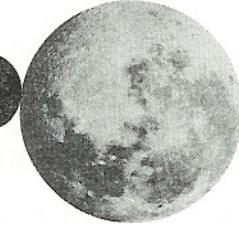


MASSACHUSETTS INSTITUTE OF TECHNOLOGY

APOLLO



GUIDANCE, NAVIGATION AND CONTROL

Approved: David G. Hoag Date: 8 Feb 72
D. G. HOAG, DIRECTOR
APOLLO GUIDANCE AND NAVIGATION PROGRAM

Approved: Ralph R. Ragan Date: 14 Feb 72
R. R. RAGAN, DEPUTY DIRECTOR
CHARLES STARK DRAPER LABORATORY

R-713

RELIABILITY HISTORY OF THE APOLLO GUIDANCE COMPUTER

by

Eldon C. Hall

JANUARY 1972

MIT

**CHARLES STARK DRAPER
LABORATORY**

CAMBRIDGE, MASSACHUSETTS, 02139

ACKNOWLEDGEMENT

This report was prepared under DSR Project 55-23890, sponsored by the Manned Spacecraft Center of the National Aeronautics and Space Administration through Contract NAS 9-4065.

The author would like to acknowledge the assistance of A.I. Green in the preparation of this report and many others, both within the Draper Labs and Raytheon, who have contributed to the collection and analysis of data.

The publication of this report does not constitute approval by the National Aeronautics and Space Administration of the findings or the conclusions contained therein. It is published only for the exchange and stimulation of ideas.

TABLE OF CONTENTS

	<u>Page</u>
1. INTRODUCTION.	2
2. DEVELOPMENT	3
2.1 COMPUTER DESIGN	3
2.2 DISPLAY AND KEYBOARD DESIGN	5
2.3 FINAL DESIGN	5
3. RELIABILITY APPROACHES.	7
3.1 FAULT DETECTION AND RESTART.	7
3.2 ELECTROMAGNETIC TOLERANCE	8
3.3 DESIGN PHILOSOPHY	8
3.4 COMPONENT DEVELOPMENT.	12
3.5 DESIGN QUALIFICATION AND PRODUCTION CONTROLS .	15
4. PROJECT EXPERIENCE	22
4.1 MANUFACTURING PROBLEMS	22
4.2 SYSTEM INTEGRATION EXPERIENCE	28
4.3 FIELD FAILURE HISTORY.	31
5. RELIABILITY STATISTICS.	39
6. SUMMARY AND CONCLUSIONS	41

RELIABILITY HISTORY OF THE APOLLO GUIDANCE COMPUTER

by
Eldon C. Hall

ABSTRACT

The APOLLO Guidance Computer was designed to provide the computation necessary for guidance, navigation and control of the Command Module and the Lunar Landing Module of the APOLLO spacecraft. The computer was designed using the technology of the early 1960's and the production was completed by 1969. During the development, production, and operational phase of the program, the computer has accumulated a very interesting history which is valuable for evaluating the technology, production methods, system integration, and the reliability of the hardware. The operational experience* in the APOLLO guidance systems includes 17 computers which flew missions and another 26 flight type computers which are still in various phases of prelaunch activity including storage, system checkout, prelaunch spacecraft checkout, etc.

These computers were manufactured and maintained under very strict quality control procedures with requirements for reporting and analyzing all indications of failure. Probably no other computer or electronic equipment with equivalent complexity has been as well documented and monitored. Since it has demonstrated a unique reliability history, it is important to evaluate the techniques and methods which have contributed to the high reliability of this computer.

*The operational experience includes missions through Apollo 15 which flew in August 1971. The compilation of all other data from this report ended 31 December 1970.

1. INTRODUCTION

The APOLLO guidance computer (AGC) is a real-time digital-control computer whose conception and development took place in the early part of 1960. The computer may be classified as a parallel, general-purpose or whole number binary computer. This class of computer is representative of most of the ground-based digital computers in existence in the late 1950s, when the precursors of the AGC were being designed. Few computers of this class had been designed by that time for the aerospace environment, and those few embodied substantial compromises in performance for the sake of conserving space, weight, and power.

