetron (microwave oscillator) tubes. An illustration of a positively-charged particle curving perpendicular to a magnetic field appears here:



⁷Some textbooks speak of a “left-hand rule” which is intended to make sense of electric charge motion (current) in terms of *electron flow*. As we know, electrons are the only real mobile charge carriers within metal conductors, and so technically “electron flow” notation is most physically accurate when describing the motion of electric charges in metallic circuits. However, the right-hand rule is a *mathematical definition* for vector cross products, the concept of the cross product arising in the late 18th century when electrical science was still in its infancy. Early explorers of electricity used the established mathematical tools of their time and applied it to their work with electric currents and magnetism. At that time, charge carriers in metal wires were assumed to be “positive” and this is how the motion of positively-charged carriers became associated with the first vector of the cross-product. As a result of this assumption which was later proven false, we have two different conventions for denoting the motion of electricity: electron-flow which is *physically* accurate (for metal wires, at least), and conventional flow which is *mathematically* accurate. This, perhaps more than any other reason, is why educational programs designed for mathematically rigorous fields (e.g. electrical engineering) exclusively use conventional flow notation rather than electron flow notation to denote the direction of current.

If the moving charge in question is not a single charged particle but rather part an electric *current* passing through a conductor parallel to the first, both conductors will experience a mutually-attracting force given by the following equation:



$$\vec{F} = I\vec{l} \times \vec{B}$$

Where,

\vec{F} = Force exerted on both conductors (Newtons)

I = Current (Amperes)

\vec{l} = Length of wire (meters)

\vec{B} = Magnetic field (Tesla, or Webers per square meter, or Newtons per Ampere-meter)

The point-charge Lorentz force equation and the two-conductor Lorentz force equation are not that different from one another. Dimensional analysis validates this: the Lorentz force on a moving charge uses that charge quantity (Coulombs) multiplied by the point-charge's velocity in meters per second to give Coulomb-meters per second for the first term:

$$Q\vec{v} = [\text{C}] \left[\frac{\text{m}}{\text{s}} \right] = \left[\frac{\text{C} \cdot \text{m}}{\text{s}} \right]$$

The Lorentz force on a current-carrying conductor uses the current (Amperes, which is Coulombs per second) multiplied by length in meters, for the same composite units of Coulomb-meters per second:

$$I\vec{l} = \left[\frac{\text{C}}{\text{s}} \right] [\text{m}] = \left[\frac{\text{C} \cdot \text{m}}{\text{s}} \right]$$

This dimensional equivalence makes conceptual sense as well: an electrically-charged particle moving through empty space is an electric current in its own right, and an electric current flowing through a conductor is just a collection of charged particles moving through space (just not *empty* space). In either case, the basis for the Lorentz force remains the same: the moving charge(s) create their own magnetic field, which reacts with the magnetic field of the original current-carrying wire to produce forces acting on both.

If the two currents flow in the same direction, their mutual forces *attract*. If the two currents flow in opposite directions, their mutual forces *repel*. This is the basis of electric motors: causing

mechanical motion by electro-magnetic attraction and repulsion. It also represents an interesting contrast with electric fields:


With electric fields, opposite *charges* attract and like *charges* repel.

With magnetic fields, opposite *poles* attract and like *poles* repel.

With parallel currents, opposite *directions* repel and like *directions* attract⁸.

Two parallel current-carrying conductors of length l and separated by a distance d will generate a mutual force proportional to both their currents:

$$F = l \frac{\mu I_1 I_2}{2\pi d}$$



⁸That is, assuming it's *like* charges moving in these directions! If the charges in question are opposite each other – for example electrons in one circuit and holes in another – then like directions will repel and opposite directions will attract!

The circular loops surrounding the current-carrying conductors in the previous illustrations represent the magnetic lines of *flux* (Φ_B) surrounding each of those conductors. The magnetic field (B) is related to magnetic flux by area (A), the field being a measurement of how densely-packed those flux lines are per unit area. For this reason, magnetic field (B) is more properly known as *magnetic flux density*:

$$\vec{B} = \frac{\Phi_B}{\vec{A}}$$

Where,

\vec{B} = Magnetic field or flux density (Tesla, Webers per square meter, or Newtons per Ampere-meter)

Φ_B = Magnetic flux (Webers)

\vec{A} = Area over which flux is distributed (square meters)

An older unit of measurement for magnetic flux density B is the *Gauss* which is much smaller than a Tesla, with one Tesla equivalent to 10,000 Gauss. To put things into perspective, the Earth's natural magnetic field has a strength of approximately one-half of one Gauss⁹.

Magnetic field strength is an inverse function of distance from any current-carrying wire, and also depends on the magnetic permeability of the space adjacent to the wire:

$$B = \frac{\mu I}{2\pi d}$$

Where,

B = Magnetic field or flux density (Tesla, Webers per square meter, or Newtons per Ampere-meter)

μ = Magnetic permeability of the surrounding space (Tesla-meters per Ampere, $4\pi \times 10^{-7}$ for empty space)

I = Current (Amperes)

d = Distance from conductor (meters)

⁹Using the online *Magnetic Field Calculator* application provided by NOAA (the National Oceanic and Atmospheric Administration) at <https://ngdc.noaa.gov/geomag/calculators/magcalc.shtml#igrfwmm>, applying the World Magnetic Model WMM modeling algorithm for years 2019-2024, the total magnetic field strength at my home is 53,584.4 nano-Tesla (53,584.4 nT or 0.535844 Gauss), and presently (May 2020) decaying at a rate of -104.1 nT per year.

The relation of magnetic *flux* to current through a conductor follows a similar equation:

$$\Phi = \frac{\mu AI}{2\pi d}$$

Where,

Φ = Magnetic flux (Webers)

μ = Magnetic permeability of the surrounding space (Tesla-meters per Ampere, $4\pi \times 10^{-7}$ for empty space)

A = Area over which flux is distributed (square meters)

I = Current (Amperes)

d = Distance from conductor (meters)

As this equation makes clear, the amount of magnetic flux surrounding a current-carrying conductor depends not only on the amount of current, but also on the sampled area, the distance from the wire, and also the surrounding material. Most¹⁰ substances (gas, liquid, solid) have permeability values greater than that of empty space, and so this means magnetic flux is usually *enhanced* by the presence of matter around the current-carrying conductor.

The total magnetic flux enclosed by a circular wire loop follows a similar equation:

$$\Phi = \frac{\pi\mu Ir}{2}$$

Where,

Φ = Magnetic flux (Webers)


μ = Magnetic permeability of the surrounding space (Tesla-meters per Ampere, $4\pi \times 10^{-7}$ for empty space)

I = Current (Amperes)

r = Radius of circular loop (meters)

¹⁰Interestingly, superconducting materials *forbid* magnetic fields inside of their bulk, and so the permeability value of any superconductor must be zero!

A common form of electromagnet known as a *solenoid* takes the form of a wire coil wrapped in such a way as to form a long¹¹ cylinder, often wrapped around a plastic frame, and often with a ferromagnetic material such as iron in the center:



The amount of magnetic flux, and the flux density, within the interior of a current-carrying solenoid are given by the following formulae:

$$\Phi = \frac{\mu N A I}{l} \qquad B = \frac{\mu N I}{l}$$

Where,

Φ = Magnetic flux (Webers)

B = Magnetic field or flux density (Tesla, Webers per square meter, or Newtons per Ampere-meter)

μ = Magnetic permeability of the surrounding space (Tesla-meters per Ampere, $4\pi \times 10^{-7}$ for empty space)

N = Number of turns of wire in the coil

A = Cross-sectional area of solenoid coil (square meters)


I = Current (Amperes)

l = Length of solenoid coil (meters)

These formulae have interesting implications for solenoid design. Note how a shorter (i.e. smaller length l) solenoid identical in all other respects will generate a stronger magnetic field for a given current. Note also how the flux density (B) remains constant with increasing cross-sectional area (A) if all other factors are equal, and that this necessarily means a greater amount of total magnetic flux (Φ) for a greater area A .

¹¹These magnetic field formulae apply perfectly to a solenoid coil that is closely-packed (i.e. each turn adjacent to the next) and infinitely long. Therefore, they only approximate real solenoid behavior. This fact may be understood by performing a thought experiment where we decrease the solenoid coil's length to zero, in which case the formulae predict an *infinite* amount of magnetism for any amount of current at all, which of course cannot be true.

Another common form of electromagnet known as a *toroid* is really just a solenoid bent in a circle so that its two ends meet¹² cylinder, often wrapped around a plastic frame, and often with a ferromagnetic material such as iron in the center. Toroids have the unusual property of *containing* their magnetic flux lines extremely well, unlike solenoids, wires, and simple coils which all radiate magnetic fields. They find application as energy-storage devices, or as electromagnets suitable for applying magnetic fields to specimens placed *inside* the toroid's cross-section:



The amount of magnetic flux, and the flux density, within the interior of a current-carrying toroid are identical to that within an otherwise identical otherwise identical solenoid having a length (l) equal to the toroid's circumference:

$$\Phi = \frac{\mu N A I}{l} \qquad B = \frac{\mu N I}{l}$$

Where,

Φ = Magnetic flux (Webers)

B = Magnetic field or flux density (Tesla, Webers per square meter, or Newtons per Ampere-meter)

μ = Magnetic permeability of the surrounding space (Tesla-meters per Ampere, $4\pi \times 10^{-7}$ for empty space)

N = Number of turns of wire in the coil

A = Cross-sectional area of toroid (square meters)

I = Current (Amperes)


l = Circumference of toroid (meters)

If we wish to substitute toroid radius (r) for circumferential length (l), the formulae become the following:

$$\Phi = \frac{\mu N A I}{2\pi r} \qquad B = \frac{\mu N I}{2\pi r}$$

¹²Again, the magnetic field formulae are only accurate for a toroidal coil that is closely-packed (i.e. each turn adjacent to the next) and infinitely long, and therefore only approximate real toroid behavior. This fact may be understood by performing an equivalent thought experiment as before where we decrease the toroid's circumference to zero and absurdly end up with *infinite* magnetism for a finite current.

Many applications of electromagnetism involve conductive *coils* wrapped around some form of ferromagnetic core material, the purpose of that core being to provide a higher-permeability pathway for the magnetic flux than would exist otherwise through air, and the purpose of the wire coil being to intensify the amount of magnetism developed by the electric current beyond what would be possible with a straight current-carrying wire. These magnetic cores typically form a closed loop, or *magnetic circuit* for the lines of magnetic flux to naturally form a closed path. A simple example appears here:



The amount of magnetic flux (Φ) present in the magnetic “circuit” formed by the iron core depends on many factors. First and foremost is the amount of electric current (in Amperes) passing through the wire coil and the number of turns that coil makes around the iron core. The product of this current and the number of turns is called the *magnetomotive force* or *mmf* of the magnetic circuit, analogous to “electromotive force” or “emf” often used as a synonym for voltage in an electric circuit. Not surprisingly, the standard metric unit of measurement for magnetomotive force is the *Ampere-turn*.

However, magnetomotive force alone does not fully describe the current’s effect on magnetism within the iron core. The total length of the magnetic circuit is also an important factor, since a longer path distributes that magnetomotive force over a greater distance. The quotient of magnetomotive force and magnetic circuit length is called *magnetic field intensity*, symbolized by the variable H and expressed in units of *Ampere-turns per meter*.

Magnetic permeability (μ) relates magnetic field intensity (H) to the magnetic flux density (B) within the core material, such that a greater permeability will result in higher flux density for any given amount of field intensity. Permeability is a property of the core material and not its geometry, mathematically defined as the ratio of flux density to field intensity: $\mu = \frac{B}{H}$

Magnetic reluctance (\mathfrak{R}) relates magnetomotive force (mmf) to magnetic flux (Φ), and is related not only to the core material’s permeability but also its geometry. It is mathematically defined as the ratio of magnetomotive force to magnetic flux: $\mathfrak{R} = \frac{\text{mmf}}{\Phi}$

If all this seems confusing, you are in good company. Not only are there many magnetic variables, some related to physical geometry and others not, but there are *two different sets of metric units* appropriate for expressing each! The older units were based on the centimeter-gram-second (CGS) version of the metric system, while the newer units are based on the meter-kilogram-second or SI (Système International) version of the metric system.

Quantity	Symbol	SI unit	CGS unit
Magnetomotive force	mmf	Ampere-turn (A-t)	Gilbert (Gb)
Flux	Φ	Weber (Wb)	Maxwell (Mx)
Field intensity	H	Ampere-turns per meter (A-t/m)	Oersted (Oe)
Flux density	B	Tesla (T)	Gauss (G)
Permeability	μ	Tesla-meters per Ampere-turn (T-m/A-t)	Gauss per Oersted (G/Oe)
Reluctance	\mathfrak{R}	Ampere-turns per Weber (A-t/Wb)	Gilberts per Maxwell (G/Mx)

Magnetomotive force (mmf) and magnetic flux (Φ) may be thought of as the “raw” measures of magnetism, with Ampere-turns and Webers being their respective SI metric units. Reluctance (\mathfrak{R}) is the ratio of the two for any given magnetic circuit with known dimensions and core material. Simply put, reluctance tells you how many Ampere-turns of magnetomotive force will be necessary to create one Weber of magnetic flux in a given space.


Magnetic field intensity (H) and magnetic flux density (B) may be thought of as the “normalized” measures of magnetism, with Ampere-turns per meter and Tesla being their respective SI metric units. H and B relate to mmf and flux by the physical dimensions of the magnetic circuit (length and cross-sectional area, respectively). Permeability is the ratio of the two for any given magnetic core material. Simply put, permeability tells you how many Tesla of magnetic field (i.e. flux density, or Webers of flux per square meter or cross-sectional core area) you will obtain for one Ampere-turn per meter of magnetic field intensity applied to a given core material.

Conversion between the newer SI and the older CGS metric units are as follows:

Quantity	Conversion equivalence
Magnetomotive force (mmf)	1 Ampere-turn = $\frac{4\pi}{10}$ Gilberts
Magnetic flux (Φ)	1 Weber = 10^8 Maxwells
Magnetic field intensity (H)	1 Ampere-turn/meter = $\frac{4\pi}{1000}$ Oersteds
Magnetic flux density (B)	1 Tesla = 10^4 Gauss
Permeability (μ)	1 Tesla-meter/Ampere-turn = $\frac{10^7}{4\pi}$ Gauss/Oersteds
Reluctance (\mathfrak{R})	1 Ampere-turn/Weber = $\frac{4\pi}{10^9}$ Gilberts/Maxwell

5.6 Derivation of electron drift velocity

Electrons typically drift at a very slow velocity through an electrical conductor, even when the current value is rather large. To illustrate, we will analyze a case where 200 Amperes of current passes through a solid copper bar with a cross-sectional area of 1 square centimeter. We must imagine this copper bar as being part of a larger circuit complete with a source, because an open-ended metal bar obviously does not comprise a *circuit* for a continuous current to pass nor does it provide any energy to motivate a current:



Our analysis will proceed as follows:

1. Calculate the number of charge carriers within one cubic centimeter of this bar (i.e. 1 centimeter of the bar's length, given a 1 cm² cross-sectional area)
2. Calculate the drift rate of the current (200 Amperes) in charge carriers per second
3. Divide the carrier drift rate (carriers/second) by the carrier density (carriers/cm³) to find the volumetric drift rate in cubic centimeters per second
4. Divide the volumetric drift rate in (cm³/sec) by the cross-sectional area (1 cm²) to find the drift velocity in linear centimeters per second

Note that most values will be shown rounded to two significant figures for the sake of brevity, but the actual calculations will be performed with many more significant figures in order to avoid unnecessary rounding errors.

First, calculating the number of charge carriers within 1 cm³ of solid copper. The density of solid copper metal is 8900 kg per cubic meter, or 8,900,000 grams per cubic meter. Converting this into grams per cubic centimeter by using the "unity fraction" method of unit cancellation:

$$\left(\frac{8900000 \text{ g}}{\text{m}^3}\right) \left(\frac{1 \text{ m}}{100 \text{ cm}}\right)^3 = \frac{8.9 \text{ g}}{\text{cm}^3}$$

The atomic weight of copper (found in a Periodic Table of the Elements) is 63.546 amu, or 63.546 grams per mole. We will use this figure to determine the number of moles of copper atoms within our 1 cm³ volume of solid copper:

$$\left(\frac{8.9 \text{ g}}{\text{cm}^3}\right) \left(\frac{1 \text{ mol}}{63.546 \text{ g}}\right) = \frac{0.14 \text{ mol}}{\text{cm}^3}$$

1 mole of *anything* is 6.022×10^{23} units, so we may calculate the number of individual copper atoms in our 1 cm^3 sample by multiplying 0.14 moles by this number:

$$\left(\frac{0.14 \text{ mol}}{\text{cm}^3}\right) \left(\frac{6.022 \times 10^{23} \text{ atoms}}{\text{mol}}\right) = \frac{8.4 \times 10^{22} \text{ atoms}}{\text{cm}^3}$$

Each copper atom possesses one “free” electron in its valence shell, which means our 1 cm^3 sample of the copper bar contains 8.4×10^{22} free charges (electrons).

One Ampere of electric current is equal to 6.2×10^{18} individual charges passing by per second of time. 200 Amperes of current is therefore equal to:

$$\left(\frac{200 \text{ A}}{1}\right) \left(\frac{6.2 \times 10^{18} \text{ electrons / sec}}{1 \text{ A}}\right) = \frac{1.2 \times 10^{21} \text{ electrons}}{\text{sec}}$$

Dividing this electron flow rate by the volumetric density of electrons within copper will yield a volumetric flow rate in cubic centimeters’ worth of electrons per second:

$$\frac{1.2 \times 10^{21} \text{ electrons / sec}}{8.4 \times 10^{22} \text{ electrons / cm}^3} = \frac{0.015 \text{ cm}^3}{\text{sec}}$$


Dividing this volumetric charge flow rate by the bar’s cross-sectional area yields the linear drift velocity in centimeters per second:

$$\frac{0.015 \text{ cm}^3 / \text{sec}}{1 \text{ cm}^2} = \frac{0.015 \text{ cm}}{\text{sec}}$$

Converting into centimeters per minute:

$$\left(\frac{0.015 \text{ cm}}{\text{sec}}\right) \left(\frac{60 \text{ sec}}{1 \text{ min}}\right) = \frac{0.88 \text{ cm}}{\text{min}}$$

As you can see, the average drift velocity of electrons through this copper bar is quite slow, even for a relatively high¹³ amount of current. At this velocity, the electrons will take over a minute’s worth of time to move just 1 centimeter along the bar’s length!




¹³For reference, 200 Amperes is the maximum amount of current most North American households are rated to consume in total.