The computer is the control and processing center of the APOLLO Guidance, Navigation and Control system. It processes data and issues discrete output and control pulses to the guidance system and other spacecraft systems. An operational APOLLO spacecraft contains two guidance computers and three DSKYs (keyboard and display unit for operator interface), with one computer and two DSKYs in the command module, and one of each in the lunar module. The computers are electrically identical, but differ in the use of computer software and interface control functions. As a control computer, some of the major functions are: alignment of the inertial measurement unit, processing of radar data, management of astronaut display and controls and generation of commands for spacecraft engine control. As a general purpose computer, the AGC solves the guidance and navigation equations required for the lunar mission.

2. DEVELOPMENT

The principal features of the electrical and mechanical design of the AGC were shaped by the nebulous constraints of the APOLLO program (unknown computational capacity, reliability, space, weight, and power) and the technology available to digital designers. The AGC evolved from these constraints and the development of mission requirements rather than from a fixed specification generated *a priori*. The desire for reliability beyond the state-of-the-art in digital computers was one of the most important driving forces which impacted the development and production of the computer. From this evolutionary process two designs resulted which were used operationally. The Block I computer was used on three unmanned spacecraft development flights, and the Block II was used on one unmanned Lunar Module flight and all manned flights. The major topics of interest are the Block II design and the techniques developed during the earlier phase which have impacted the computer design and reliability.

2.1 COMPUTER DESIGN

The first version of the Block I computer emerged in late 1962 with integrated circuit logic, wired-in (fixed) program memory, coincident-current erasable memory, and discrete-component circuits for the oscillator, power supplies, certain built-in test circuits, interfaces, and memory electronics. The final Block I computer was packaged using welded interconnections within modules which were interconnected with automatic wire-wrap.

This design had very limited capabilities due to the constraint on physical size and the desire for high reliability. The instruction repertoire, word length, and number of erasable memory cells were very limited. Provision was made, however, for a moderately large amount of fixed memory for

instructions and constants. A high density memory of the read-only type, called a rope memory, had been developed earlier to meet the goals of small physical size and high reliability and was carried over into the design of the APOLLO computer.

The rope memory, being a transformer type, depends for its information storage on the patterns with which its sensing wires are woven at the time of manufacture. Once a rope memory is built, its information content is fixed and is unalterable by electrical excitation. The high density and the information retention characteristics were the features that made it attractive for the AGC. Other technological developments which supported the AGC development were: 1. in semiconductor technology, where silicon transistors progressed to planar forms, then epitaxial form, and eventually to monolithic integrated circuits, 2. in coincident-current memories with low temperature coefficient lithium-ferrite cores for operation over a broad temperature range, 3. in packaging techniques, with the introduction of welded interconnection, multilayer printed circuit, and machine wirewrapping. These developments allowed significant reductions in volume and weight while coincidentally enhancing reliability. These packaging techniques were reduced to practice and had been used by MIT/DL in the development of the POLARIS guidance computer.

Integrated circuits were in development by the semiconductor industry during the late 1950s under Air Force sponsorship. In late 1961, MIT/DL evaluated a number of integrated circuits for the APOLLO guidance computer. An integrated circuit equivalent of the prototype APOLLO computer was constructed and tested in mid-1962 to discover any problems the circuits might exhibit when used in large numbers. Reliability, power consumption, noise generation, and noise susceptibility were the primary subjects of concern in the use of integrated circuits in the AGC. The performance of the units under evaluation was sufficient to justify their exclusive use for the logic section of the computer.

2.2 DISPLAY AND KEYBOARD DESIGN

As an adjunct to the APOLLO guidance computer, a display and keyboard unit was required as an information interface with the crew. The original design was made during the latter stages of development of the first version of the Block I computer, at which time neon numeric indicator tubes of the "Nixie" variety were used to generate three 4-digit displays for information, plus three 2-digit displays for identification. These were the minimum considered necessary, and they provided the capability of displaying three-vectors with sufficient precision for crew operations. The 2-digit indicators were used to display numeric codes for verbs, nouns, and program numbers. The verb-noun format permitted communication in language with syntax similar to that of spoken language. Examples of verbs were "display", "monitor", "load", and "proceed", and examples of nouns were "time", "gimbal angles", "error indications", and "star identification number." A keyboard was incorporated along with the display to allow the entering of numbers and codes for identifying them.

2.3 FINAL DESIGN

The Block II computer design (see Figure 1), resulting from the changes in technology and better definition of mission requirements since the Block I design, roughly doubled the speed, raised between 1.5 and 2 times the memory capacity, increased input/output capability, decreased size, and decreased power consumption. In addition the mechanical design included features which provided for moisture proofing and easy access to the six fixed memory modules. The design intent was to permit changing the memory inflight if the mission required more memory.

The final DSKY design incorporated three 5-digit registers and three 2-digit registers using segmented electroluminescent numeric displays, a 19-element keyboard with characters lighted with electroluminescent panels, and a

FIGURE 1

AGC CHARACTERISTICS

PERFORMANCE CHARACTERISTICS	BLOCK I	BLOCK II
Word Length	15 Bits + Parity	15 Bits + Parity
Number System	One's Complement	One's Complement
Fixed Memory Registers	24,576 Words	36,864 Words
Erasable Memory Registers	1,024 Words	2,048 Words
Number of Normal Instructions	11	34
Number of Involuntary Instructions (Interrupt, Increment, etc.)	8	10
Number of Interrupt Options	5	10
Number of Interface Counters	20	29
Number of Interface Circuits	143	227
Computer Clock Accuracy	0.3 ppm	0.3 ppm
Memory Cycle Time	11.7 μ sec	11.7 μ sec
Counter Increment Time	11.7 μ sec	11.7 μ sec
Addition Time	23.4 μ sec	23.4 μ sec
Multiplication Time	117 μ sec	46.8 μ sec
Divide Time	187.2 μ sec	70.2 μ sec
Double Precision Addition Time	1.65 millisec (subroutine)	35.1 μ sec
Number of Logic Gates	4,100	5,600
Volume	1.21 cubic ft.	0.97 cubic ft.
Weight	87 pounds	70 pounds
Power Consumption	85 watts	55 watts

14-legend caution and status display lighted with filamentary bulbs. The displays were switched under control of the computer using a matrix of 120 miniature relays some of which were latching in order to provide memory for the display elements.

3. RELIABILITY APPROACHES

Many approaches were taken to assure that the computer would realize the reliability requirements of the mission. The requirement for the AGC was a mission success probability of $(P_S)=0.998$. Early approaches which were studied included: 1. built-in test for fault detection, 2. in-flight repair, 3. dual computers with manual switchover, 4. a powered-down mode of operation called standby, 5. electrical and mechanical designs that left large margins above expected operating conditions, 6. an emphasis on reliability of components, testing procedures, and manufacturing. Of these approaches the concept of in-flight repair and dual computers was discarded after the configuration of the spacecraft was modified to provide for crew safety backups in the case of guidance failures. The mission success probability for the AGC remained the same however.

3.1 FAULT DETECTION AND RESTART

The computer's ability to detect faults using built-in test circuits was provided since it was known that digital equipment was very sensitive to transient disturbances and that a method of recovery from transient faults was very desirable. In the early designs these circuits and the self-checking software were necessary to accomplish the fault location required for in-flight repair. The circuits and the software were simplified for the final Block II AGC. Typical built-in tests include: a RUPT lock (too long in interrupt mode), TC trap (transfer of control to self address), parity fail (a parity bit is appended in every word in memory and is tested on all transfers to CPU), night watchman

alarm (a specified location has not been referenced often enough), and power fail (the voltage has dropped below a predetermined level). The circuits comprise two categories: those that are derived logically, and those that are derived using analog-type detection circuitry. The former circuitry is distributed within the logic modules of the computer and the latter in the alarm module.

The outputs of these fault detection circuits generate a computer restart, that is, transfer of control to a fixed program address. In addition, an indicator display is turned on. If the fault is transient in nature, the restart will succeed and the restart display can be cleared by depressing the reset (RSET) key. If the fault is a hard failure, the restart display will persist and a switch to a backup mode of operation is indicated.

3.2 ELECTROMAGNETIC TOLERANCE

In addition to the circuits to detect faults, considerable design effort and testing was expended in order to make the computer very tolerant to externally generated transient conditions and electromagnetic interference (EMI). For example, one test technique which was used to evaluate the shielding and grounding was the use of electrostatic discharges into the computer case and cabling of the system. After considerable testing and some significant changes in methods of grounding, the computer tolerated spark discharges to the case and cabling without failure. This desire for EMI tolerance had an impact on the cable shielding, the routing of wires within the computer, the interface circuit design, the power supply design, and the signal grounding internal to the computer.