5.7 Time delay circuits and Thévenin's/Norton's Theorems

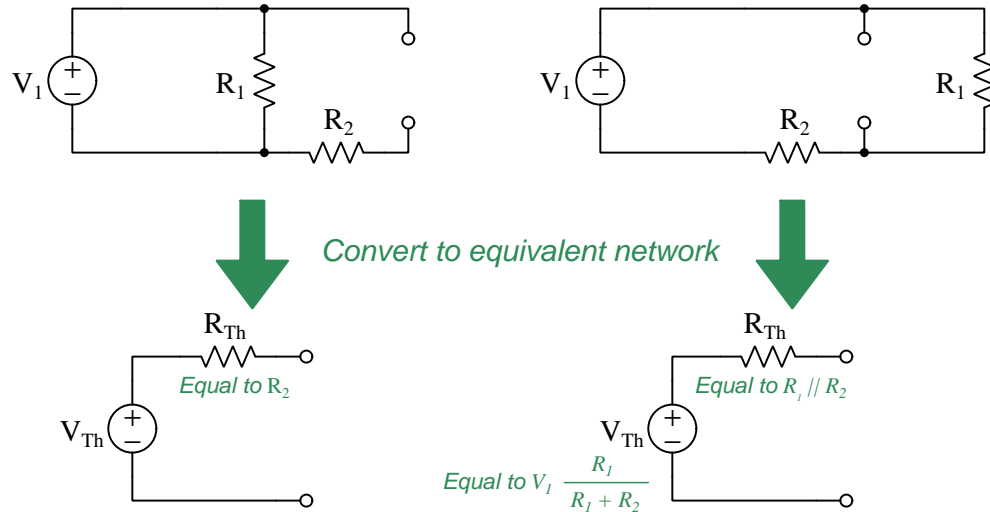
Circuits containing resistors along with capacitors or inductors typically result in voltage and current values that vary over time. Determining starting and final values for these parameters, as well as calculating the “time constant” ($\tau = RC = \frac{L}{R}$) can be challenging especially if the circuit in question consists of more than just a single resistor, single source, and single reactive (C or L).

Take for example the following two RC circuits, each one containing one capacitor but two resistors:



If the capacitor begins in a completely de-energized state with the switch in the open position, capacitor voltage for each circuit is guaranteed to begin at 0 Volts because energy stored in any capacitance always manifests as a voltage across that capacitance. When we close the switch, we would expect the capacitor in both circuits to begin energizing as they absorb energy from the voltage source V_1 , and this means their voltages will rise over time. However, what is not as clear upon first inspection is what the final voltage will be for either capacitor after the switch has been left in its closed state for a long time, as well as how we might calculate the time constant $\tau = RC$ given the fact each circuit has *two* resistors and not just one.

Obviously, an RC circuit containing just one resistor and one capacitor instead of two resistors and one capacitor would be simpler to analyze, so why don't we just apply either Thévenin's Theorem or Norton's Theorem to each circuit in order to reduce it to just a single source and single resistance? Choosing the capacitor to be the load, we will convert the rest of each circuit into its Thévenin equivalent:




Recall that determining either the Thévenin or the Norton equivalent resistance of any complex network consists of disabling all sources in the original circuit and determining resistance as measured from load terminal to load terminal. In the left-hand circuit the disabled V_1 is replaced by a short which completely bypasses R_1 and leaves only R_2 remaining. In the right-hand circuit the disabled V_1 (again, replaced by a short) causes R_1 and R_2 to be in parallel with each other from the perspective of the two load terminals. Thus, the time constant for the left-hand circuit is simply $\tau = R_2 C$ while the time constant for the right-hand circuit will be $\tau = \frac{C}{\frac{1}{R_1} + \frac{1}{R_2}}$.

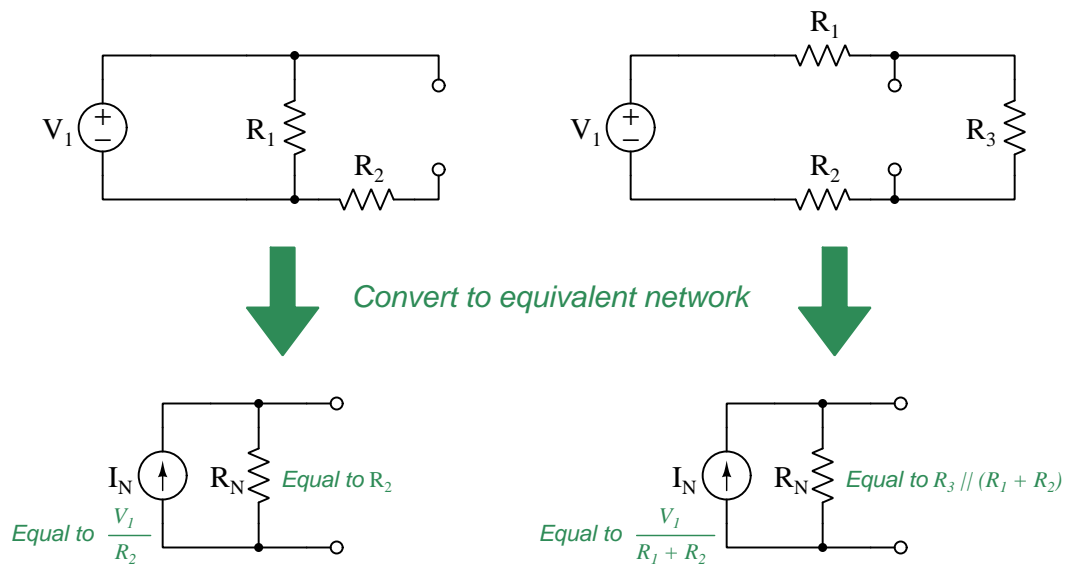
Recall that determining the Thévenin equivalent voltage of any complex network consists of calculating voltage from one load terminal to the other with all original sources active. Here we see that the left-hand circuit's open-circuit load terminal voltage must simply be V_1 , while the right-hand circuit's open-circuit terminal voltage will be the result of V_1 powering a two-resistor voltage divider.

Re-attaching the capacitor C to the load terminals of the Thévenin equivalent networks results in a pair of simpler circuits to analyze, where the time constant in each case is $\tau = R_{Th} C$ and the final capacitor voltage after full energization is V_{Th} .

Norton's Theorem, of course, works just as well as Thévenin's Theorem for this purpose. To illustrate, we will apply Norton's Theorem to a pair of resistor-inductor circuits to determine how multiple resistors affect the inductors' energization:



Removing the inductor from each circuit and converting the remaining networks into Norton equivalents:



The process for determining Norton equivalent resistance is exactly the same as it is for Thévenin equivalent resistance: simply disable all original sources and determine what an ohmmeter would register if connected between the two load terminals. In the left-hand circuit we see once again that R_1 is of no effect to the Norton equivalent value of R_N , while in the right-hand circuit R_N is a series-parallel combination of the three original resistors.

Recall that determining the Norton equivalent current of any complex network consists of calculating current passing through a shorting wire placed between the two load terminals with all the original sources active. In the left-hand circuit we see that this short-circuit current is the limited only by R_2 and so $I_N = \frac{V_1}{R_2}$. In the right-hand circuit we see only R_1 and R_2 serving to limit current, and so $I_N = \frac{V_1}{R_1 + R_2}$.

Re-attaching the inductor L to the load terminals of the Norton equivalent networks results in a pair of simpler circuits to analyze, where the time constant in each case is $\tau = \frac{L}{R_N}$ and the final inductor current after full energization is I_N .

Chapter 6

Programming References

A powerful tool for mathematical modeling is text-based *computer programming*. This is where you type coded commands in text form which the computer is able to interpret. Many different text-based languages exist for this purpose, but we will focus here on just two of them, *C++* and *Python*.

6.1 Programming in C++

One of the more popular text-based computer programming languages is called *C++*. This is a *compiled* language, which means you must create a plain-text file containing C++ code using a program called a *text editor*, then execute a software application called a *compiler* to translate your “source code” into instructions directly understandable to the computer. Here is an example of “source code” for a very simple C++ program intended to perform some basic arithmetic operations and print the results to the computer’s console:

```
#include <iostream>
using namespace std;

int main (void)
{
    float x, y;

    x = 200;
    y = -560.5;

    cout << "This simple program performs basic arithmetic on" << endl;
    cout << "the two numbers " << x << " and " << y << " and then" << endl;
    cout << "displays the results on the computer's console." << endl;

    cout << endl;

    cout << "Sum = " << x + y << endl;
    cout << "Difference = " << x - y << endl;
    cout << "Product = " << x * y << endl;
    cout << "Quotient of " << x / y << endl;

    return 0;
}
```

Computer languages such as C++ are designed to make sense when read by human programmers. The general order of execution is left-to-right, top-to-bottom just the same as reading any text document written in English. Blank lines, indentation, and other “whitespace” is largely irrelevant in C++ code, and is included only to make the code more pleasing¹ to view.

¹Although not included in this example, *comments* preceded by double-forward slash characters (*//*) may be added to source code as well to provide explanations of what the code is supposed to do, for the benefit of anyone reading it. The compiler application will ignore all comments.

Let's examine the C++ source code to explain what it means:

- `#include <iostream>` and `using namespace std;` are set-up instructions to the compiler giving it some context in which to interpret your code. The code specific to your task is located between the brace symbols (`{` and `}`, often referred to as “curly-braces”).
- `int main (void)` labels the “Main” function for the computer: the instructions within this function (lying between the `{` and `}` symbols) it will be commanded to execute. Every complete C++ program contains a `main` function at minimum, and often additional functions as well, but the `main` function is where execution always begins. The `int` declares this function will return an *integer* number value when complete, which helps to explain the purpose of the `return 0;` statement at the end of the `main` function: providing a numerical value of zero at the program's completion as promised by `int`. This returned value is rather incidental to our purpose here, but it is fairly standard practice in C++ programming.
- Grouping symbols such as parentheses and {braces} abound in C, C++, and other languages (e.g. Java). Parentheses typically group data to be processed by a function, called *arguments* to that function. Braces surround lines of executable code belonging to a particular function.
- The `float` declaration reserves places in the computer's memory for two *floating-point* variables, in this case the variables' names being `x` and `y`. In most text-based programming languages, variables may be named by single letters or by combinations of letters (e.g. `xyz` would be a single variable).
- The next two lines assign numerical values to the two variables. Note how each line terminates with a semicolon character (`;`) and how this pattern holds true for most of the lines in this program. In C++ semicolons are analogous to periods at the ends of English sentences. This demarcation of each line's end is necessary because C++ ignores whitespace on the page and doesn't “know” otherwise where one line ends and another begins.
- All the other instructions take the form of a `cout` command which prints characters to the “standard output” stream of the computer, which in this case will be text displayed on the console. The double-less-than symbols (`<<`) show data being sent *toward* the `cout` command. Note how verbatim text is enclosed in quotation marks, while variables such as `x` or mathematical expressions such as `x - y` are not enclosed in quotations because we want the computer to display the numerical values represented, not the literal text.
- Standard arithmetic operations (add, subtract, multiply, divide) are represented as `+`, `-`, `*`, and `/`, respectively.
- The `endl` found at the end of every `cout` statement marks the end of a line of text printed to the computer's console display. If not for these `endl` inclusions, the displayed text would resemble a run-on sentence rather than a paragraph. Note the `cout << endl;` line, which does nothing but create a blank line on the screen, for no reason other than esthetics.

After saving this *source code* text to a file with its own name (e.g. `myprogram.cpp`), you would then *compile* the source code into an *executable* file which the computer may then run. If you are using a console-based compiler such as *GCC* (very popular within variants of the Unix operating system², such as Linux and Apple’s OS X), you would type the following command and press the Enter key:

```
g++ -o myprogram.exe myprogram.cpp
```

This command instructs the *GCC* compiler to take your source code (`myprogram.cpp`) and create with it an executable file named `myprogram.exe`. Simply typing `./myprogram.exe` at the command-line will then execute your program:

```
./myprogram.exe
```

If you are using a graphic-based C++ development system such as Microsoft Visual Studio³, you may simply create a new console application “project” using this software, then paste or type your code into the example template appearing in the editor window, and finally run your application to test its output.

As this program runs, it displays the following text to the console:

```
This simple program performs basic arithmetic on  
the two numbers 200 and -560.5 and then  
displays the results on the computer’s console.
```

```
Sum = -360.5  
Difference = 760.5  
Product = -112100  
Quotient of -0.356824
```

As crude as this example program is, it serves the purpose of showing how easy it is to write and execute simple programs in a computer using the C++ language. As you encounter C++ example programs (shown as source code) in any of these modules, feel free to directly copy-and-paste the source code text into a text editor’s screen, then follow the rest of the instructions given here (i.e. save to a file, compile, and finally run your program). You will find that it is generally easier to

²A very functional option for users of Microsoft Windows is called *Cygwin*, which provides a Unix-like console environment complete with all the customary utility applications such as *GCC*!

³Using Microsoft Visual Studio community version 2017 at the time of this writing to test this example, here are the steps I needed to follow in order to successfully compile and run a simple program such as this: (1) Start up Visual Studio and select the option to create a New Project; (2) Select the Windows Console Application template, as this will perform necessary set-up steps to generate a console-based program which will save you time and effort as well as avoid simple errors of omission; (3) When the editing screen appears, type or paste the C++ code within the `main()` function provided in the template, deleting the “Hello World” `cout` line that came with the template; (4) Type or paste any preprocessor directives (e.g. `#include` statements, `namespace` statements) necessary for your code that did not come with the template; (5) Lastly, under the Debug drop-down menu choose either Start Debugging (F5 hot-key) or Start Without Debugging (Ctrl-F5 hotkeys) to compile (“Build”) and run your new program. Upon execution a console window will appear showing the output of your program.

learn computer programming by closely examining others' example programs and modifying them than it is to write your own programs starting from a blank screen.

6.2 Programming in Python

Another text-based computer programming language called *Python* allows you to type instructions at a terminal prompt and receive immediate results without having to compile that code. This is because Python is an *interpreted* language: a software application called an *interpreter* reads your source code, translates it into computer-understandable instructions, and then executes those instructions in one step.

The following shows what happens on my personal computer when I start up the Python interpreter on my personal computer, by typing `python3`⁴ and pressing the Enter key:

```
Python 3.7.2 (default, Feb 19 2019, 18:15:18)
[GCC 4.1.2] on linux
Type "help", "copyright", "credits" or "license" for more information.
>>>
```

The `>>>` symbols represent the prompt within the Python interpreter “shell”, signifying readiness to accept Python commands entered by the user.

Shown here is an example of the same arithmetic operations performed on the same quantities, using a Python interpreter. All lines shown preceded by the `>>>` prompt are entries typed by the human programmer, and all lines shown without the `>>>` prompt are responses from the Python interpreter software:

```
>>> x = 200
>>> y = -560.5
>>> x + y
-360.5
>>> x - y
760.5
>>> x * y
-112100.0
>>> x / y
-0.35682426404995538
>>> quit()
```

⁴Using version 3 of Python, which is the latest at the time of this writing.

More advanced mathematical functions are accessible in Python by first entering the line `from math import *` which “imports” these functions from Python’s math *library* (with functions identical to those available for the C programming language, and included on any computer with Python installed). Some examples show some of these functions in use, demonstrating how the Python interpreter may be used as a scientific calculator:

```
>>> from math import *
>>> sin(30.0)
-0.98803162409286183
>>> sin(radians(30.0))
0.49999999999999994
>>> pow(2.0, 5.0)
32.0
>>> log10(10000.0)
4.0
>>> e
2.7182818284590451
>>> pi
3.1415926535897931
>>> log(pow(e,6.0))
6.0
>>> asin(0.7071068)
0.78539819000368838
>>> degrees(asin(0.7071068))
45.000001524425265
>>> quit()
```

Note how trigonometric functions assume angles expressed in *radians* rather than *degrees*, and how Python provides convenient functions for translating between the two. Logarithms assume a base of e unless otherwise stated (e.g. the `log10` function for common logarithms).

The interpreted (versus compiled) nature of Python, as well as its relatively simple syntax, makes it a good choice as a person’s first programming language. For complex applications, interpreted languages such as Python execute slower than compiled languages such as C++, but for the very simple examples used in these learning modules speed is not a concern.

Another Python math library is `cmath`, giving Python the ability to perform arithmetic on complex numbers. This is very useful for AC circuit analysis using *phasors*⁵ as shown in the following example. Here we see Python's interpreter used as a scientific calculator to show series and parallel impedances of a resistor, capacitor, and inductor in a 60 Hz AC circuit:

```
>>> from math import *
>>> from cmath import *
>>> r = complex(400,0)
>>> f = 60.0
>>> xc = 1/(2 * pi * f * 4.7e-6)
>>> zc = complex(0,-xc)
>>> xl = 2 * pi * f * 1.0
>>> zl = complex(0,xl)
>>> r + zc + zl
(400-187.38811239154882j)
>>> 1/(1/r + 1/zc + 1/zl)
(355.837695813625+125.35793777619385j)
>>> polar(r + zc + zl)
(441.717448903332, -0.4381072059213295)
>>> abs(r + zc + zl)
441.717448903332
>>> phase(r + zc + zl)
-0.4381072059213295
>>> degrees(phase(r + zc + zl))
-25.10169387356105
```

When entering a value in rectangular form, we use the `complex()` function where the arguments are the real and imaginary quantities, respectively. If we had opted to enter the impedance values in polar form, we would have used the `rect()` function where the first argument is the magnitude and the second argument is the angle in radians. For example, we could have set the capacitor's impedance (`zc`) as $X_C \angle -90^\circ$ with the command `zc = rect(xc,radians(-90))` rather than with the command `zc = complex(0,-xc)` and it would have worked the same.

Note how Python defaults to rectangular form for complex quantities. Here we defined a 400 Ohm resistance as a complex value in rectangular form ($400 + j0 \Omega$), then computed capacitive and inductive reactances at 60 Hz and defined each of those as complex (phasor) values ($0 - jX_c \Omega$ and $0 + jX_l \Omega$, respectively). After that we computed total impedance in series, then total impedance in parallel. Polar-form representation was then shown for the series impedance ($441.717 \Omega \angle -25.102^\circ$). Note the use of different functions to show the polar-form series impedance value: `polar()` takes the complex quantity and returns its polar magnitude and phase angle in *radians*; `abs()` returns just the polar magnitude; `phase()` returns just the polar angle, once again in radians. To find the polar phase angle in degrees, we nest the `degrees()` and `phase()` functions together.

The utility of Python's interpreter environment as a scientific calculator should be clear from these examples. Not only does it offer a powerful array of mathematical functions, but also unlimited

⁵A "phasor" is a voltage, current, or impedance represented as a complex number, either in rectangular or polar form.

assignment of variables as well as a convenient text record⁶ of all calculations performed which may be easily copied and pasted into a text document for archival.

It is also possible to save a set of Python commands to a text file using a text editor application, and then instruct the Python interpreter to execute it at once rather than having to type it line-by-line in the interpreter's shell. For example, consider the following Python program, saved under the filename `myprogram.py`:

```
x = 200
y = -560.5

print("Sum")
print(x + y)

print("Difference")
print(x - y)

print("Product")
print(x * y)

print("Quotient")
print(x / y)
```

As with C++, the interpreter will read this source code from left-to-right, top-to-bottom, just the same as you or I would read a document written in English. Interestingly, whitespace *is* significant in the Python language (unlike C++), but this simple example program makes no use of that.

To execute this Python program, I would need to type `python myprogram.py` and then press the Enter key at my computer console's prompt, at which point it would display the following result:

```
Sum
-360.5
Difference
760.5
Product
-112100.0
Quotient
-0.35682426405
```


As you can see, syntax within the Python programming language is simpler than C++, which is one reason why it is often a preferred language for beginning programmers.

⁶Like many command-line computing environments, Python's interpreter supports "up-arrow" recall of previous entries. This allows quick recall of previously typed commands for editing and re-evaluation.

If you are interested in learning more about computer programming in *any* language, you will find a wide variety of books and free tutorials available on those subjects. Otherwise, feel free to learn by the examples presented in these modules.

6.3 Modeling inverse exponential growth and decay using C++

Many different types of physical systems, including capacitor and inductor circuits, often exhibit a phenomenon known as *inverse exponential* growth and decay, where variables asymptotically⁷ approach final values over time. Consider the following two circuits, one capacitive and the other inductive, both components receiving an input of energy from a source:



Capacitors store energy as a function of voltage (manifest as an *electric field*), while inductors store energy as a function of current (manifest as a *magnetic field*). In both circuits, the reactive component receives energy from the source over time, which explains the capacitor's rising voltage and the inductor's rising current. The other variable in each circuit (i.e. current for the capacitor and voltage for the inductor) is a function of the *rate* at the energy-related variable grows, and so these variables approach zero as the energy-related variables level off at some terminal value.

In each case, the decreasing variable is said to *decay to zero* over time, and its path has the same shape as an inverse exponential function, e^{-t} . As time progresses (i.e. t grows larger), the value of e^{-t} diminishes. The full formula for each circuit includes a term for the initial (peak) value as well as the *time constant* (τ) of the circuit:

$$I_C = I_0 e^{-\frac{t}{\tau}}$$

$$V_L = V_0 e^{-\frac{t}{\tau}}$$

In each case, the increasing variable is said to experience *inverse exponential growth*, and its shape appears to be a mirror-image of the decaying variable. Mathematically, the inverse growth function is simply the complement of the decay function, $1 - e^{-t}$. As time progresses (i.e. larger t), the value of e^{-t} approaches zero and the value of $1 - e^{-t}$ approaches 1. The amount of time equal to one time constant (τ) is the time required for the variable in question to go 63.2% ($1 - e^{-1}$) of

⁷An asymptote is a mathematical concept, consisting of a fixed value which some other variable approaches closer and closer but never fully attains.

the way from its starting value to its final (asymptote) value. For any variable decaying to zero, this is the time required to reach 36.8% (e^{-1}) of its starting value.

These exponential functions are not limited to describing capacitor and inductor circuits, but are actually useful for a wide range of physical systems. The decay of a hot object's temperature⁸, for example, follows the same e^{-t} trajectory.

We may employ a simple computer program to repeatedly calculate percentage values for such rising and decaying variables as time (t) progresses from zero to ten (measured in whole-multiples of time constants). This is a very good application for a computer, as it relieves us of both the tedium and the many opportunities for error faced with repeatedly calculating e^{-t} and $1 - e^{-t}$.

The "source code" for this computer program, written in the C++ language, is shown here containing all the instructions necessary to tell the computer how to perform these repeated calculations and display the results on the console:

```
#include <iostream>
#include <cmath>
using namespace std;

int main (void)
{
    float initial, x, t;

    initial = 100.0;      // This sets the initial value to 100%

    cout.precision(4);
    cout.setf(ios::fixed, ios::floatfield);

    cout << "Time \t\t" << "Falling \t" << "Rising" << endl;

    for (t = 0.0; t <= 10 ; t = t + 1.0)
    {
        x = initial * exp(-t);
        cout << t << " tau \t";
        cout << x << "% \t";
        cout << 100.0 - x << "%" << endl;
    }

    return 0;
}
```

⁸In fact, this is referred to as *Newton's Cooling Law*.

Let's analyze the lines of code within this short C++ program, touching on the following programming principles:

- Reading/execution order
- Preprocessor directives (`#include`) and `namespace`
- The `main` function
- Delimiter characters (e.g. `{ }` ;)
- Variable declarations
- `float` and `int` data types and variable names
- Variable initialization and assignment (e.g. `=`)
- The `cout` instruction
- `for` loops
- Mathematical functions
- Return

As with most text-based programming languages, C and C++ alike are read by the compiler software and executed in the same order one reads an English-language document: left to right character by character, and top to bottom line by line.

The first three lines (`#include` and `using namespace`) direct the compiler software on how to interpret much of what follows. The `using` line tells the compiler to use the “standard” namespace⁹ convention for C++, while the two previous lines (called preprocessor directives) request the contents of the `iostream` and `math` header files to be included when compiling the code into an executable file. For example, the `cout` instruction is defined within `iostream` while the `exp` function is defined within `math`.

All C++ programs have a `main` function marking the starting point of execution. All of our circuit-modeling code appears between the `main` function's curly-brace (`{ }`) symbols, those symbols denoting the boundaries of the `main` function. This particular `main` function is defined here as returning an *integer* number value (`int`) when complete, and it requires no input of information (`void`).

In C, C++, and some other languages such as Java, parentheses and “curly-brace” symbols serve to *group* certain types of information together, much the same as parentheses and brackets group terms together in mathematical formulae. Parentheses immediately following a function name enclose *arguments* to that function: i.e. data that function must act upon. Curly-brace symbols define the starting and ending points for lines of code to be executed within a given function.

⁹Namespaces are a concept found in C++ but not in its predecessor language C. Without getting too detailed, namespaces allow advanced programmers to redefine the meaning of certain key instructions and functions within the language for their own purposes. In this simple program we are doing no such thing, and so we instruct the compiler to use the *standard* (`std`) namespace of C++.

The line beginning with `float` reserves space in the computer's memory for three *floating-point* variables. "Floating-point" is a type of digital numeration useful for expressing non-integer values both large and small. C++ permits variable names ranging from single letters (e.g. `x`) to entire words (e.g. `initial`).

This declaration line ends with a semicolon delimiter (`;`) telling the compiler where the line ends, which is important because neither C nor C++ takes heed of "whitespace" in the source file. You might think of the C/C++ semicolon as being the programming-language equivalent of a period at the end of an English sentence. If not for these semicolons, the compiler would treat all the lines as being a single, continuous line in the same way you would read unterminated lines in a page of text as part of a long sentence. You will notice most of the lines of code end with semicolons, exceptions being `main` (because its curly-brace symbols serve the same purpose) and the `#include` directives which are technically set-up instructions for the compiler and not executed at run-time.

Immediately following the declaration line is a line of code *initializing* one of those variables with a numerical value, in this case setting the `initial` value to `100.0` which will represent 100% of source voltage or current in a capacitive or inductive circuit. In C and C++, a single "equals" symbol (`=`) signifies *assignment* rather than the conventional mathematical meaning of *equality*. Whenever you see a single equals symbol in C or C++, think "*set equal to*".

Several lines of code in this short program begin with `cout` which is the C++ instruction for outputting text to the computer's *console* display. When `cout` is accompanied by the "put to" operator (`<<`) it means that instruction will output either verbatim text or text-control characters (enclosed in quotation marks) or the numerical values of any variables referenced by `cout`. Text-control characters include `\t` (tab) instructing a horizontal shift in text and `\n` (end-line) signifying the end of line line and the beginning of another as displayed on the computer's console.

Those `cout` lines lacking "put to" (`<<`) operators instruct subsequent instances of `cout` how to behave. For example, `cout.precision(4)` tells `cout` to display every numerical output at four decimal places to the right of the decimal point. In the next line of code, `cout.setf(ios::fixed, ios::floatfield)` instructs `cout` to use a fixed-point notation when displaying numbers. Both of these formatting instructions will make the program's displayed text look neater on the computer's console, but neither is strictly necessary if all we care about is function and not appearance.

All the interesting operations occur within the `for` loop, where time (`t`) is incremented from 0 to 10 in steps of 1, and the computer calculates e^{-t} with the function `exp(-t)`. In C and C++, a `for` loop instructs the computer to repeat certain lines of code (enclosed in another set of curly-braces) while a particular variable changes from one value to another. In this particular case the variable in question is `t`, starting at a value of `0.0` and incrementing by `1.0` with each new iteration of the `for` loop. The loop repeats so long as `t` is equal to or less than `10.0`.

Three `cout` lines contained within the `for` loop work together to generate a single line of printed text. The first `cout` prints the time variable (`t`), the next `cout` prints the value of the "falling" variable `x`, and the last `cout` prints the value of the "rising" variable which is simply the complement of `x`. A single `\n` control character is printed with that last `cout` instruction to terminate the line and prepare for a new line of text with the next iteration of the `for` loop.

Finally, the last line of code in this program returns an integer value of zero as the `main` function successfully completes. This is not strictly necessary, but is a good programming practice.

When compiled and executed, this program prints the percentage values of the rising and falling quantities over a timespan of zero to ten time constants (i.e. 0 to 10τ). It is instructive to compare this text output to the lines of code previously shown, and to follow the execution path of the code in your mind. For example, a careful examination of the code should reveal to you why we only see one line reading `Time Falling Rising` but several lines of numbers:

Time	Falling	Rising
0.0000 tau	100.0000%	0.0000%
1.0000 tau	36.7879%	63.2121%
2.0000 tau	13.5335%	86.4665%
3.0000 tau	4.9787%	95.0213%
4.0000 tau	1.8316%	98.1684%
5.0000 tau	0.6738%	99.3262%
6.0000 tau	0.2479%	99.7521%
7.0000 tau	0.0912%	99.9088%
8.0000 tau	0.0335%	99.9665%
9.0000 tau	0.0123%	99.9877%
10.0000 tau	0.0045%	99.9955%

At the very start ($t = 0$) the falling variable is at its maximum value of 100% and the rising variable is at zero. After one time constants' worth of time ($t = \tau$) the falling variable has fallen well over half-way to a value of approximately 36.8% of its initial value. At that same time the rising variable has risen well over half-way to just over 63.2%. After ten time constants' worth of time ($t = 10\tau$) we can see both variables are a small fraction of a percent away from their final values.

The first point in time where we see the variables come within one percent of their final values is at $t = 5\tau$. This is why students are commonly taught that the variables have “settled” for all practical purposes after *five* time constants. It is important to understand, however, that there is nothing truly significant about 5τ . *Five* just happens to be the first whole-number multiple of τ where the variables approach their final values to within the completely arbitrary threshold of 1%. Just to be clear on this point: the rising and falling variables indeed settle to within 1% of their final values after five time constants' worth of time has passed, but their exact value at 5τ is not 1%, nor do they reach the 1% threshold after exactly 5τ , and they certainly are not *fully* settled by then. The rest of this section is devoted to overturning these misconceptions by modeling the inverse exponential functions in greater detail.

We could alter our program to take smaller steps in time, incrementing t by steps of 0.4 instead than by steps of 1, if we wished to take a closer look at exactly when these falling and rising variables come within 1% of their final values. We need only alter one line of code to do this, the conditions controlling execution of the `for` loop:

```
for (t = 0.0; t < 10 ; t = t + 0.4)
```

Now when we re-compile and run this program, we obtain a much more precise table of numbers:

Time	Falling	Rising
0.0000 tau	100.0000%	0.0000%
0.4000 tau	67.0320%	32.9680%
0.8000 tau	44.9329%	55.0671%
1.2000 tau	30.1194%	69.8806%
1.6000 tau	20.1897%	79.8103%
2.0000 tau	13.5335%	86.4665%
2.4000 tau	9.0718%	90.9282%
2.8000 tau	6.0810%	93.9190%
3.2000 tau	4.0762%	95.9238%
3.6000 tau	2.7324%	97.2676%
4.0000 tau	1.8316%	98.1684%
4.4000 tau	1.2277%	98.7723%
4.8000 tau	0.8230%	99.1770%
5.2000 tau	0.5517%	99.4483%
5.6000 tau	0.3698%	99.6302%
6.0000 tau	0.2479%	99.7521%
6.4000 tau	0.1662%	99.8338%
6.8000 tau	0.1114%	99.8886%
7.2000 tau	0.0747%	99.9253%
7.6000 tau	0.0500%	99.9500%
8.0000 tau	0.0335%	99.9665%
8.4000 tau	0.0225%	99.9775%
8.8000 tau	0.0151%	99.9849%
9.2000 tau	0.0101%	99.9899%
9.6000 tau	0.0068%	99.9932%
10.0000 tau	0.0045%	99.9955%

This simulation proves 5τ is not the true point at which we get within 1% of final value, since we see the values come within 1% after only 4.8 time constants.

We may focus our computer simulation even tighter by altering the `for` loop conditions once again for finer resolution (and more limited range, so as to not generate an enormously long table of values):

```
for (t = 4.55; t < 4.7 ; t = t + 0.01)
```

The result, when compiled and run again is as follows:

Time	Falling	Rising
4.5500 tau	1.0567%	98.9433%
4.5600 tau	1.0462%	98.9538%
4.5700 tau	1.0358%	98.9642%
4.5800 tau	1.0255%	98.9745%
4.5900 tau	1.0153%	98.9847%
4.6000 tau	1.0052%	98.9948%
4.6100 tau	0.9952%	99.0048%
4.6200 tau	0.9853%	99.0147%
4.6300 tau	0.9755%	99.0245%
4.6400 tau	0.9658%	99.0342%
4.6500 tau	0.9562%	99.0438%
4.6600 tau	0.9466%	99.0534%
4.6700 tau	0.9372%	99.0628%
4.6800 tau	0.9279%	99.0721%
4.6900 tau	0.9187%	99.0813%

With this closer view we see the values come within 1% of final closer to *4.61* time constants. In other words, 5τ is a rather coarse rounding-off of the true time value for settling within 1%.

Of course, the *best* way to determine how long it takes for these variables to come within 1% of their final values is to algebraically solve for t with the decay function set equal to 0.01 (i.e. 1%):

$$0.01 = e^{-t}$$

$$\ln 0.01 = -t$$

$$t = -\ln 0.01$$

$$t = 4.605170186$$

It should be clear by now that the frequently-taught rule of “settled within five time constants” is a rounded-off estimation based on an arbitrary threshold and should be regarded as such. In any timing applications requiring precision, the rule of 5τ just isn’t good enough and you will need to actually calculate t .

6.4 Modeling inverse exponential growth and decay using Python

This very simple Python script shows the value of the e^{-x} and $1 - e^{-x}$ functions, using an infinite loop to obtain different values of x from the human user of the script:

```
from math import *

while True:
    x = float(input("x = "))
    print("e^-x =", exp(-x))
    print("1 - e^-x =", 1 - exp(-x))
    print(" ")
```

To execute this script on a computer with Python installed, save the script as a plain-text file to some filename such as `script.py` and then from the command line interface type the command `python3 script.py` to execute that script using version 3 of the Python language.

Here is an example showing the script's output with user-entered values of 0, 1, 2, 3.5, and 6 for x :

```
x = 0
e^-x = 1.0
1 - e^-x = 0.0

x = 1
e^-x = 0.36787944117144233
1 - e^-x = 0.6321205588285577

x = 2
e^-x = 0.1353352832366127
1 - e^-x = 0.8646647167633873

x = 3.5
e^-x = 0.0301973834223185
1 - e^-x = 0.9698026165776815

x = 6
e^-x = 0.0024787521766663585
1 - e^-x = 0.9975212478233336
```


The mathematical instruction in Python for raising e to some exponent is the *exponential* function `exp()`. Whatever value is inside the parentheses of this function is the exponent. Therefore, `exp(-1)` instructs the computer to calculate e^{-1} .

6.5 Modeling an energizing inductor using C++

Here is an example C++ program intended to calculate voltage and current for an energizing inductor:

```
#include <iostream>
#include <cmath>
using namespace std;

int main (void)
{
    float v1, r1, l1, i, vr1, vl1, tau, t, incr;

    t = 0.0;
    incr = 20e-6;
    v1 = 10.00;
    r1 = 3.3e3;
    l1 = 100e-3;

    cout << "+---V1---R1---L1----+ " << endl;
    cout << "|                               | " << endl;
    cout << "+-----+ " << endl;

    cout << "V1 = " << v1 << " Volts" << endl;
    cout << "R1 = " << r1 << " Ohms" << endl;
    cout << "L1 = " << l1 << " Henrys" << endl;

    tau = l1 / r1;

    // This "FOR" loop repeats until the resistor's
    // voltage reaches 99% of the source voltage
    for (t = 0.00 ; vr1 < (0.99 * v1) ; t = t + incr)
    {
        vl1 = v1 * exp(-t / tau);
        vr1 = v1 - vl1;
        i = vr1 / r1;

        cout << endl;
        cout << "Time = " << t << " Seconds" << endl;
        cout << " I = " << i << " Amperes" << endl;
        cout << " VR1 = " << vr1 << " Volts" << endl;
        cout << " VL1 = " << vl1 << " Volts" << endl;
    }
}
```

```

    return 0;
}

```

Note the use of a `for` “loop” in the code, commanding the computer to repeat a set of instructions so long as some condition is met. In this case, the condition is that the resistor’s voltage is less than 99% of the source voltage, as described within the text of the code by the *comment* lines (those two lines of text preceded by double forward-slash characters). Note also the use of the *exponential* function (`exp`) used to compute $e^{-\frac{t}{\tau}}$, which requires an additional “include” statement at the beginning of the code (`#include <cmath>`) which gives the compiler access to descriptions of mathematical functions more advanced than those defined within `iostream`.