3.3 DESIGN PHILOSOPHY

The electrical, mechanical, and thermal designs for the AGC followed a philosophy of overdesign, that is, one of providing capability in excess of identified requirements.

In the area of electrical design, the general philosophy was to make circuits as simple as possible, restrict the operating speed, minimize the component power consumption, and provide adequate operating margins when subjected to extremes of power supply voltages and thermal environments.

Standardization of circuit types was maximized at the expense of total component count. The use of several different types of circuit elements which would tend to reduce the total component count was avoided.

All components and circuits were designed with very comfortable operating margins. These included: first, computer operating speeds which were constrained to be well within the state-of-the-art of components and circuits; second, circuits which were designed for low power operation, not only for the purpose of conserving the total power, but also to keep the component power dissipation within very comfortable margins. The designers were constantly confronted with a conflict between operating speed, power consumption, and tolerance to voltage margins. Despite the requirement to minimize total power consumption, the resulting electrical design tolerated wide variations in power supply voltage.

In the area of mechanical design, the Block II computer utilizes modular construction and wire wrapping for the interconnections of the modules. The computer consists of two major subassemblies or trays (Trays A and B) containing modules and interconnecting wiring. The trays with the covers and gaskets provide mechanical support, thermal control via the spacecraft cold plate, environmental seal and shielding from electromagnetic interference. The rope modules are plugged into the structure from outside the sealed case. This permits program changes without breaking the environmental seal.

The module construction is basically welded cordwood type using standard components and integrated circuits. In the case of the 24 logic modules, the integrated circuit gates packaged in flatpacks are welded to multilayer boards

for interconnection between gates. The module frames provide mechanical support and thermal control for the components in addition to tray interface connector and jacking screws.

The modules are partitioned between the two trays such that the logic, interface, and power supply are in Tray A. The memory, memory electronics, analog alarm circuits, and oscillator are in Tray B, in addition to the connectors and mechanical support for the tray mounting the six rope modules.

The interconnecting wiring in the trays is accomplished by machine controlled wire wrapping for all interconnections. This technique provides a well controlled and easily reproduced method for making the large numbers of interconnections required. In the computer there are about 15,000 connector pins with an average of more than two connections per pin. After the wiring is complete, the tray is potted to provide mechanical support for the interconnecting wires and connector pins.

In the area of thermal design, the temperature control of the computer was achieved through conduction to the cold plate structure of the spacecraft. Radiational cooling was minimized by the choice of finishes to meet the requirements of spacecraft thermal control. Under some conditions, the surfaces surrounding the computer were at a higher temperature than the computer, thus causing additional heat loads instead of providing radiational cooling. In every case however, analysis indicated the effects of thermal radiation could be ignored in the thermal design of the computer.

Since the total power consumption of the computer is relatively low, the thermal control was mainly one of distributing the heat load in the computer and providing conduction paths to the cold plate. Module locations in the two trays (A and B) were carefully selected. The two power supplies were located at one wall in Tray A, where a short path and extra metal could be provided for the heat conduction to the cold plate. The E-memory, memory drivers, and

sense amplifiers are located in the center of Tray B to provide temperature tracking of the temperature compensating circuits and the memory cores. Conduction paths were provided from the electrical components to the base of the modules and then into the wirewrap plate, where the heat fans out to the sides of the trays, and thus down the walls of the Tray A cover to interface with the cold plate in the CM and with cold rails in the LM. In the case of the two switching transistors (NPN and PNP), thermal design included specifying a special package. The package was the standard TO-18 case size but with a solid metal header for decreased junction-to-case temperature rise. At the time of the Block II mechanical design, the solid metal header was not available in the TO-18 case size but had been used by semiconductor manufacturers on other similar cases. Thus the thermal design provided conduction from the element dissipating heat, such as the transistor chip, through all the mechanical interfaces to the cold plate.

The goals of the thermal design effort were: first, to ensure that the temperature of components and especially semiconductors remained below 100°C under worst-case conditions. The second goal was to provide a reasonably uniform thermal environment between modules like the memory electronics and logic modules. A temperature gradient between logic modules would reduce the operating margins of the logic. Thermal measurements on the finished computer have verified that these goals were met. The measured temperature difference between logic modules was less than 5°C and therefore negligible. The temperature rise through the structure to the hottest components was low enough to maintain junction temperatures well below 100°C.