When compiled and executed, this program generates the following output:

```

+---V1---R1---L1---+
|                   |
+-----+
V1 = 10 Volts
R1 = 3300 Ohms
L1 = 0.1 Henrys

Time = 0 Seconds
  I = 0 Amperes
  VR1 = 0 Volts
  VL1 = 10 Volts

Time = 2e-05 Seconds
  I = 0.00146409 Amperes
  VR1 = 4.83149 Volts
  VL1 = 5.16851 Volts

Time = 4e-05 Seconds
  I = 0.0022208 Amperes
  VR1 = 7.32865 Volts
  VL1 = 2.67135 Volts

Time = 6e-05 Seconds
  I = 0.00261191 Amperes
  VR1 = 8.61931 Volts
  VL1 = 1.38069 Volts

Time = 8e-05 Seconds
  I = 0.00281406 Amperes
  VR1 = 9.28639 Volts

```

VL1 = 0.713613 Volts

Time = 1e-04 Seconds

I = 0.00291854 Amperes

VR1 = 9.63117 Volts

VL1 = 0.368832 Volts

Time = 0.00012 Seconds

I = 0.00297254 Amperes

VR1 = 9.80937 Volts

VL1 = 0.190631 Volts

Time = 0.00014 Seconds

I = 0.00300045 Amperes

VR1 = 9.90147 Volts

VL1 = 0.098528 Volts

6.6 Solving RC and LR circuit differential equations using C++

A *differential equation* is a mathematical equation containing both a variable and at least one of its derivatives (i.e. its rate-of-change). Differential equations are extremely useful for describing a wide range of physical phenomena including the energization and de-energization of energy-storing components such as capacitors and inductors. While differential equations require calculus techniques to solve analytically, their solutions may be approximated quite closely using repeated calculations of simple arithmetic. Computers excel at doing such tasks, and so here we will view multiple programs written in C++ to solve for voltages and currents in resistor-capacitor and resistor-inductor networks.

Each of these programs outputs data in comma-separated variable (CSV) text format, suitable for plotting using mathematical visualization software such as `gnuplot` or spreadsheets such as Microsoft `Excel`.

6.6.1 Differential equation solver for an RC circuit and voltage source

```

#include <iostream>
#include <cmath>
using namespace std;

/* Schematic:
+-----R-----+
|                 |
Vs                 C   Ic = C dVc/dt
|                 |
+-----+
*/

int main (void)
{
    double Ic, C = 33e-6, R = 10e3, Vr, Vs = 0.0, Vc = 10.0, dVc, t, dt = 0.0001;

    cout << "Time , Vc" << endl;

    for (t = 0 ; t < 2.0 ; t = t + dt)
    {
        Vr = Vs - Vc;           // Kirchhoff's Voltage Law
        Ic = Vr / R;           // Ohm's Law
        dVc = dt * (Ic / C);   // I = C dV/dt "Ohm's Law" for capacitors
        Vc = Vc + dVc;         // Integrating capacitor voltage from small changes
        cout << t << " , " << Vc << endl;
    }

    return 0;
}

```

With the initialized values shown in the `double` line of this program, it will compute capacitor voltages at 0.1 millisecond intervals assuming an initial capacitor voltage of 10 Volts and discharging toward zero volts (i.e. the voltage source value `Vs` is set to 0.0). Of course, any combination of `Vs` and initial `Vc` values will work perfectly well, as this program will show the gradual progression from initial to final voltage over time.

6.6.2 Differential equation solver for an RC circuit and current source

```

#include <iostream>
#include <cmath>
using namespace std;

/* Schematic:
+-----+-----+
|       |       |
| Is    | R     | C   I = C dVc/dt
|       |       |
+-----+-----+
*/

int main (void)
{
    double Ic, C = 33e-6, R = 10e3, Is = 1e-3, Ir, Vc = 0.0, dVc, t, dt = 0.0001;

    cout << "Time , Vc" << endl;

    for (t = 0 ; t < 2.0 ; t = t + dt)
    {
        Ir = Vc / R;           // Ohm's Law
        Ic = Is - Ir;         // Kirchhoff's Current Law
        dVc = dt * (Ic / C); // I = C dV/dt "Ohm's Law" for capacitors
        Vc = Vc + dVc;       // Integrating capacitor voltage from small changes
        cout << t << " , " << Vc << endl;
    }

    return 0;
}

```

With the initialized values shown in the `double` line of this program, it will compute capacitor voltages at 0.1 millisecond intervals assuming an initial capacitor voltage of 0 Volts and charging toward 10 Volts (i.e. the source current of 1 milliAmpere passing entirely through the 10 k Ω resistor once the capacitor reaches full voltage). Of course, any combination of source current (I_s) and initial capacitor voltage V_c values will work perfectly well, as this program will show the gradual progression from initial to final voltage over time.

6.6.3 Differential equation solver for an LR circuit and voltage source

```

#include <iostream>
#include <cmath>
using namespace std;

/* Schematic:
+-----R-----+
|                 |
Vs                 L   VL = L dIL/dt
|                 |
+-----+
*/

int main (void)
{
    double VL, L = 50.0, R = 100.0, Vs = 100.0, Vr, IL = 0.0, dIL, t, dt = 0.0001;

    cout << "Time , IL" << endl;

    for (t = 0 ; t < 2.0 ; t = t + dt)
    {
        Vr = IL * R;           // Ohm's Law
        VL = Vs - Vr;       // Kirchhoff's Voltage Law
        dIL = dt * (VL / L); // V = L dI/dt "Ohm's Law" for capacitors
        IL = IL + dIL;     // Integrating inductor current from small changes
        cout << t << " , " << IL << endl;
    }

    return 0;
}

```

With the initialized values shown in the `double` line of this program, it will compute inductor currents at 0.1 millisecond intervals assuming zero initial inductor current and charging toward 1 Ampere (i.e. the voltage source value `Vs` of 100 Volts divided by 100 Ohms). Of course, any combination of `Vs` and initial `IL` values will work perfectly well, as this program will show the gradual progression from initial to final current over time.

6.6.4 Differential equation solver for an LR circuit and current source

```

#include <iostream>
#include <cmath>
using namespace std;

/* Schematic:
+-----+-----+
|       |       |
Is      R      L   Vl = L dI1/dt
|       |       |
+-----+-----+
*/

int main (void)
{
    double Vl, L = 50.0, R = 100.0, Is = 100e-3, Ir, I1 = 0.0, dI1, t, dt = 0.0001;

    cout << "Time , I1" << endl;

    for (t = 0 ; t < 2.0 ; t = t + dt)
    {
        Ir = Is - I1;          // Kirchhoff's Current Law
        Vl = Ir * R;          // Ohm's Law
        dI1 = dt * (Vl / L); // V = L dI/dt "Ohm's Law" for capacitors
        I1 = I1 + dI1;        // Integrating inductor current from small changes
        cout << t << " , " << I1 << endl;
    }

    return 0;
}

```

With the initialized values shown in the `double` line of this program, it will compute inductor currents at 0.1 millisecond intervals assuming zero initial inductor current and charging toward 100 milliAmperes. Of course, any combination of `Is` and initial `I1` values will work perfectly well, as this program will show the gradual progression from initial to final current over time.

Chapter 7

Questions

This learning module, along with all others in the ModEL collection, is designed to be used in an inverted instructional environment where students independently read¹ the tutorials and attempt to answer questions on their own *prior* to the instructor's interaction with them. In place of lecture², the instructor engages with students in Socratic-style dialogue, probing and challenging their understanding of the subject matter through inquiry.

Answers are not provided for questions within this chapter, and this is by design. Solved problems may be found in the Tutorial and Derivation chapters, instead. The goal here is *independence*, and this requires students to be challenged in ways where others cannot think for them. Remember that you always have the tools of *experimentation* and *computer simulation* (e.g. SPICE) to explore concepts!

The following lists contain ideas for Socratic-style questions and challenges. Upon inspection, one will notice a strong theme of *metacognition* within these statements: they are designed to foster a regular habit of examining one's own thoughts as a means toward clearer thinking. As such these sample questions are useful both for instructor-led discussions as well as for self-study.

¹Technical reading is an essential academic skill for any technical practitioner to possess for the simple reason that the most comprehensive, accurate, and useful information to be found for developing technical competence is in textual form. Technical careers in general are characterized by the need for continuous learning to remain current with standards and technology, and therefore any technical practitioner who cannot read well is handicapped in their professional development. An excellent resource for educators on improving students' reading prowess through intentional effort and strategy is the book *Reading For Understanding – How Reading Apprenticeship Improves Disciplinary Learning in Secondary and College Classrooms* by Ruth Schoenbach, Cynthia Greenleaf, and Lynn Murphy.

²Lecture is popular as a teaching method because it is easy to implement: any reasonably articulate subject matter expert can talk to students, even with little preparation. However, it is also quite problematic. A good lecture always makes complicated concepts seem easier than they are, which is bad for students because it instills a false sense of confidence in their own understanding; reading and re-articulation requires more cognitive effort and serves to verify comprehension. A culture of teaching-by-lecture fosters a debilitating dependence upon direct personal instruction, whereas the challenges of modern life demand independent and critical thought made possible only by gathering information and perspectives from afar. Information presented in a lecture is ephemeral, easily lost to failures of memory and dictation; text is forever, and may be referenced at any time.

GENERAL CHALLENGES FOLLOWING TUTORIAL READING

- Summarize as much of the text as you can in one paragraph of your own words. A helpful strategy is to explain ideas as you would for an intelligent child: as simple as you can without compromising too much accuracy.
- Simplify a particular section of the text, for example a paragraph or even a single sentence, so as to capture the same fundamental idea in fewer words.
- Where did the text make the most sense to you? What was it about the text's presentation that made it clear?
- Identify where it might be easy for someone to misunderstand the text, and explain why you think it could be confusing.
- Identify any new concept(s) presented in the text, and explain in your own words.
- Identify any familiar concept(s) such as physical laws or principles applied or referenced in the text.
- Devise a proof of concept experiment demonstrating an important principle, physical law, or technical innovation represented in the text.
- Devise an experiment to disprove a plausible misconception.
- Did the text reveal any misconceptions you might have harbored? If so, describe the misconception(s) and the reason(s) why you now know them to be incorrect.
- Describe any useful problem-solving strategies applied in the text.
- Devise a question of your own to challenge a reader's comprehension of the text.

GENERAL FOLLOW-UP CHALLENGES FOR ASSIGNED PROBLEMS

- Identify where any fundamental laws or principles apply to the solution of this problem, especially before applying any mathematical techniques.
- Devise a thought experiment to explore the characteristics of the problem scenario, applying known laws and principles to mentally model its behavior.
- Describe in detail your own strategy for solving this problem. How did you identify and organized the given information? Did you sketch any diagrams to help frame the problem?
- Is there more than one way to solve this problem? Which method seems best to you?
- Show the work you did in solving this problem, even if the solution is incomplete or incorrect.
- What would you say was the most challenging part of this problem, and why was it so?
- Was any important information missing from the problem which you had to research or recall?
- Was there any extraneous information presented within this problem? If so, what was it and why did it not matter?
- Examine someone else's solution to identify where they applied fundamental laws or principles.
- Simplify the problem from its given form and show how to solve this simpler version of it. Examples include eliminating certain variables or conditions, altering values to simpler (usually whole) numbers, applying a limiting case (i.e. altering a variable to some extreme or ultimate value).
- For quantitative problems, identify the real-world meaning of all intermediate calculations: their units of measurement, where they fit into the scenario at hand. Annotate any diagrams or illustrations with these calculated values.
- For quantitative problems, try approaching it qualitatively instead, thinking in terms of “increase” and “decrease” rather than definite values.
- For qualitative problems, try approaching it quantitatively instead, proposing simple numerical values for the variables.
- Were there any assumptions you made while solving this problem? Would your solution change if one of those assumptions were altered?
- Identify where it would be easy for someone to go astray in attempting to solve this problem.
- Formulate your own problem based on what you learned solving this one.

GENERAL FOLLOW-UP CHALLENGES FOR EXPERIMENTS OR PROJECTS

- In what way(s) was this experiment or project easy to complete?
- Identify some of the challenges you faced in completing this experiment or project.

- Show how thorough documentation assisted in the completion of this experiment or project.
- Which fundamental laws or principles are key to this system's function?
- Identify any way(s) in which one might obtain false or otherwise misleading measurements from test equipment in this system.
- What will happen if (component X) fails (open/shorted/etc.)?
- What would have to occur to make this system unsafe?

7.1 Conceptual reasoning

These questions are designed to stimulate your analytic and synthetic thinking³. In a Socratic discussion with your instructor, the goal is for these questions to prompt an extended dialogue where assumptions are revealed, conclusions are tested, and understanding is sharpened. Your instructor may also pose additional questions based on those assigned, in order to further probe and refine your conceptual understanding.

Questions that follow are presented to challenge and probe your understanding of various concepts presented in the tutorial. These questions are intended to serve as a guide for the Socratic dialogue between yourself and the instructor. Your instructor’s task is to ensure you have a sound grasp of these concepts, and the questions contained in this document are merely a means to this end. Your instructor may, at his or her discretion, alter or substitute questions for the benefit of tailoring the discussion to each student’s needs. The only absolute requirement is that each student is challenged and assessed at a level equal to or greater than that represented by the documented questions.

It is far more important that you convey your *reasoning* than it is to simply convey a correct answer. For this reason, you should refrain from researching other information sources to answer questions. What matters here is that *you* are doing the thinking. If the answer is incorrect, your instructor will work with you to correct it through proper reasoning. A correct answer without an adequate explanation of how you derived that answer is unacceptable, as it does not aid the learning or assessment process.

You will note a conspicuous lack of answers given for these conceptual questions. Unlike standard textbooks where answers to every other question are given somewhere toward the back of the book, here in these learning modules students must rely on other means to check their work. The best way by far is to debate the answers with fellow students and also with the instructor during the Socratic dialogue sessions intended to be used with these learning modules. Reasoning through challenging questions with other people is an excellent tool for developing strong reasoning skills.

Another means of checking your conceptual answers, where applicable, is to use circuit simulation software to explore the effects of changes made to circuits. For example, if one of these conceptual questions challenges you to predict the effects of altering some component parameter in a circuit, you may check the validity of your work by simulating that same parameter change within software and seeing if the results agree.

³*Analytical* thinking involves the “disassembly” of an idea into its constituent parts, analogous to dissection. *Synthetic* thinking involves the “assembly” of a new idea comprised of multiple concepts, analogous to construction. Both activities are high-level cognitive skills, extremely important for effective problem-solving, necessitating frequent challenge and regular practice to fully develop.

7.1.1 Reading outline and reflections

“Reading maketh a full man; conference a ready man; and writing an exact man” – Francis Bacon

Francis Bacon’s advice is a blueprint for effective education: reading provides the learner with knowledge, writing focuses the learner’s thoughts, and critical dialogue equips the learner to confidently communicate and apply their learning. Independent acquisition and application of knowledge is a powerful skill, well worth the effort to cultivate. To this end, students should read these educational resources closely, journal their own reflections on the reading, and discuss in detail their findings with classmates and instructor(s). You should be able to do all of the following after reading any instructional text:

Briefly **SUMMARIZE THE TEXT** in the form of a journal entry documenting your learning as you progress through the course of study. Share this summary in dialogue with your classmates and instructor. Journaling is an excellent self-test of thorough reading because you cannot clearly express what you have not read or did not comprehend.

Demonstrate **ACTIVE READING STRATEGIES**, including verbalizing your impressions as you read, simplifying long passages to convey the same ideas using fewer words, annotating text and illustrations with your own interpretations, working through mathematical examples shown in the text, cross-referencing passages with relevant illustrations and/or other passages, identifying problem-solving strategies applied by the author, etc. Technical reading is a special case of problem-solving, and so these strategies work precisely because they help solve any problem: paying attention to your own thoughts (metacognition), eliminating unnecessary complexities, identifying what makes sense, paying close attention to details, drawing connections between separated facts, and noting the successful strategies of others.

Identify **IMPORTANT THEMES**, especially **GENERAL LAWS** and **PRINCIPLES**, expounded in the text and express them in the simplest of terms as though you were teaching an intelligent child. This emphasizes connections between related topics and develops your ability to communicate complex ideas to anyone.

Form **YOUR OWN QUESTIONS** based on the reading, and then pose them to your instructor and classmates for their consideration. Anticipate both correct and incorrect answers, the incorrect answer(s) assuming one or more plausible misconceptions. This helps you view the subject from different perspectives to grasp it more fully.

Devise **EXPERIMENTS** to test claims presented in the reading, or to disprove misconceptions. Predict possible outcomes of these experiments, and evaluate their meanings: what result(s) would confirm, and what would constitute disproof? Running mental simulations and evaluating results is essential to scientific and diagnostic reasoning.

Specifically identify any points you found **CONFUSING**. The reason for doing this is to help diagnose misconceptions and overcome barriers to learning.

7.1.2 Foundational concepts

Correct analysis and diagnosis of electric circuits begins with a proper understanding of some basic concepts. The following is a list of some important concepts referenced in this module's full tutorial. Define each of them in your own words, and be prepared to illustrate each of these concepts with a description of a practical example and/or a live demonstration.

Thought experiments as a problem-solving strategy

Magnetic field

Energy

Conservation of Energy

Inductance

Permeability

Inductor

Electrical source

Electrical load

"Ohm's Law" for inductance

Current stabilization

Ohm's Law

Kirchhoff's Current Law

Time constant

Magnetic hysteresis

Magnetic saturation

Inductor current limit

Parasitic capacitance


Inductor physical size

Series versus parallel inductances

Mutual inductance

7.1.3 Magnetic flux lines

Draw the pattern of the magnetic field produced by electric current through a straight wire and through a wire coil:



Challenges

- Demonstrate how the “right hand rule” applies to these scenarios.

7.1.4 Wirewound resistor

One way to manufacture a precision resistor is to precisely cut a length of wire having a known amount of resistance per unit length, and package that wire into a tightly-wound coil in order to conserve space. However, wrapping a wire into a coil shape will create inductance, and in the case of a precision resistor we *only* want resistance.

Devise a method for wrapping resistive wire in such a manner that it will exhibit negligible inductance.

Suppose you needed to determine whether or not a wirewound resistor was non-inductive prior to installing it in a precision circuit, but you had no part number or other reference with which to research the design or specifications of this resistor. All you have is the resistor, a multimeter, and basic electronic lab equipment (power supplies, breadboards, etc.). Devise a procedure for testing the inductance (or non-inductance) of this resistor.

Challenges

- If you were tasked with forming a custom-made precision resistor in this manner, how would you proceed? What information would you need to know, and what steps would you follow to build the resistor?

7.1.5 Impossibility of storing an energized inductor

Capacitors and inductors are both capable of storing energy, of acting as either electrical sources or electrical loads. For example, it is possible to energize (“charge”) a capacitor, then place that capacitor on a shelf for a long period of time, and later demonstrate that some energy is still retained inside the capacitor.

Why is this practically impossible to do with an inductor: “charge” it up, place it on a shelf, and some time later measure the presence of stored energy within that inductor?


Challenges

- Suppose finances and technology were not barriers. Could such an experiment be run to prove the concept of storing energy within an inductor for some indefinite period of time?

7.1.6 Inductor as source versus load


Suppose an inductor is connected directly to an adjustable-current source, and the current of that source is steadily *increased* over time. We know that an increasing current through an inductor will produce a magnetic field of increasing strength. Does this increase in magnetic field constitute an *accumulation* of energy in the inductor, or a *release* of energy from the inductor? In this scenario, does the inductor act as a *load* or as a *source* of electrical energy?

Current *increasing*



Now, suppose the adjustable current source is steadily *decreased* over time. We know this will result in a magnetic field of decreasing strength in the inductor. Does this decrease in magnetic field constitute an *accumulation* of energy in the inductor, or a *release* of energy from the inductor? In this scenario, does the inductor act as a *load* or as a *source* of electrical energy?

Current *decreasing*



For each of these scenarios, label the direction of current in the circuit.


Challenges

- Identify factors influencing the voltage's *magnitude* in each circuit.

7.1.7 Inductor charging circuit

A practical strategy for quantitatively analyzing inductor circuits is to first *qualitatively* determine the trajectories of energy, voltage, and current values over time based on first principles (e.g. Ohm's Law, Kirchhoff's Laws, and the fact that inductive energy storage is a function of inductor current). Once you have done this, you will better know which equation(s) to use for quantitative analysis. A common mistake made in these types of problems is to plug numbers into equations without first reasoning *why* those equations are appropriate, so the point of this exercise is to build sound problem-solving habits!

Try qualitatively determining the following parameters in this simple resistor-inductor circuit at three different times: (1) just before the switch closes, (2) at the instant the switch contacts touch, and (3) after the switch has been closed for a long time. Assume that the inductor begins in a completely de-energized state:



Express your answers qualitatively: “maximum,” “minimum,” or perhaps “zero” if you know that to be the case.

Before the switch closes:

$$E_L \text{ (energy stored in inductor)} =$$

$$V_L =$$

$$V_R =$$

$$V_{\text{switch}} =$$

$$I =$$

At the instant of switch closure:

$$E_L \text{ (energy stored in inductor)} =$$

$$V_L =$$

$$V_R =$$

$$V_{\text{switch}} =$$

$$I =$$

Long after the switch has closed:

E_L (energy stored in inductor) =

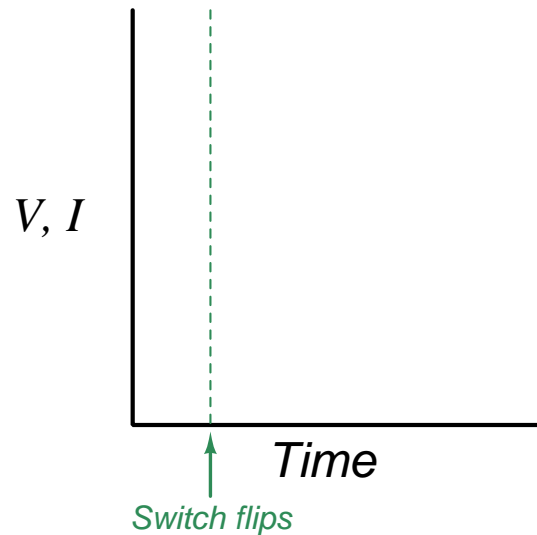
$V_L =$

$V_R =$

$V_{switch} =$

$I =$

Also sketch a qualitative graph showing V_L and I in this circuit over time:



Note: you may find the following equations conceptually useful:

$$E_L = \frac{1}{2}LI^2 \quad V = L\frac{dI}{dt}$$

Also, identify how the following foundational concepts apply to this circuit:

- How does the Conservation of Energy help explain the inductor's current at the instant of switch closure?
- How does Kirchhoff's Voltage Law apply to the circuit before the switch closes?
- How does Kirchhoff's Voltage Law apply to the circuit after the switch closes?

Challenges


- Explain how the *Conservation of Energy* relates to changes in *voltage* and *current* for an inductor.

- In this experiment, is the inductor absorbing energy or releasing energy? How can we tell?
- Explain what will happen in the circuit when the switch is re-opened after the switch has been closed for a long period of time.
- Explain how drawing multiple versions of a circuit and annotating conditions before and after switch actuation is a helpful problem-solving strategy for any time-delay circuit.

7.1.8 A shocking experience

Many years ago, I decided to experiment with electromagnetism by making an electromagnet out of a spool of wire. I placed a steel bolt through the center of the spool so as to have a core of high permeability, and passed current from a battery through the wire to make a magnetic field. Not having any “jumper” wires, I held the wire ends of the spool in contact with the 9-volt battery terminals, one in each hand.

The electromagnet worked just fine, and I was able to move some steel paper clips with the magnetic field generated by it. However, when I broke the circuit by releasing one of the wire ends from the battery terminal it was touching, I received a small electric shock! Shown here is a schematic diagram of me, in the circuit:



At the time, I didn't understand how inductance worked. I only understood how to make magnetism with electricity, but I didn't realize a coil of wire could generate (high voltage!) electricity from its own magnetic field. I did know, however, that the 9 Volts output by the battery was much too weak to shock me (yes, I touched the battery terminals directly to verify this fact), so *something* in the circuit must have generated a voltage greater than 9 Volts.

If you had been there to explain what just happened to the younger version of me, what would you say?


Challenges

- Describe a safer way to conduct this experiment.
- Do you suppose this experiment could have proven lethal under the right circumstances? Explain why or why not.

7.1.9 Delayed relay action


Electromechanical relays are extremely useful devices, but they have their idiosyncrasies. One of them is a consequence of the fact that the relay's coil acts as an inductor, storing energy in its magnetic field.

In the following circuit, a relay is used to switch power to a large electric motor, while being controlled by a light-duty pushbutton switch:



The problem here is every time the pushbutton switch is released, the contacts arc significantly. This happens because the inductor releases all of its stored energy as a high-voltage spark across the opening contacts. *Inductive kickback* is the phrase commonly used to describe this effect, and over time it will prematurely destroy the switch.

An electronics technician understands the nature of the problem and proposes a solution. By connecting a light bulb in parallel with the relay coil, the coil's energy now has a safer place to dissipate whenever the pushbutton switch contacts open:



Instead of that stored energy manifesting itself as a high-voltage arc at the switch, it powers the light bulb for a brief time after the switch opens, dissipating in a non-destructive manner.

However, the addition of the light bulb introduces a new, unexpected problem to the circuit. Now, when the pushbutton switch is released, the relay delays for a fraction of a second before disengaging. This causes the motor to “overshoot” its position instead of stopping when it is supposed to.

Explain why this happens, with reference to the LR time constant of this circuit before and after the addition of the lamp.

Challenges


- ???.
- ???.
- ???.

7.1.10 Explaining the meaning of calculations

An unfortunate tendency among beginning students in any quantitative discipline is to perform calculations without regard for the real-world meanings of the values, and also to follow mathematical formulae without considering the general principles embodied in each. To ignore concepts while performing calculations is a serious error for a variety of reasons, not the least of which being an increased likelihood of computing results that turn out to be nonsense.

In the spirit of honoring concepts, I present to you a quantitative problem where all the calculations have been done for you, but all variable labels, units, and other identifying data have been stripped away. Your task is to *assign proper meaning* to each of the numbers, identify the correct unit of measurement in each case, apply any appropriate metric prefixes to those values, explain the significance of each value by describing where it “fits” into the circuit being analyzed, and identify the general principle employed at each step.

Here is the schematic diagram of the resistor-inductor circuit:



Here are all the calculations performed in order from first to last. Conditions for these calculations are that the switch began in the position shown and was left that way for a long time, after which it was toggled to the opposite position at time $t = 0$ seconds ; all voltages and currents were calculated at $t = 20$ microseconds after toggling:

1. $\frac{18}{1000} = 18 \times 10^{-3}$
2. $\frac{330 \times 10^{-3}}{1000 + 10000} = 3 \times 10^{-5}$
3. $(18 \times 10^{-3}) \times (1000 + 10000) = 198.00$
4. $e^{\frac{-20 \times 10^{-6}}{30 \times 10^{-6}}} = 0.513417$
5. $198 \times 0.513417 = 101.657$
6. $(18 \times 10^{-3}) \times 0.513417 = 9.24151 \times 10^{-3}$

Explain what each value means in the circuit, identify its unit of measurement and appropriate metric prefix, and identify the general principle used to compute it!

Challenges

- How is it possible for the inductor to momentarily develop more voltage than the DC source itself outputs?

7.2 Quantitative reasoning

These questions are designed to stimulate your computational thinking. In a Socratic discussion with your instructor, the goal is for these questions to reveal your mathematical approach(es) to problem-solving so that good technique and sound reasoning may be reinforced. Your instructor may also pose additional questions based on those assigned, in order to observe your problem-solving firsthand.

Mental arithmetic and estimations are strongly encouraged for all calculations, because without these abilities you will be unable to readily detect errors caused by calculator misuse (e.g. keystroke errors).

You will note a conspicuous lack of answers given for these quantitative questions. Unlike standard textbooks where answers to every other question are given somewhere toward the back of the book, here in these learning modules students must rely on other means to check their work. My advice is to use circuit simulation software such as SPICE to check the correctness of quantitative answers. Refer to those learning modules within this collection focusing on SPICE to see worked examples which you may use directly as practice problems for your own study, and/or as templates you may modify to run your own analyses and generate your own practice problems.

Completely worked example problems found in the Tutorial may also serve as “test cases⁴” for gaining proficiency in the use of circuit simulation software, and then once that proficiency is gained you will never need to rely⁵ on an answer key!

⁴In other words, set up the circuit simulation software to analyze the same circuit examples found in the Tutorial. If the simulated results match the answers shown in the Tutorial, it confirms the simulation has properly run. If the simulated results disagree with the Tutorial’s answers, something has been set up incorrectly in the simulation software. Using every Tutorial as practice in this way will quickly develop proficiency in the use of circuit simulation software.

⁵This approach is perfectly in keeping with the instructional philosophy of these learning modules: *teaching students to be self-sufficient thinkers*. Answer keys can be useful, but it is even more useful to your long-term success to have a set of tools on hand for checking your own work, because once you have left school and are on your own, there will no longer be “answer keys” available for the problems you will have to solve.

7.2.1 Miscellaneous physical constants

Note: constants shown in **bold** type are *exact*, not approximations. Values inside of parentheses show one standard deviation (σ) of uncertainty in the final digits: for example, the magnetic permeability of free space value given as $1.25663706212(19) \times 10^{-6}$ H/m represents a center value (i.e. the location parameter) of $1.25663706212 \times 10^{-6}$ Henrys per meter with one standard deviation of uncertainty equal to $0.0000000000019 \times 10^{-6}$ Henrys per meter.

Avogadro's number (N_A) = **6.02214076** $\times 10^{23}$ **per mole** (mol⁻¹)

Boltzmann's constant (k) = **1.380649** $\times 10^{-23}$ **Joules per Kelvin** (J/K)

Electronic charge (e) = **1.602176634** $\times 10^{-19}$ **Coulomb** (C)

Faraday constant (F) = **96,485.33212...** $\times 10^4$ **Coulombs per mole** (C/mol)

Magnetic permeability of free space (μ_0) = **1.25663706212(19)** $\times 10^{-6}$ Henrys per meter (H/m)

Electric permittivity of free space (ϵ_0) = **8.8541878128(13)** $\times 10^{-12}$ Farads per meter (F/m)

Characteristic impedance of free space (Z_0) = **376.730313668(57)** Ohms (Ω)

Gravitational constant (G) = **6.67430(15)** $\times 10^{-11}$ cubic meters per kilogram-seconds squared (m³/kg-s²)

Molar gas constant (R) = **8.314462618...** **Joules per mole-Kelvin** (J/mol-K) = 0.08205746(14) liters-atmospheres per mole-Kelvin

Planck constant (h) = **6.62607015** $\times 10^{-34}$ **joule-seconds** (J-s)

Stefan-Boltzmann constant (σ) = **5.670374419...** $\times 10^{-8}$ **Watts per square meter-Kelvin⁴** (W/m²·K⁴)

Speed of light in a vacuum (c) = **299,792,458 meters per second** (m/s) = 186282.4 miles per second (mi/s)

Note: All constants taken from NIST data "Fundamental Physical Constants – Complete Listing", from <http://physics.nist.gov/constants>, National Institute of Standards and Technology (NIST), 2018 CODATA Adjustment.

7.2.2 Introduction to spreadsheets

A powerful computational tool you are encouraged to use in your work is a *spreadsheet*. Available on most personal computers (e.g. Microsoft Excel), *spreadsheet* software performs numerical calculations based on number values and formulae entered into cells of a grid. This grid is typically arranged as lettered columns and numbered rows, with each cell of the grid identified by its column/row coordinates (e.g. cell B3, cell A8). Each cell may contain a string of text, a number value, or a mathematical formula. The spreadsheet automatically updates the results of all mathematical formulae whenever the entered number values are changed. This means it is possible to set up a spreadsheet to perform a series of calculations on entered data, and those calculations will be re-done by the computer any time the data points are edited in any way.

For example, the following spreadsheet calculates average speed based on entered values of distance traveled and time elapsed:

	A	B	C	D
1	Distance traveled	46.9	Kilometers	
2	Time elapsed	1.18	Hours	
3	Average speed	= B1 / B2	km/h	
4				
5				

Text labels contained in cells A1 through A3 and cells C1 through C3 exist solely for readability and are not involved in any calculations. Cell B1 contains a sample distance value while cell B2 contains a sample time value. The formula for computing speed is contained in cell B3. Note how this formula begins with an “equals” symbol (=), references the values for distance and speed by lettered column and numbered row coordinates (B1 and B2), and uses a forward slash symbol for division (/). The coordinates B1 and B2 function as *variables*⁶ would in an algebraic formula.

When this spreadsheet is executed, the numerical value 39.74576 will appear in cell B3 rather than the formula = B1 / B2, because 39.74576 is the computed speed value given 46.9 kilometers traveled over a period of 1.18 hours. If a different numerical value for distance is entered into cell B1 or a different value for time is entered into cell B2, cell B3’s value will automatically update. All you need to do is set up the given values and any formulae into the spreadsheet, and the computer will do all the calculations for you.

Cell B3 may be referenced by other formulae in the spreadsheet if desired, since it is a variable just like the given values contained in B1 and B2. This means it is possible to set up an entire chain of calculations, one dependent on the result of another, in order to arrive at a final value. The arrangement of the given data and formulae need not follow any pattern on the grid, which means you may place them anywhere.

⁶Spreadsheets may also provide means to attach text labels to cells for use as variable names (Microsoft Excel simply calls these labels “names”), but for simple spreadsheets such as those shown here it’s usually easier just to use the standard coordinate naming for each cell.

Common⁷ arithmetic operations available for your use in a spreadsheet include the following:

- Addition (+)
- Subtraction (-)
- Multiplication (*)
- Division (/)
- Powers (^)
- Square roots (sqrt())
- Logarithms (ln() , log10())

Parentheses may be used to ensure⁸ proper order of operations within a complex formula. Consider this example of a spreadsheet implementing the *quadratic formula*, used to solve for roots of a polynomial expression in the form of $ax^2 + bx + c$:

$$x = \frac{-b \pm \sqrt{b^2 - 4ac}}{2a}$$

	A	B
1	x_1	= (-B4 + sqrt((B4^2) - (4*B3*B5))) / (2*B3)
2	x_2	= (-B4 - sqrt((B4^2) - (4*B3*B5))) / (2*B3)
3	a =	9
4	b =	5
5	c =	-2

This example is configured to compute roots⁹ of the polynomial $9x^2 + 5x - 2$ because the values of 9, 5, and -2 have been inserted into cells B3, B4, and B5, respectively. Once this spreadsheet has been built, though, it may be used to calculate the roots of *any* second-degree polynomial expression simply by entering the new a , b , and c coefficients into cells B3 through B5. The numerical values appearing in cells B1 and B2 will be automatically updated by the computer immediately following any changes made to the coefficients.

⁷Modern spreadsheet software offers a bewildering array of mathematical functions you may use in your computations. I recommend you consult the documentation for your particular spreadsheet for information on operations other than those listed here.

⁸Spreadsheet programs, like text-based programming languages, are designed to follow standard order of operations by default. However, my personal preference is to use parentheses even where strictly unnecessary just to make it clear to any other person viewing the formula what the intended order of operations is.

⁹Reviewing some algebra here, a *root* is a value for x that yields an overall value of zero for the polynomial. For this polynomial ($9x^2 + 5x - 2$) the two roots happen to be $x = 0.269381$ and $x = -0.82494$, with these values displayed in cells B1 and B2, respectively upon execution of the spreadsheet.

Alternatively, one could break up the long quadratic formula into smaller pieces like this:

$$y = \sqrt{b^2 - 4ac} \quad z = 2a$$

$$x = \frac{-b \pm y}{z}$$

	A	B	C
1	x_1	= (-B4 + C1) / C2	= sqrt((B4^2) - (4*B3*B5))
2	x_2	= (-B4 - C1) / C2	= 2*B3
3	a =	9	
4	b =	5	
5	c =	-2	

Note how the square-root term (y) is calculated in cell C1, and the denominator term (z) in cell C2. This makes the two final formulae (in cells B1 and B2) simpler to interpret. The positioning of all these cells on the grid is completely arbitrary¹⁰ – all that matters is that they properly reference each other in the formulae.

Spreadsheets are particularly useful for situations where the same set of calculations representing a circuit or other system must be repeated for different initial conditions. The power of a spreadsheet is that it automates what would otherwise be a tedious set of calculations. One specific application of this is to simulate the effects of various components within a circuit failing with abnormal values (e.g. a shorted resistor simulated by making its value nearly zero; an open resistor simulated by making its value extremely large). Another application is analyzing the behavior of a circuit design given new components that are out of specification, and/or aging components experiencing drift over time.

¹⁰My personal preference is to locate all the “given” data in the upper-left cells of the spreadsheet grid (each data point flanked by a sensible name in the cell to the left and units of measurement in the cell to the right as illustrated in the first distance/time spreadsheet example), sometimes coloring them in order to clearly distinguish which cells contain entered data versus which cells contain computed results from formulae. I like to place all formulae in cells below the given data, and try to arrange them in logical order so that anyone examining my spreadsheet will be able to figure out *how* I constructed a solution. This is a general principle I believe all computer programmers should follow: *document and arrange your code to make it easy for other people to learn from it.*

7.2.3 Number of time constants

Determine the number of time constants (τ) that 7.5 seconds is equal to in each of the following resistor-inductor circuits:

- $R = 100 \Omega$, $L = 50 \text{ mH}$; 7.5 sec =
- $R = 45 \Omega$, $L = 2.2 \text{ H}$; 7.5 sec =
- $R = 1 \text{ k}\Omega$, $L = 725 \text{ mH}$; 7.5 sec =
- $R = 4.7 \text{ k}\Omega$, $L = 325 \text{ mH}$; 7.5 sec =
- $R = 6.2 \Omega$, $L = 25 \text{ H}$; 7.5 sec =

Challenges

- For each of these LR networks, calculate the percentage of decay for voltage or current in a de-energizing condition following 7.5 seconds after disconnection of the source.