Basic to the success of the APOLLO guidance computer was the realization that conventional reliability practices were not sufficient to meet the reliability requirement for the computer. An early estimate using fairly optimistic component failure rates and component counts, showed the resulting computer failure rate to be well above that which would be required to meet the computers

apportionment of the mission success probability ($P_S = 0.998$). Under these conditions designers could use redundancy techniques or develop more reliable components and manufacturing procedures in order to improve the reliability. In the case of the APOLLO computer various methods of accomplishing the redundancy were studied. However none could be used and still meet the power, size and weight requirements of the APOLLO mission. The elimination of redundancy provided the motivation for improving reliability at all levels of design, specification, manufacturing and testing. The tight assembly, inspection and test procedures during the manufacturing process detected many problems, each of which was closely monitored, and for which corrective actions were developed. The resulting emphasis on quality has paid off by decreasing the actual failure rates of the computer considerably below the original estimates, even though the component count increased after the original reliability estimates were made.

3.4 COMPONENT DEVELOPMENT

During the early stages of the computer design, an effort was made to constrain the number of different components to a selected few, thereby concentrating the engineering effort required in the area of component development. These constraints were rigidly adhered to and were a constant source of complaints from the circuit design engineers because they felt the limited number of component types constricted their designs excessively. Not only the types of parts were limited but also the range of values. For example, resistors were limited to one type and to a tightly restricted number of different values. The constraints were reviewed frequently and relaxed as new requirements were justified, but the existence of the constraints accomplished a greater than normal degree of standardization. The benefits that resulted from the effort to standardize were: (1) a reduction in the level of activity needed to specify the components and the level needed to develop testing methods that were capable of continuously monitoring the quality of the components, (2) a reduction in the efforts required to track the manufacturing problems that

were related to a component defect or testing procedure, and (3) more important to the reliability of the component was the large volume of procurements that provided increased competition between vendors and greater motivation to meet the reliability requirements.

Component selection was started in parallel with the development of circuit designs. Initially the design engineers were required to specify the general characteristics of the required components and the possible vendors for the component. Then, after a vendor was selected, sample purchases and engineering tests were made. One of the earliest and most important reliability tests was an internal visual examination of the component in order to identify the construction processes used. This visual examination identified weaknesses in the design, helped determine the type of tests that could be used to qualify the part, and provided information necessary to establish process controls. Additional engineering tests, both environmental and electrical, provided the information as feedback to the vendor for product improvement. This process of iteration varied in magnitude for different types of components. Parts like resistors and some condensors required little or no development activity, as only the type of component and the vendor needed to be selected. At the other extreme, the semiconductor components required development activity that lasted well into the design and production of the Block II computer.

The most prominent example of the activity involved in component selection and the value of standardization in minimizing the activity required was the development of the integrated circuit NOR gate. The Block I logic design was accomplished with only one type. The initial Block II design also used one type but had to be changed to two types as a result of logic coupling in the substrate between the two independent gates on the single chip. The resulting types (a dual logic gate and a dual expander gate) differed only in interconnection pattern on the chip. Therefore the manufacturing and testing of the gates were otherwise identical, and the engineering effort could be concentrated on the development of a single device.

To select standard transistors and diodes was probably more difficult because of the wider variety of applications. The NPN transistor was a good example of this problem because the range of application varied from the very low current high frequency operation in the oscillator to the high current memory drivers and high voltage relay drivers. This range of applications stressed the state-of-the-art in transistor manufacturing, since it required a reasonably high voltage, high current type transistor. But it also required high gain at low currents as well as fast switching and low leakage. This range of applications was satisfied by the development (or selection) of a transistor chip with adequate electrical characteristics that could be mounted in a metal-base TO-18 header. The case configuration was selected as the result of thermal design considerations. The metal-base TO-18 header provided a package configuration with a low junction-to-case thermal resistance. Transistors for a relatively few special circuit applications, such as the oscillator, which required high gain at low current, could be selected during computer assembly from the distribution of parameters available in a procurement lot. This standardized the transistor production, qualification, and testing up to module fabrication. To select a standard PNP transistor was a problem similar to the NPN. Diodes were standardized to one type and selected for special application like the matching of forward voltage drop in the rope sensing circuits.