7.2.4 Inverse exponential functions

The following two expressions are frequently used to calculate values of changing variables (voltage and current) in RC and LR timing circuits:

$$e^{-\frac{t}{\tau}} \quad \text{or} \quad 1 - e^{-\frac{t}{\tau}}$$

One of these expressions describes the percentage that a changing value in an RC or LR circuit *has gone* from the starting time. The other expression describes how far that same variable *has left to go* before it reaches its ultimate value (at $t = \infty$).

The question is, which expression represents which quantity? This is often a point of confusion, because students have a tendency to try to correlate these expressions to the quantities by rote memorization. Does the expression $e^{-\frac{t}{\tau}}$ represent the amount a variable has changed, or how far it has left to go until it stabilizes? What about the other expression $1 - e^{-\frac{t}{\tau}}$? More importantly, *how can we figure this out so we don't have to rely on memory?*



Challenges

- A useful mathematical problem-solving technique is to set the independent variable in this case, t , to some simple value such as 0 or 1 to see what happens. What would setting $t = 0$ simulate in either of these functions?

7.2.5 Inductor color codes

Small inductors often look like resistors, even to the extent that their values are designated by colored bands. For example, an inductor sporting the colors Blu, Red, Brn, Gld would have a value of $620 \mu\text{H}$, $\pm 5\%$.

Determine the values of the following inductors (express the tolerance as a percentage), based on their color codes:

- Red, Grn, Brn, Gld
- Wht, Org, Red, Sil
- Grn, Gry, Blk
- Vio, Blu, Org, Gld

Challenges

- Identify a significant difference between the standard resistor color code and the inductor color code.

7.2.6 Voltage versus rate-of-current-change

A 5 H inductor is subjected to an electric current that changes at a rate of 4.5 Amperes per second. How much voltage will be dropped by the inductor?

Now suppose two 5 H inductors connected in series are subjected to the same rate-of-change of current. How much voltage will be dropped across the series combination?


Now suppose two 5 H inductors connected in parallel are subjected to the same rate-of-change of current. How much voltage will be dropped across the parallel combination?

Challenges

- Based on the voltages calculated for each inductor network, prove that L adds in series and diminishes in parallel.

7.2.7 Inductance of three-inductor network

Calculate the total inductance in this collection of inductors, as measured between the two wires:




Challenges

- If inductor L_2 fails open, what is the effect on total inductance?
- If inductor L_3 fails shorted, what is the effect on total inductance?

7.2.8 Unknown inductor value

How large must inductor L_x be in order to provide a total inductance of 2.5 H in this network of inductors?



Challenges

- The physical proximity and orientation of these inductors to each other actually matters. Explain why.

7.2.9 Resistance for specified time constant

Calculate the resistor value which when connected in series with a 75 mH inductor will provide a *time constant* (τ) of 20 microseconds. Express your answer in the form of a five-band precision resistor color code (with a tolerance of $\pm 0.25\%$).

Challenges

- What advantage does a five-band resistor color code enjoy over a four-band?

7.2.10 Cylindrical inductor wire turns

Calculate how many turns of wire must be wrapped around a hollow, non-magnetic (air) solenoid core 2 cm in diameter and 10 cm in length in order to create an inductance of 22 mH. You may use the permeability of free space (μ_0) for the μ value of the air core.

Next, calculate the required number of turns to produce the same inductance with a solid iron core of the same dimensions, assuming that the iron has a relative permeability (μ_r) of 4000 (i.e. the magnetic permeability of iron is 4000 times greater than that of open space).

Challenges

- A common mistake is to calculate a value of 0.04184 turns for the iron-core inductor. Explain how someone might make this mistake.

7.2.11 Toroidal inductor wire turns


Calculate how many turns of wire must be wrapped around a toroidal, iron-powder inductor core to make it exhibit 2.5 μH of inductance. Assume the powdered-iron toroidal core is a Micrometals part number T50-66 with an A_L value of 29.0 nanoHenrys per turn-squared.

Challenges


- Identify sources of parasitic (or “stray”) resistance and capacitance in this inductor.

7.2.12 SPDT inductor switching circuit

Calculate the final value of current through the inductor with the switch in the left-hand position (assuming that many time constants' worth of time have passed):



Now, assume that the switch is *instantly* moved to the right-hand position. How much voltage will the inductor initially drop?




Explain why this voltage is so very different from the supply voltage. What practical uses might a circuit such as this have?

Challenges

- Suppose this circuit were built and tested, and it was found that the voltage developed across the inductor at the moment the switch moved to the right-hand position far exceeded the value you calculated. Identify some possible problems in the circuit which could account for this excessive voltage.

7.2.13 Tabulating inductor charging

Determine the inductor voltage and inductor current at the specified times (time $t = 0$ microseconds being the exact moment the switch contacts close):



- Time = $0 \mu\text{s}$; $V_L = \underline{\hspace{2cm}}$; $I_L = \underline{\hspace{2cm}}$
- Time = $10 \mu\text{s}$; $V_L = \underline{\hspace{2cm}}$; $I_L = \underline{\hspace{2cm}}$
- Time = $20 \mu\text{s}$; $V_L = \underline{\hspace{2cm}}$; $I_L = \underline{\hspace{2cm}}$
- Time = $30 \mu\text{s}$; $V_L = \underline{\hspace{2cm}}$; $I_L = \underline{\hspace{2cm}}$
- Time = $40 \mu\text{s}$; $V_L = \underline{\hspace{2cm}}$; $I_L = \underline{\hspace{2cm}}$
- Time = $50 \mu\text{s}$; $V_L = \underline{\hspace{2cm}}$; $I_L = \underline{\hspace{2cm}}$

Challenges

- Explain why it is a good problem-solving strategy to first sketch a graph showing the expected trajectories of voltage and current over time before performing any calculations.
- How would the voltage and current values be affected if this circuit had *two* 25 mH inductors in parallel with each other instead of just the one in the circuit?
- How would the voltage and current values be affected if this circuit had *two* 25 mH inductors in series with each other instead of just the one in the circuit?


7.2.14 Time required to reach specified currents

Calculate the amount of time it takes for the current through a 33 mH inductor to grow from 0 milliAmperes to 50 milliAmperes, if powered by a 12 Volt battery connected in series through a 200 Ω resistor.

Calculate the amount of time it takes for a 100 mH inductor to de-energize from 5 Amperes to 2 Amperes while connected to nothing but a 1.8 Ω resistor.

A 3.5 H inductor begins in an energized state of 75 milliAmperes, and de-energizes through a 1 k Ω resistor. How long will it take before the inductor's voltage will fall to a relatively safe value (30 Volts or less)?

Determine the amount of time needed for the resistor voltage (V_R) to rise to the specified levels after the switch is thrown to the "energize" position, assuming an inductor with no internal resistance:




- 1 Volt =
- 2 Volts =
- 4 Volts =
- 6 Volts =
- 8 Volts =

Challenges

- What will happen to all of these time values if the resistor in each circuit is increased in value?

7.2.15 Measuring inductance by time delay

An inductor of unknown inductance is measured with an ohmmeter and found to have a coil resistance of 85Ω . It is then connected in series with a 970Ω resistor, energized by a 19-Volt DC source, and then voltage across the 970Ω resistor is measured and plotted over time:



Calculate the approximate size of the inductor in milliHenrys based on the information you see here.

Challenges

- A common tendency among students new to electric circuit analysis is to first “think mathematically” rather than first “think conceptually”. Before plugging any numbers into any formulae, one should always identify electrical principles such as sources and loads, directions of current, and voltage polarities, usually sketching a diagram of a circuit to see where all these

concepts fit. Demonstrate how to “think conceptually” when solving this problem, rather than immediately trying to compute values.

- Why are all the voltage values on this plot negative rather than positive?
- How would the plot of voltage over time change if the resistor were made smaller?
- What happened to the energy previously stored in the inductor?

7.2.16 Power line reactors

Volume 18, Number 1 issue of *The Electric Journal* (circa 1921) contains an article written by G. G. Grissinger of the Westinghouse Electric and Manufacturing Company entitled “High-Speed Circuit Breakers” in which we find the following quote on page 65:

It is the function of the high-speed circuit breaker shown in Fig. 1, to open an electric circuit after a short-circuit occurs, so quickly, that the current will be unable to reach a dangerously high value. To the eye the ordinary carbon circuit breaker opens a short-circuit with great rapidity, yet the current is able to reach the maximum value before the circuit breaker starts to open. Oscillograms show that the current on short-circuit, starting from zero value, builds up at the rate of from one million to three million amperes per second, depending upon the constants of the circuit. This means that a value of 10000 amperes would be reached in from 0.01 to 0.003 seconds.


Assuming a peak voltage of 44 kiloVolts, how large of an inductor connected in series with the line would be required to limit the build-up rate of fault current to a mere 500,000 Amperes per second?

Challenges

- Explain the meaning of the following phrase from this article: “*depending upon the constants of the circuit*”.
- Power circuit breakers of that era used contacts made of carbon rather than metal. Explain why.

7.2.17 SPICE analysis of an energizing inductor

The following resistor-inductor circuit will be simulated by a computer program called SPICE. The text-based “netlist” used to describe the circuit to SPICE is shown immediately to the right of the schematic diagram. Numbered points on the diagram identify “nodes” used by SPICE to declare how the components connect to one another:



```
* SPICE circuit
v1 1 0 dc 15
r1 1 2 5100
l1 2 0 33e-3 ic=0
.tran 2e-6 30e-6 uic
.width out=80
.plot tran v(2)
.end
```

In this analysis, SPICE will plot inductor voltage (between node 2 and ground) over a timespan of 30 microseconds in increments of 2 microseconds each. The inductor begins (i.e. “initial condition” $ic=$) with zero current through it, which means no stored energy. When run, the simulation results are as follows:

Legend: + = v(2)

time	v(2)	0.00e+00	1.00e+01	2.00e+01
0.000e+00	0.000e+00	+	.	.
2.000e-06	1.101e+01	.	. +	.
4.000e-06	8.087e+00	.	+	.
6.000e-06	5.935e+00	.	+	.
8.000e-06	4.358e+00	.	+	.
1.000e-05	3.197e+00	.	+	.
1.200e-05	2.347e+00	.	+	.
1.400e-05	1.723e+00	.	+	.
1.600e-05	1.264e+00	.	+	.
1.800e-05	9.277e-01	.	+	.
2.000e-05	6.812e-01	.	+	.
2.200e-05	4.998e-01	.	+	.
2.400e-05	3.668e-01	.	+	.
2.600e-05	2.693e-01	.	+	.
2.800e-05	1.976e-01	.	+	.
3.000e-05	1.449e-01	.	+	.
time	v(2)	0.00e+00	1.00e+01	2.00e+01


Verify any of the simulated inductor voltage values using your own manual calculations.

Challenges

- Explain how you may use a SPICE simulation such as this for your own practice, instead of relying on questions and answers provided to you in a textbook.
- Explain how you could use this simulation to practice calculating inductor or resistor values from the simulated results. In other words, imagine this was real data from a real inductor-resistor circuit experiment, and you needed to compute either the resistor value or the inductor value from all the other data.

7.2.18 SPICE analysis of a de-energizing inductor

The following resistor-inductor circuit will be simulated by a computer program called SPICE. The text-based “netlist” used to describe the circuit to SPICE is shown immediately to the right of the schematic diagram. Numbered points on the diagram identify “nodes” used by SPICE to declare how the components connect to one another:



```
* SPICE circuit
vamm1 1 0
r1 1 2 30
l1 0 2 0.5 ic=4
.tran 5e-3 75e-3 uic
.width out=80
.plot tran i(vamm1)
.end
```

In this analysis, SPICE will plot circuit current (measured using the “dummy” voltage source V_{amm1}) over a timespan of 75 milliseconds in increments of 5 milliseconds each. The inductor begins (i.e. “initial condition” $ic=$) with 4 Amperes of current through it, with no other energy source in the circuit. When run, the simulation results are as follows:

Legend: + = vamm1#branch

time	vamm1#br	0.00e+00	1.00e+00	2.00e+00	3.00e+00	4.00e+00
0.000e+00	0.000e+00	+
5.000e-03	2.963e+00	.	.	.	+	.
1.000e-02	2.196e+00	.	.	+	.	.
1.500e-02	1.626e+00	.	.	+	.	.
2.000e-02	1.205e+00	.	.	+	.	.
2.500e-02	8.923e-01	.	+	.	.	.
3.000e-02	6.609e-01	.	+	.	.	.
3.500e-02	4.897e-01	.	+	.	.	.
4.000e-02	3.625e-01	.	+	.	.	.
4.500e-02	2.685e-01	.	+	.	.	.
5.000e-02	1.990e-01	.	+	.	.	.
5.500e-02	1.473e-01	.	+	.	.	.
6.000e-02	1.091e-01	.	+	.	.	.
6.500e-02	8.085e-02	.	+	.	.	.
7.000e-02	5.986e-02	.	+	.	.	.
7.500e-02	4.431e-02	.	+	.	.	.
time	vamm1#br	0.00e+00	1.00e+00	2.00e+00	3.00e+00	4.00e+00

Verify any of the simulated current values using your own manual calculations.

Challenges

- Explain how you may use a SPICE simulation such as this for your own practice, instead of relying on questions and answers provided to you in a textbook.
- Explain how you could use this simulation to practice calculating inductor or resistor values from the simulated results. In other words, imagine this was real data from a real inductor-resistor circuit experiment, and you needed to compute either the resistor value or the inductor value from all the other data.

7.3 Diagnostic reasoning

These questions are designed to stimulate your deductive and inductive thinking, where you must apply general principles to specific scenarios (deductive) and also derive conclusions about the failed circuit from specific details (inductive). In a Socratic discussion with your instructor, the goal is for these questions to reinforce your recall and use of general circuit principles and also challenge your ability to integrate multiple symptoms into a sensible explanation of what's wrong in a circuit. Your instructor may also pose additional questions based on those assigned, in order to further challenge and sharpen your diagnostic abilities.

As always, your goal is to fully *explain* your analysis of each problem. Simply obtaining a correct answer is not good enough – you must also demonstrate sound reasoning in order to successfully complete the assignment. Your instructor's responsibility is to probe and challenge your understanding of the relevant principles and analytical processes in order to ensure you have a strong foundation upon which to build further understanding.

You will note a conspicuous lack of answers given for these diagnostic questions. Unlike standard textbooks where answers to every other question are given somewhere toward the back of the book, here in these learning modules students must rely on other means to check their work. The best way by far is to debate the answers with fellow students and also with the instructor during the Socratic dialogue sessions intended to be used with these learning modules. Reasoning through challenging questions with other people is an excellent tool for developing strong reasoning skills.

Another means of checking your diagnostic answers, where applicable, is to use circuit simulation software to explore the effects of faults placed in circuits. For example, if one of these diagnostic questions requires that you predict the effect of an open or a short in a circuit, you may check the validity of your work by simulating that same fault (substituting a very high resistance in place of that component for an open, and substituting a very low resistance for a short) within software and seeing if the results agree.

7.3.1 Insulation breakdown


One mode of failure for an inductor is breakdown (degradation) of the insulation covering the wire used in its windings. Identify how such a failure might affect the operation of an inductor, and how it might be detected through appropriate testing.

Challenges

- Do you suppose this sort of inductor failure might be caused by excessive voltage or excessive current?

7.3.2 Incorrect voltage/current curves

Explain what is wrong with these voltage and current plots for an inductor energized by a constant-voltage source through a series resistance:



Furthermore, try to explain *why* someone might make this mistake when predicting voltage and current.

Challenges

- Devise a “thought experiment” whereby this error would become obviously absurd (i.e. *reductio ad absurdum*).

7.3.3 Inductor proximity


A practical consideration affecting inductors (but not capacitors) is physical proximity: how close multiple inductors are located near each other in a circuit. Explain why inductors might affect each other (and how!), whereas capacitors do not.

Challenges

- One interesting technique for combating this phenomenon is to orient nearby inductors *perpendicular* to each other rather than parallel or in-line with each other. Explain why this strategy works.

7.3.4 Parasitic inductance on a PCB

Components soldered into printed circuit boards often possess “stray” inductance, also known as *parasitic* inductance. Observe this resistor, soldered in to a circuit board:



Where does the parasitic inductance come from? What is it about a resistor, mounted to a circuit board, that creates a (very) small amount of inductance? How is it possible to minimize this inductance, in case it is detrimental to the circuit’s operation?

Challenges

- Fast-switching digital logic circuits, even though operating on very low levels of current, nevertheless may experience extremely high *rates of change* due to the speed (frequency) of the signal oscillations. It is not uncommon in modern digital systems to witness rates as high as 500 A/ns (500 Amperes per nanosecond!). How much parasitic inductance would be necessary to drop 1 Volt at this high rate of change?

7.3.5 Inductor saturation

What happens to the inductance of an inductor as its core becomes *saturated*? Does the inductance value increase, decrease, or remain the same? Explain your answer.

Challenges

- What condition(s) might lead to inductor saturation?
- Are there any core materials immune from saturation?

Appendix A

Problem-Solving Strategies

The ability to solve complex problems is arguably one of the most valuable skills one can possess, and this skill is particularly important in any science-based discipline.

- Study principles, not procedures. Don't be satisfied with merely knowing how to compute solutions – learn *why* those solutions work.
- Identify what it is you need to solve, identify all relevant data, identify all units of measurement, identify any general principles or formulae linking the given information to the solution, and then identify any “missing pieces” to a solution. Annotate all diagrams with this data.
- Sketch a diagram to help visualize the problem. When building a real system, always devise a plan for that system and analyze its function *before* constructing it.
- Follow the units of measurement and meaning of every calculation. If you are ever performing mathematical calculations as part of a problem-solving procedure, and you find yourself unable to apply each and every intermediate result to some aspect of the problem, it means you don't understand what you are doing. Properly done, every mathematical result should have practical meaning for the problem, and not just be an abstract number. You should be able to identify the proper units of measurement for each and every calculated result, and show where that result fits into the problem.
- Perform “thought experiments” to explore the effects of different conditions for theoretical problems. When troubleshooting real systems, perform *diagnostic tests* rather than visually inspecting for faults, the best diagnostic test being the one giving you the most information about the nature and/or location of the fault with the fewest steps.
- Simplify the problem until the solution becomes obvious, and then use that obvious case as a model to follow in solving the more complex version of the problem.
- Check for exceptions to see if your solution is incorrect or incomplete. A good solution will work for *all* known conditions and criteria. A good example of this is the process of testing scientific hypotheses: the task of a scientist is not to find support for a new idea, but rather to *challenge* that new idea to see if it holds up under a battery of tests. The philosophical

principle of *reductio ad absurdum* (i.e. disproving a general idea by finding a specific case where it fails) is useful here.