A few circuit applications could not be met using these standard parts. Most instances were in the power supplies, where very high power and current were required. Comparing the effort of specifying, evaluating, qualifying, and monitoring a low usage component to that of a high usage component illustrates the advantages of standardization. As an example, consider the high current switching transistor used in the pulse width modulated power supply. This component is a single usage item but had vendor and application troubles several times during the computer production. Individual problems with this device consumed as much analysis effort as comparable problems with the high usage component.

3.5 DESIGN QUALIFICATION AND PRODUCTION CONTROLS

To produce a reliable computer and ensure that it has, in fact, met its design objectives regarding reliability, it was necessary to institute a regime of design and production qualification, as well as quality and process controls, both for component production and for assembled units. Testing was required at many levels of assembly to ensure that design objectives and specifications were met. In addition, all components, modules, and one complete computer were subjected to a series of qualification tests. In the case of component procurement, process controls were established, but the use of captive or special high-quality production lines to achieve control was avoided.

3.5.1 Component Qualification

Components were qualified differently depending on their criticality and production maturity. A specification control drawing (SCD) was prepared; a nominal amount of engineering evaluation was conducted; the parts were released for production procurement; and then subjected to the component flight qualification program. These parts had no screen and burn-in requirement other than that which was specified in the SCD. Critical parts, like the integrated circuits and high usage transistors, followed the more rigorous procedure of engineering qualification and production screening. The DSKY relay and the standard diode followed a procedure between these two extremes where the engineering evaluation and qualification were minimized, but a tightly controlled screen procedure was introduced as a requirement fairly late in the program.

All parts were subjected to testing or data analysis sufficient to establish that the part was qualified for in-flight operation. The qualification of critical components like the integrated circuits required considerable development, since the technology was new and very little history had been developed that would lead to a knowledge of the component reliability.

The engineering qualification process of the critical parts began with an assessment of the vendor's ability to supply devices, the institution of component standardization in designs, the generation of specification control drawings and the preliminary study of device failure modes. Qualification procurements that supplied parts for the engineering qualification testing and engineering evaluations established confidence in the manufacturer's device processing and provided data on the device failure modes. Conclusions from the failure mode analyses were supplied to the manufacturer who then applied corrective action. This cyclic procedure was continued until the most obvious problems were eliminated. Knowledge of the failure modes and methods of exciting the failure modes were used to design the test environments and rejection criteria of the component screening procedures.

The design of the qualification testing procedure considered the conditions of the component application and the most likely failure mechanisms. Because these tests used small sample sizes, approximately 100 from each manufacturer, only those mechanisms with a reasonably high probability of excitation could be detected, even though the tests and failure analysis were carefully conducted. It was also extremely important that all qualification and engineering testing be performed on devices fabricated from processes as near identical to computer production as possible. The qualification method that was used subjected the devices from various vendors to environmental extremes beyond usage conditions in an attempt to identify failure modes that could occur in normal applications. This method, commonly called the step stress technique, was used in most cases but, since the same lot of devices was subjected to different stress levels serially, care had to be exercised in the analysis of failures in order to determine which test condition caused the failure. Based on the results of step stress tests, vendors were selected, and test conditions for screen and burn-in were verified.

3.5.2 Production Procurement

Engineering qualification and evaluation tests determined those vendors capable of supplying the semiconductor part without serious reliability problems. Qualification tests alone were insufficient to determine the ability of a vendor to control his process and continue to deliver a quality product. Large volume production procurement of a high reliability part requires continuous monitoring and process control to insure that the quality demonstrated in the qualification tests is maintained during the production cycle. The requirement for this continued monitoring of vendor quality and processes was written into the procurement and processing specifications.

A Flight Processing Specification (FPS) was developed in response to apparent and real reliability needs. The need for the FPS or its equivalent evolved from a great deal of data and also from sobering history. At the outset of the program there were many component problems. One instance occurred when the reliability group stated that some parts should not be used in fabricating computers. However, because of production schedule pressures, the faulty components were used, and, as predicted, the modules with these defective parts developed failures and had to be scrapped. This constant conflict between production schedules and reliability required that the reliability be better defined with a quantitative measure of the quality before the component was released to production. A reliability specification similar to the SCD was required. Then, the quality of parts, on a lot basis, could be evaluated from quantitative data. The FPS became the tool that generated quantitative data for determining the quality of a lot of components. It became apparent after considerable experience that the FPS forced component part process control without explicitly stating it, while the NASA quality specification* stated process control without the ability to enforce it. That is, the NASA quality specification required that processes would be documented and not changed

* NPC 200-2 "The Quality Program Provision for Space System Contracts", April 1962.

without approval. However the FPS provided vendor motivation because lots would be rejected, if the vendor lost control of the process in such a way that the change was reflected in the visual inspection of product quality.