- Work “backward” from a hypothetical solution to a new set of given conditions.
- Add quantities to problems that are qualitative in nature, because sometimes a little math helps illuminate the scenario.
- Sketch graphs illustrating how variables relate to each other. These may be quantitative (i.e. with realistic number values) or qualitative (i.e. simply showing increases and decreases).
- Treat quantitative problems as qualitative in order to discern the relative magnitudes and/or directions of change of the relevant variables. For example, try determining what happens if a certain variable were to increase or decrease before attempting to precisely calculate quantities: how will each of the dependent variables respond, by increasing, decreasing, or remaining the same as before?
- Consider limiting cases. This works especially well for qualitative problems where you need to determine which direction a variable will change. Take the given condition and magnify that condition to an extreme degree as a way of simplifying the direction of the system’s response.
- Check your work. This means regularly testing your conclusions to see if they make sense. This does *not* mean repeating the same steps originally used to obtain the conclusion(s), but rather to use some other means to check validity. Simply repeating procedures often leads to *repeating the same errors* if any were made, which is why alternative paths are better.

Appendix B

Instructional philosophy

“The unexamined circuit is not worth energizing” – Socrates (if he had taught electricity)

These learning modules, although useful for self-study, were designed to be used in a formal learning environment where a subject-matter expert challenges students to digest the content and exercise their critical thinking abilities in the answering of questions and in the construction and testing of working circuits.

The following principles inform the instructional and assessment philosophies embodied in these learning modules:

- The first goal of education is to enhance clear and independent thought, in order that every student reach their fullest potential in a highly complex and inter-dependent world. Robust reasoning is *always* more important than particulars of any subject matter, because its application is universal.
- Literacy is fundamental to independent learning and thought because text continues to be the most efficient way to communicate complex ideas over space and time. Those who cannot read with ease are limited in their ability to acquire knowledge and perspective.
- Articulate communication is fundamental to work that is complex and interdisciplinary.
- Faulty assumptions and poor reasoning are best corrected through challenge, not presentation. The rhetorical technique of *reductio ad absurdum* (disproving an assertion by exposing an absurdity) works well to discipline student’s minds, not only to correct the problem at hand but also to learn how to detect and correct future errors.
- Important principles should be repeatedly explored and widely applied throughout a course of study, not only to reinforce their importance and help ensure their mastery, but also to showcase the interconnectedness and utility of knowledge.

These learning modules were expressly designed to be used in an “inverted” teaching environment¹ where students first read the introductory and tutorial chapters on their own, then individually attempt to answer the questions and construct working circuits according to the experiment and project guidelines. The instructor never lectures, but instead meets regularly with each individual student to review their progress, answer questions, identify misconceptions, and challenge the student to new depths of understanding through further questioning. Regular meetings between instructor and student should resemble a Socratic² dialogue, where questions serve as scalpels to dissect topics and expose assumptions. The student passes each module only after consistently demonstrating their ability to logically analyze and correctly apply all major concepts in each question or project/experiment. The instructor must be vigilant in probing each student’s understanding to ensure they are truly *reasoning* and not just *memorizing*. This is why “Challenge” points appear throughout, as prompts for students to think deeper about topics and as starting points for instructor queries. Sometimes these challenge points require additional knowledge that hasn’t been covered in the series to answer in full. This is okay, as the major purpose of the Challenges is to stimulate analysis and synthesis on the part of each student.

The instructor must possess enough mastery of the subject matter and awareness of students’ reasoning to generate their own follow-up questions to practically any student response. Even completely correct answers given by the student should be challenged by the instructor for the purpose of having students practice articulating their thoughts and defending their reasoning. Conceptual errors committed by the student should be exposed and corrected not by direct instruction, but rather by reducing the errors to an absurdity³ through well-chosen questions and thought experiments posed by the instructor. Becoming proficient at this style of instruction requires time and dedication, but the positive effects on critical thinking for both student and instructor are spectacular.

An inspection of these learning modules reveals certain unique characteristics. One of these is a bias toward thorough explanations in the tutorial chapters. Without a live instructor to explain concepts and applications to students, the text itself must fulfill this role. This philosophy results in lengthier explanations than what you might typically find in a textbook, each step of the reasoning process fully explained, including footnotes addressing common questions and concerns students raise while learning these concepts. Each tutorial seeks to not only explain each major concept in sufficient detail, but also to explain the logic of each concept and how each may be developed

¹In a traditional teaching environment, students first encounter new information via *lecture* from an expert, and then independently apply that information via *homework*. In an “inverted” course of study, students first encounter new information via *homework*, and then independently apply that information under the scrutiny of an expert. The expert’s role in lecture is to simply *explain*, but the expert’s role in an inverted session is to *challenge, critique*, and if necessary *explain* where gaps in understanding still exist.

²Socrates is a figure in ancient Greek philosophy famous for his unflinching style of questioning. Although he authored no texts, he appears as a character in Plato’s many writings. The essence of Socratic philosophy is to leave no question unexamined and no point of view unchallenged. While purists may argue a topic such as electric circuits is too narrow for a true Socratic-style dialogue, I would argue that the essential thought processes involved with scientific reasoning on *any* topic are not far removed from the Socratic ideal, and that students of electricity and electronics would do very well to challenge assumptions, pose thought experiments, identify fallacies, and otherwise employ the arsenal of critical thinking skills modeled by Socrates.

³This rhetorical technique is known by the Latin phrase *reductio ad absurdum*. The concept is to expose errors by counter-example, since only one solid counter-example is necessary to disprove a universal claim. As an example of this, consider the common misconception among beginning students of electricity that voltage cannot exist without current. One way to apply *reductio ad absurdum* to this statement is to ask how much current passes through a fully-charged battery connected to nothing (i.e. a clear example of voltage existing without current).

from “first principles”. Again, this reflects the goal of developing clear and independent thought in students’ minds, by showing how clear and logical thought was used to forge each concept. Students benefit from witnessing a model of clear thinking in action, and these tutorials strive to be just that.

Another characteristic of these learning modules is a lack of step-by-step instructions in the Project and Experiment chapters. Unlike many modern workbooks and laboratory guides where step-by-step instructions are prescribed for each experiment, these modules take the approach that students must learn to closely read the tutorials and apply their own reasoning to identify the appropriate experimental steps. Sometimes these steps are plainly declared in the text, just not as a set of enumerated points. At other times certain steps are implied, an example being assumed competence in test equipment use where the student should not need to be told *again* how to use their multimeter because that was thoroughly explained in previous lessons. In some circumstances no steps are given at all, leaving the entire procedure up to the student.

This lack of prescription is not a flaw, but rather a feature. Close reading and clear thinking are foundational principles of this learning series, and in keeping with this philosophy all activities are designed to *require* those behaviors. Some students may find the lack of prescription frustrating, because it demands more from them than what their previous educational experiences required. This frustration should be interpreted as an unfamiliarity with autonomous thinking, a problem which must be corrected if the student is ever to become a self-directed learner and effective problem-solver. Ultimately, the need for students to read closely and think clearly is more important both in the near-term and far-term than any specific facet of the subject matter at hand. If a student takes longer than expected to complete a module because they are forced to outline, digest, and reason on their own, so be it. The future gains enjoyed by developing this mental discipline will be well worth the additional effort and delay.

Another feature of these learning modules is that they do not treat topics in isolation. Rather, important concepts are introduced early in the series, and appear repeatedly as stepping-stones toward other concepts in subsequent modules. This helps to avoid the “compartmentalization” of knowledge, demonstrating the inter-connectedness of concepts and simultaneously reinforcing them. Each module is fairly complete in itself, reserving the beginning of its tutorial to a review of foundational concepts.

This methodology of assigning text-based modules to students for digestion and then using Socratic dialogue to assess progress and hone students’ thinking was developed over a period of several years by the author with his Electronics and Instrumentation students at the two-year college level. While decidedly unconventional and sometimes even unsettling for students accustomed to a more passive lecture environment, this instructional philosophy has proven its ability to convey conceptual mastery, foster careful analysis, and enhance employability so much better than lecture that the author refuses to ever teach by lecture again.

Problems which often go undiagnosed in a lecture environment are laid bare in this “inverted” format where students must articulate and logically defend their reasoning. This, too, may be unsettling for students accustomed to lecture sessions where the instructor cannot tell for sure who comprehends and who does not, and this vulnerability necessitates sensitivity on the part of the “inverted” session instructor in order that students never feel discouraged by having their errors exposed. *Everyone* makes mistakes from time to time, and learning is a lifelong process! Part of the instructor’s job is to build a culture of learning among the students where errors are not seen as shameful, but rather as opportunities for progress.

To this end, instructors managing courses based on these modules should adhere to the following principles:

- Student questions are always welcome and demand thorough, honest answers. The only type of question an instructor should refuse to answer is one the student should be able to easily answer on their own. Remember, *the fundamental goal of education is for each student to learn to think clearly and independently*. This requires hard work on the part of the student, which no instructor should ever circumvent. Anything done to bypass the student's responsibility to do that hard work ultimately limits that student's potential and thereby does real harm.
- It is not only permissible, but encouraged, to answer a student's question by asking questions in return, these follow-up questions designed to guide the student to reach a correct answer through their own reasoning.
- All student answers demand to be challenged by the instructor and/or by other students. This includes both correct and incorrect answers – the goal is to practice the articulation and defense of one's own reasoning.
- No reading assignment is deemed complete unless and until the student demonstrates their ability to accurately summarize the major points in their own terms. Recitation of the original text is unacceptable. This is why every module contains an "Outline and reflections" question as well as a "Foundational concepts" question in the Conceptual reasoning section, to prompt reflective reading.
- No assigned question is deemed answered unless and until the student demonstrates their ability to consistently and correctly apply the concepts to *variations* of that question. This is why module questions typically contain multiple "Challenges" suggesting different applications of the concept(s) as well as variations on the same theme(s). Instructors are encouraged to devise as many of their own "Challenges" as they are able, in order to have a multitude of ways ready to probe students' understanding.
- No assigned experiment or project is deemed complete unless and until the student demonstrates the task in action. If this cannot be done "live" before the instructor, video-recordings showing the demonstration are acceptable. All relevant safety precautions must be followed, all test equipment must be used correctly, and the student must be able to properly explain all results. The student must also successfully answer all Challenges presented by the instructor for that experiment or project.

Students learning from these modules would do well to abide by the following principles:

- No text should be considered fully and adequately read unless and until you can express every idea *in your own words, using your own examples*.
- You should always articulate your thoughts as you read the text, noting points of agreement, confusion, and epiphanies. Feel free to print the text on paper and then write your notes in the margins. Alternatively, keep a journal for your own reflections as you read. This is truly a helpful tool when digesting complicated concepts.
- Never take the easy path of highlighting or underlining important text. Instead, *summarize* and/or *comment* on the text using your own words. This actively engages your mind, allowing you to more clearly perceive points of confusion or misunderstanding on your own.
- A very helpful strategy when learning new concepts is to place yourself in the role of a teacher, if only as a mental exercise. Either explain what you have recently learned to someone else, or at least *imagine* yourself explaining what you have learned to someone else. The simple act of having to articulate new knowledge and skill forces you to take on a different perspective, and will help reveal weaknesses in your understanding.
- Perform each and every mathematical calculation and thought experiment shown in the text on your own, referring back to the text to see that your results agree. This may seem trivial and unnecessary, but it is critically important to ensuring you actually understand what is presented, especially when the concepts at hand are complicated and easy to misunderstand. Apply this same strategy to become proficient in the use of *circuit simulation software*, checking to see if your simulated results agree with the results shown in the text.
- Above all, recognize that learning is hard work, and that a certain level of frustration is unavoidable. There are times when you will struggle to grasp some of these concepts, and that struggle is a natural thing. Take heart that it will yield with persistent and varied⁴ effort, and never give up!

Students interested in using these modules for self-study will also find them beneficial, although the onus of responsibility for thoroughly reading and answering questions will of course lie with that individual alone. If a qualified instructor is not available to challenge students, a workable alternative is for students to form study groups where they challenge⁵ one another.

To high standards of education,

Tony R. Kuphaldt

⁴As the old saying goes, “Insanity is trying the same thing over and over again, expecting different results.” If you find yourself stumped by something in the text, you should attempt a different approach. Alter the thought experiment, change the mathematical parameters, do whatever you can to see the problem in a slightly different light, and then the solution will often present itself more readily.

⁵Avoid the temptation to simply share answers with study partners, as this is really counter-productive to learning. Always bear in mind that the answer to any question is far less important in the long run than the method(s) used to obtain that answer. The goal of education is to empower one’s life through the improvement of clear and independent thought, literacy, expression, and various practical skills.

Appendix C

Tools used

I am indebted to the developers of many open-source software applications in the creation of these learning modules. The following is a list of these applications with some commentary on each.

You will notice a theme common to many of these applications: a bias toward *code*. Although I am by no means an expert programmer in any computer language, I understand and appreciate the flexibility offered by code-based applications where the user (you) enters commands into a plain ASCII text file, which the software then reads and processes to create the final output. Code-based computer applications are by their very nature *extensible*, while WYSIWYG (What You See Is What You Get) applications are generally limited to whatever user interface the developer makes for you.

The GNU/Linux computer operating system

There is so much to be said about Linus Torvalds' **Linux** and Richard Stallman's **GNU** project. First, to credit just these two individuals is to fail to do justice to the *mob* of passionate volunteers who contributed to make this amazing software a reality. I first learned of **Linux** back in 1996, and have been using this operating system on my personal computers almost exclusively since then. It is *free*, it is completely *configurable*, and it permits the continued use of highly efficient **Unix** applications and scripting languages (e.g. shell scripts, Makefiles, **sed**, **awk**) developed over many decades. **Linux** not only provided me with a powerful computing platform, but its open design served to inspire my life's work of creating open-source educational resources.

Bram Moolenaar's Vim text editor

Writing code for any code-based computer application requires a *text editor*, which may be thought of as a word processor strictly limited to outputting plain-ASCII text files. Many good text editors exist, and one's choice of text editor seems to be a deeply personal matter within the programming world. I prefer **Vim** because it operates very similarly to **vi** which is ubiquitous on **Unix/Linux** operating systems, and because it may be entirely operated via keyboard (i.e. no mouse required) which makes it fast to use.

Donald Knuth's \TeX typesetting system

Developed in the late 1970's and early 1980's by computer scientist extraordinaire Donald Knuth to typeset his multi-volume magnum opus *The Art of Computer Programming*, this software allows the production of formatted text for screen-viewing or paper printing, all by writing plain-text code to describe how the formatted text is supposed to appear. \TeX is not just a markup language for documents, but it is also a Turing-complete programming language in and of itself, allowing useful algorithms to be created to control the production of documents. Simply put, *\TeX is a programmer's approach to word processing*. Since \TeX is controlled by code written in a plain-text file, this means anyone may read that plain-text file to see exactly how the document was created. This openness afforded by the code-based nature of \TeX makes it relatively easy to learn how other people have created their own \TeX documents. By contrast, examining a beautiful document created in a conventional WYSIWYG word processor such as Microsoft **Word** suggests nothing to the reader about *how* that document was created, or what the user might do to create something similar. As Mr. Knuth himself once quipped, conventional word processing applications should be called WYSIAYG (What You See Is *All* You Get).

Leslie Lamport's \LaTeX extensions to \TeX

Like all true programming languages, \TeX is inherently extensible. So, years after the release of \TeX to the public, Leslie Lamport decided to create a massive extension allowing easier compilation of book-length documents. The result was \LaTeX , which is the markup language used to create all ModEL module documents. You could say that \TeX is to \LaTeX as **C** is to **C++**. This means it is permissible to use any and all \TeX commands within \LaTeX source code, and it all still works. Some of the features offered by \LaTeX that would be challenging to implement in \TeX include automatic index and table-of-content creation.

Tim Edwards' **Xcircuit** drafting program

This wonderful program is what I use to create all the schematic diagrams and illustrations (but not photographic images or mathematical plots) throughout the ModEL project. It natively outputs PostScript format which is a true vector graphic format (this is why the images do not pixellate when you zoom in for a closer view), and it is so simple to use that I have never had to read the manual! Object libraries are easy to create for **Xcircuit**, being plain-text files using PostScript programming conventions. Over the years I have collected a large set of object libraries useful for drawing electrical and electronic schematics, pictorial diagrams, and other technical illustrations.

Gimp graphic image manipulation program

Essentially an open-source clone of Adobe's `PhotoShop`, I use `Gimp` to resize, crop, and convert file formats for all of the photographic images appearing in the `MODEL` modules. Although `Gimp` does offer its own scripting language (called `Script-Fu`), I have never had occasion to use it. Thus, my utilization of `Gimp` to merely crop, resize, and convert graphic images is akin to using a sword to slice bread.

SPICE circuit simulation program

`SPICE` is to circuit analysis as `TEX` is to document creation: it is a form of markup language designed to describe a certain object to be processed in plain-ASCII text. When the plain-text "source file" is compiled by the software, it outputs the final result. More modern circuit analysis tools certainly exist, but I prefer `SPICE` for the following reasons: it is *free*, it is *fast*, it is *reliable*, and it is a fantastic tool for *teaching* students of electricity and electronics how to write simple code. I happen to use rather old versions of `SPICE`, version 2g6 being my "go to" application when I only require text-based output. `NGSPICE` (version 26), which is based on Berkeley `SPICE` version 3f5, is used when I require graphical output for such things as time-domain waveforms and Bode plots. In all `SPICE` example netlists I strive to use coding conventions compatible with all `SPICE` versions.

Andrew D. Hwang's ePiX mathematical visualization programming library

This amazing project is a `C++` library you may link to any `C/C++` code for the purpose of generating PostScript graphic images of mathematical functions. As a completely free and open-source project, it does all the plotting I would otherwise use a Computer Algebra System (CAS) such as `Mathematica` or `Maple` to do. It should be said that `ePiX` is *not* a Computer Algebra System like `Mathematica` or `Maple`, but merely a mathematical *visualization* tool. In other words, it won't determine integrals for you (you'll have to implement that in your own `C/C++` code!), but it can graph the results, and it does so beautifully. What I really admire about `ePiX` is that it is a `C++` programming library, which means it builds on the existing power and toolset available with that programming language. Mr. Hwang could have probably developed his own stand-alone application for mathematical plotting, but by creating a `C++` library to do the same thing he accomplished something much greater.

gnuplot mathematical visualization software

Another open-source tool for mathematical visualization is `gnuplot`. Interestingly, this tool is *not* part of Richard Stallman’s GNU project, its name being a coincidence. For this reason the authors prefer “gnu” *not* be capitalized at all to avoid confusion. This is a much “lighter-weight” alternative to a spreadsheet for plotting tabular data, and the fact that it easily outputs directly to an X11 console or a file in a number of different graphical formats (including PostScript) is very helpful. I typically set my `gnuplot` output format to default (X11 on my Linux PC) for quick viewing while I’m developing a visualization, then switch to PostScript file export once the visual is ready to include in the document(s) I’m writing. As with my use of `Gimp` to do rudimentary image editing, my use of `gnuplot` only scratches the surface of its capabilities, but the important points are that it’s *free* and that it *works well*.

Python programming language

Both Python and C++ find extensive use in these modules as instructional aids and exercises, but I’m listing Python here as a *tool* for myself because I use it almost daily as a *calculator*. If you open a Python interpreter console and type `from math import *` you can type mathematical expressions and have it return results just as you would on a hand calculator. Complex-number (i.e. *phasor*) arithmetic is similarly supported if you include the complex-math library (`from cmath import *`). Examples of this are shown in the Programming References chapter (if included) in each module. Of course, being a fully-featured programming language, Python also supports conditionals, loops, and other structures useful for calculation of quantities. Also, running in a console environment where all entries and returned values show as text in a chronologically-ordered list makes it easy to copy-and-paste those calculations to document exactly how they were performed.

Appendix D

Creative Commons License

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.

Appendix E

References

“2000 Series High Temperature Powder Cores for Power Applications”, Issue C, Micrometals Iron Powder Cores, Anaheim, CA, February 2007.

“2022 Iron Powder Products Catalog”, Micrometals Iron Powder Cores, Anaheim, CA, March 2019.

Boylestad, Robert L., *Introductory Circuit Analysis*, 9th Edition, Prentice Hall, Upper Saddle River, NJ, 2000.

Bucher, Elmer E., *The Wireless Experimenter’s Manual*, Wireless Press Incorporated, New York, 1920.

“Inductor Color Codes”, Token Electronics Industry Company Limited, New Taipei City, Taiwan ROC, 16 January 2017.

“Inductors Instructional Guide”, document VMN-SG2139-1203, Vishay Intertechnology Incorporated, Shelton, CT.

Morgan, Alfred P., *Wireless Telegraphy and Telephony*, Fourth Edition, The Norman W. Henley Publishing Company, New York, April 1920.

“Part Number: T50-66C”, Revision 20190524, Micrometals Iron Powder Cores, Anaheim, CA, 30 May 2019.

Van Stephenson, Gary, *US Patent 10,957,473 B2*, “Dual Superconducting Magnetic Energy Storage”, application 2 November 2018, patent granted 23 March 2021.

“What Are The Different Types of Inductors?”, Agile Magnetics, Concord, NH.

Appendix F

Version history

This is a list showing all significant additions, corrections, and other edits made to this learning module. Each entry is referenced by calendar date in reverse chronological order (newest version first), which appears on the front cover of every learning module for easy reference. Any contributors to this open-source document are listed here as well.

26-27 March 2025 – added comments to the Introduction chapter urging instructors to have students show their in-class work solving problems. Also added more examples to the “Example: RC time-delay network values” Case Tutorial section.

20-21 March 2025 – added a new Case Tutorial section called “Example: LR time-delay network values” with multiple circuit examples for student practice. Also added a new Conceptual Reasoning question asking students to explain the conceptual meaning of all calculations done for an LR time-delay circuit.

5 March 2025 – added another photograph to the “Early examples of inductors” Historical Reference showing a spiral inductor.

3 March 2025 – corrected a grammatical error (“we will may”) courtesy of Caleb Wing.

10 January 2025 – added a new Case Tutorial section, this one on applying Norton’s Theorem to networks containing multiple resistors in order to simplify to a single-resistor single-inductor LR network.

14 November 2024 – added a new bullet-point to the list of challenging concepts in the Introduction chapter, on proper graphing of trajectories.

8 November 2024 – added a new Programming Reference section on solving differential equations using C++.

23-24 October 2024 – added a new Technical References section on applying Thévenin’s and Norton’s Theorems to time-delay circuits. Also added a hint to the “Inductor charging circuit” Conceptual Reasoning question.

18 October 2024 – completely re-wrote the “Inverse exponential growth/decay” section of the Tutorial to begin with a discussion of radioactive half-life.

17 October 2024 – edited the “Inductor charging circuit” Conceptual Reasoning question to emphasize why this is a good pre-analysis strategy for any quantitative analysis.

16 September 2024 – edited the Introduction chapter to have a separate section for challenging concepts.

21 July 2024 – created a new Historical References section on old inductor technology.

2 July 2024 – added an enumerated list of conceptual analysis steps to the “Example: current source energizing an inductor” Case Tutorial section, showing students exactly how and why to approach problems such as this.

1 July 2024 – minor edits to analogies used within the Tutorial. Also corrected mathematical errors in the limit equations shown in that chapter.

26-28 June 2024 – divided the Introduction chapter into two sections, one for students and one for instructors, and added content to the instructor section recommending learning outcomes and measures. Also made minor edits to the Tutorial’s discussion on inductance acting as source, load, or short. Also edited image_0455.

5 May 2024 – added a graphing requirement to the “Inductor charging circuit” Conceptual Reasoning question.

12-13 March 2024 – corrected a typographical error in a table of inverse-exponential calculated values, courtesy of Matthew Oleson. Corrected another error in a question where I specified a starting time of 0 ms rather than $0\mu\text{s}$.

10 March 2024 – added a list of challenging concepts to the Introduction chapter.

4 February 2024 – simplified explanations of how to determine what grows versus what decays based on an analysis of stored energy. Also added some clarifying details to some of the Case Tutorial examples.

14 November 2023 – added a new Case Tutorial section on empirically determining signal rates of change.

24 October 2023 – added Challenge question to the “Inductor charging circuit” Conceptual Reasoning question.

20 October 2023 – modified the “Inverse exponential growth/decay” section of the Tutorial chapter to include a table of pre-calculated values for e^x and $1 - e^x$ showing these two formulae as the exponential term x becomes increasingly negative.

25 July 2023 – added a footnote to the Tutorial discussing magnetic saturation in more detail.

7 February 2023 – added caveats to Case Tutorial examples about the inductors being ideal (with no winding resistance).

27 November 2022 – placed questions at the top of the itemized list in the Introduction chapter prompting students to devise experiments related to the tutorial content.

17-19 October 2022 – elaborated in the Tutorial on how to incorporate initial and/or final values to the inverse-exponential formulae. Also added a Programming References section showing an extremely simple Python3 script for modeling inverse-exponential functions. Also added a new Tutorial section showing a worked example of an LR charging/discharging circuit.

22 July 2022 – minor edits to Tutorial giving a more thorough definition of a “time constant”. Also expanded upon a question in the Introduction chapter, designed to prompt more reasoning based on foundational concepts.

19 July 2022 – significant change made to image_6312 which is in the Case Tutorial section on inductive energy storage.

13 July 2022 – moved text and images from the “interesting relay behavior” Case Tutorial section into its own file for inclusion in other modules.

13 June 2022 – added a new Case Tutorial section showing energy storage calculations for series versus parallel inductor networks.

23 April 2022 – added illustrations from a patent on a superconducting energy storage device showing DC currents in a superconducting inductor while charging, discharging, and in an idle state.

22 April 2022 – added photograph of a variable inductor.

31 March 2022 – added “Magnetic field quantities” to the Derivations and Technical References chapter.

16 March 2022 – added some content to the Tutorial discussing the calculation of inductance based on permeability, coil area, turns, and coil length. Also added some instructor notes.

28 December 2021 – added a Technical Reference section on inductor parasitics.

6-8 December 2021 – added Tutorial section on making your own inductors, and expanded the Quantitative Reasoning question “Inductor wire turns” to include a problem using a toroidal powdered-iron core with a specified A_L value. Also added Case Tutorial section on testing custom-wound inductors.

20-21 October 2021 – deleted the Projects and Experiments chapter. Also added some clarifying text in the Tutorial about how the Conservation of Energy relates to voltage and current for these devices.

26-27 July 2021 – added Case Tutorial section showing interesting current profile through the

inductance of an electromechanical relay.

23 July 2021 – added foundational concept review to the “Inductor charging circuit” Conceptual Reasoning question.

16 July 2021 – added footnote to the Tutorial explaining what permeability is.

9 July 2021 – replaced some TeX-style italicizing markup with LaTeX-style.

4 July 2021 – divided Tutorial chapter into sections and made some minor edits to the text for clarity.

2 June 2021 – added a Case Tutorial chapter with sections showing how to store and release energy from inductors.

8 May 2021 – commented out or deleted empty chapters.

18 April 2021 – changed all lower-case Greek letter “phi” symbols (ϕ) to upper-case (Φ).

13 April 2021 – added “energy” as a parameter for students to qualitatively assess in the “Inductor charging circuit” Conceptual Reasoning question.

6 April 2021 – added some commentary to Tutorial on units of measurement for time constant (τ).

5 April 2021 – added a new Technical Reference section on inductor parameters and core types.

18 March 2021 – corrected one instance of “volts” that should have been capitalized “Volts”.

9 February 2021 – added a Technical Reference section on inductor labeling.

26 October 2020 – added Quantitative questions, and also fixed error in “Measuring inductance by time delay” where the graphical image showed an entirely incorrect plot for resistor voltage drop.

20 October 2020 – added footnote to Tutorial about inductor current remaining steady in a zero-voltage condition only when there is no resistance in the circuit.

9 September 2020 – added more instructor notes in “Inductor wire turns” question.

4 September 2020 – incorporated the “Inverse exponential functions” question content into the Tutorial, explaining the formulae for decaying and growing quantities in full detail.

29 August 2020 – significantly edited the Introduction chapter to make it more suitable as a pre-study guide and to provide cues useful to instructors leading “inverted” teaching sessions. Also, elaborated a bit more in the Tutorial on parasitic properties. Also, added some Case Tutorial examples showing circuits energized by voltage and current sources, respectively.

17 August 2020 – minor edit to the Tutorial text in order to shorten a paragraph, for typesetting purposes.

17 April 2020 – added a Technical Reference section on IEC standard component values (file `techref_IEC60063.latex`).

6 April 2020 – minor typographical error corrections, courtesy of Ron Felix.

7 March 2020 – added a Quantitative Reasoning problem based on an article from *The Electric Journal* circa 1921.

4 February 2020 – added some Quantitative Reasoning problems.

2 January 2020 – removed from from C++ code execution output, to clearly distinguish it from the source code listing which is still framed.

1 January 2020 – added explanatory text for the C++ code in `reference_invexp_decay.latex` and changed `main ()` to `main (void)` in another C++ programming example.

10 December 2019 – added more questions. Also added units of measurement to the explanation of the inductance formula.

2 November 2019 – added more questions.

26 August 2019 – added mention of difficult concepts to the Introduction.

14 April 2019 – added comment about “bifilar” winding in one of the Conceptual questions.

13 April 2019 – improved the explanation of what happens in LC circuits with inverse-exponential trends of voltage and current. Also introduced the term “speed” to refer to rate-of-change for inductor current (i.e. current “speed” = $\frac{dI}{dt}$). Also corrected some copy-and-paste errors from the `mod_capacitor` module where the text read “capacitor” when it should have read “inductor”.

10 April 2019 – corrected an error identified by a student of mine (Danny Starnes). In the “Measuring inductance by time delay” question, the sentence should read “connected in series” rather than “connected in parallel”.

8 April 2019 – minor edit to wording within an illustration (replaced “creates” with “means” to avoid the implication of a unidirectional causality).

24 March 2019 – added a Quantitative Reasoning questions.

13 March 2019 – added reference to inverse exponential functions in the Tutorial.

12 February 2019 – added a number of new entries to the Foundational Concepts subsection.

7 January 2019 – added illustrative examples to the Tutorial, showing what happens when an inductor is energized by a constant voltage. Also added a Derivations and Technical References chapter, with an exposition of exponential decay functions in it.

17 December 2018 – added an example C++ program showing calculations for an energizing

resistor-inductor circuit.

November 2018 – minor edits to Simplified Tutorial. Added experiment to measure drop-out time of electromechanical relay with varying voltage-suppressing load resistance.

October 2018 – clarified a number of details, with constructive criticism from Ron Felix.

August 2018 – wrote an Introduction where none existed before.

July 2018 – document first created.

Index

- τ , 109
- “Ohm’s Law” for an inductor, 47

- Adding quantities to a qualitative problem, 170
- Air-core inductor, 77
- Ampacity, 58
- Ampere, unit, 44
- Annotating diagrams, 169
- Armature, 28
- Asymptote, 109

- B-H curve, 77

- C++, 100, 110
- Capacitance, 34, 58
- Capacitor, 69
- Cathode-ray tube, 84
- CGS, 92
- Checking for exceptions, 170
- Checking your work, 170
- Choke, 44
- Code, computer, 177
- Coercivity, 78
- Color code, inductor, 74
- Compiler, C++, 100
- Complex number, 32
- Component values, IEC standard, 72
- Computer programming, 99, 110
- Condenser, 69
- Conservation of Energy, 31, 46
- Conservation of Mass, 46
- Conventional flow notation, 84
- Core, 44
- Core, magnetic, 64
- Coulomb, 83
- Cross product, 84
- Cross-product, 83

- CRT, 84
- Curie temperature, 76
- Current probe, 38

- Decay, exponential, 109
- Detector, 69
- Dimensional analysis, 85, 169
- Diode, 69
- Distance versus speed, 40

- Eddy current, 81
- Edwards, Tim, 178
- Electric field, 43, 109
- Electromagnetic induction, 77, 78
- Electromotive force, 91
- Electron flow notation, 84
- emf, 91
- Energy, kinetic, 43, 44
- Energy, potential, 44
- Equivalent Series Resistance, 80
- ESR, 80
- Exponential decay, 109

- Ferrite, 78
- Ferromagnetic, 3
- Ferromagnetic core, 64
- Field intensity, magnetic, 91
- Field, electric, 43, 109
- Field, magnetic, 43, 109
- Flux density, magnetic, 87

- Gauss, 87
- Graph values to solve a problem, 170
- Greenleaf, Cynthia, 127

- Hard magnetic material, 77
- Henry, 74
- Henry, unit, 3, 44

- How to teach with these modules, 172
- Hwang, Andrew D., 179
- Hysteresis, 57
- Hysteresis curve, 77
- Identify given data, 169
- Identify relevant principles, 169
- IEC 60063 standard, 72
- IEC standard component values, 72
- Impedance, 32
- Inductance, 3, 43, 64, 74
- Inductance factor, 61
- Inductance index, 61
- Induction, 77, 78
- Inductive reactance, 32
- Inductor, 3, 64
- Inductor color code, 74
- Instructions for projects and experiments, 173
- Intermediate results, 169
- Interpreter, Python, 104
- Inverse exponential curve, 51
- Inverse exponential decay, 109
- Inverted instruction, 172
- Java, 101, 111
- Joule's Law, 48, 77
- Joule, unit, 44
- Kirchhoff's Current Law, 16, 24, 51
- Kirchhoff's Voltage Law, 19, 23, 56
- Knuth, Donald, 178
- Laminated core, 78
- Lamport, Leslie, 178
- Left-hand rule, 84
- Limiting cases, 170
- Line reactor, 50
- Load, 46
- Lorentz force, 83
- Magnetic circuit, 91
- Magnetic field, 43, 109
- Magnetic field intensity, 91
- Magnetic flux density, 87
- Magnetic saturation, 57, 61, 77, 78
- Magnetically "hard" versus "soft" materials, 77
- Magnetomotive force, 91
- Magnetostriction, 77
- Maxwell, James Clerk, 63
- Metacognition, 132
- Metric system, CGS, 92
- Metric system, SI, 92
- mmf, 91
- Mole, 93
- Moolenaar, Bram, 177
- Murphy, Lynn, 127
- Mutual inductance, 59
- Newton, 83, 85
- Newton's Cooling Law, 110
- NOAA, 87
- Norton's Theorem, 25
- Ohm's Law, 16, 19, 22, 24, 51, 56
- Open-source, 177
- Parallel inductance, 59
- Parasitic, 58
- Parasitic effect, 38, 80
- Parasitic property, 34
- Particle accelerator, 84
- Periodic Table of the Elements, 93
- Permeability, 3, 44, 57, 87
- Pitch, 38
- Problem-solving: annotate diagrams, 169
- Problem-solving: check for exceptions, 170
- Problem-solving: checking work, 170
- Problem-solving: dimensional analysis, 169
- Problem-solving: graph values, 170
- Problem-solving: identify given data, 169
- Problem-solving: identify relevant principles, 169
- Problem-solving: interpret intermediate results, 169
- Problem-solving: limiting cases, 170
- Problem-solving: qualitative to quantitative, 170
- Problem-solving: quantitative to qualitative, 170
- Problem-solving: reductio ad absurdum, 170
- Problem-solving: simplify the system, 169
- Problem-solving: thought experiment, 59, 89, 90, 169
- Problem-solving: track units of measurement, 169
- Problem-solving: visually represent the system, 169

- Problem-solving: work in reverse, 170
- Programming, computer, 99, 110
- Python, 104

- Qualitatively approaching a quantitative problem, 170

- Reactance, inductive, 32
- Reactor, 44
- Reading Apprenticeship, 127
- Reductio ad absurdum, 170–172
- Relay, 28
- Remanence, 78
- Right-hand rule, 84
- Rise over run, 38

- Saturation, magnetic, 57, 61, 77, 78
- Schoenbach, Ruth, 127
- Scientific method, 132
- Series inductance, 59
- Short, 46
- SI, 92
- Simplifying a system, 169
- Skin effect, 32, 77
- Slope, 38
- SMES, 49
- Smith chart, 33, 35
- Socrates, 171
- Socratic dialogue, 172
- Soft magnetic material, 77
- Solenoid, 60, 89
- Source, 46
- Source code, 100
- SPDT switch, 51
- Speed versus distance, 40
- SPICE, 127
- SPST, 55
- Stallman, Richard, 177
- Standard component values, IEC, 72
- Superconducting Magnetic Energy Storage system, 49
- Superconductor, 47, 49, 88

- Tangent line, 38
- Tank circuit, 80
- Tesla, 83, 85, 87, 89, 90

- Thought experiment, 59, 89, 90, 169
- Time constant, 109
- Toroid, 90
- Torvalds, Linus, 177

- Units of measurement, 169

- Variometer, 68
- Vector, 83
- Vector cross-product, 84
- Visualizing a system, 169
- Voltage, 91

- Weber, 87
- Whitespace, C++, 100, 101
- Whitespace, Python, 107
- WMM, 87
- Work in reverse to solve a problem, 170
- World Magnetic Model, 87
- WYSIWYG, 177, 178