From a position of technical director for the APOLLO system, the only means available to ensure the required reliability was to impose the flight process specifications as a contractual requirement. One benefit of this requirement was that the APOLLO managers became aware of component reliability and actually used the data as a quantitative tool in the management decisions. The main purpose of the FPS was to establish a firm non-varying procedure that would provide data whose significance could be easily understood. One major drawback in most reliability procedures is that without a firm non-varying procedure, it becomes impossible to assess the importance of isolated failures or component anomalies. There must be complete knowledge of the order of testing, the method of testing, and the method of reporting failures to evaluate the significance of the single failure.

Another side effect was briefly discussed previously. APOLLO experience showed that component reliability could be compromised when a higher priority was placed on production schedules, and there was no requirement for documentation that identified the compromise. The reliability required by the NASA quality specification, although imposed upon the contractor, did not provide the detailed reliability procedures necessary to make the requirement effective. This is not a criticism of the NASA quality specification. It would be impossible to write a specification that would detail all things for all components. The details of a general specification are the responsibility of the prime contractors. The flight processing specification did indeed contain the detailed description of how to execute the requirements of the NASA quality specification.

In general the FPS approach turned out to be such an iron clad document that no deviation was possible without a waiver. Although a deluge of controversy

followed, and pressure was applied to loosen the requirements, it was felt that every conceivable effort should be expended to provide highest possible quality components for production. A good procedure, therefore, would highlight component problems and not success. If the FPS was to be a good management tool, the deviations and problems must appear for management decision via the waiver route. In contrast, loosening the requirements would create fewer waivers and would create the condition where the requirement for reliability was paid for, but not documented, and not necessarily realized. The waiver, indicating the lack of reliability, became part of the data package for a computer and provided documentation for judging the reliability of the computer years after the components were tested.

In the flight processing procedure, the devices, procured by lots, proceed through the screen and burn-in test sequence to determine whether the lot is qualified for flight. That is, the FPS procedure is a lot-by-lot flight qualification in contrast to the more normal procedure, where a part or vendor is qualified by testing a typical production run rather than depending upon process control to ensure that the quality is maintained.

After completion of screen and burn-in tests, the lot is stored until failure analysis is completed. After failed units are catalogued, analyzed, and classified to complete the lot assessment, a written report is prepared and, if the lot passed, the devices that passed all tests are identified with a new part number as a flight qualified part and sent to module assembly. A semiconductor part with the flight qualification part number is the only part that can be used in flight qualified computer assemblies. From failure analysis, rejected parts proceed to reject storage, where they will be available for future study. Failed lots are rejected, unless special analysis and consideration qualifies the part for flight computer production by waiver. The waiver was required to be authorized by NASA and to accompany the computer as part of the data package.

The accumulated data, from the screen and burn-in sequences and failure analysis, were used to evaluate vendor production capability, device quality, reliability, and continued status as a qualified supplier.

In particular, the flight process specifications specify the following:

1. The operational stress, environmental stress, and the test sequence. This testing procedure is referred to as the screen and burn-in process.
2. The electrical parameter tests to be performed during the screen and burn-in procedure.
3. Definitions of failures. Failures have been defined as catastrophic, several categories of noncatastrophic, and induced.
4. Disposition of failures. The conditions are defined for removing failures from the screen and burn-in procedure and forwarding them to failure analysis or storage if failure analysis is not necessary.
5. Failure mode classification. Failure modes are classified in groups according to screenability and detectability of the failure mode.
6. Maximum acceptable number of failures per classification.
7. Maximum acceptable number of failures for non-electrical tests such as leak test, lead fatigue, etc.
8. A report for each flight qualified lot. The report must contain the complete history of the lot with the specific data and analysis required for flight qualification.

9. Rejection criteria for internal visual inspection. They are applied by the device manufacturer during a 100% preseal inspection for removal of defective parts, and by the customer on a sample basis as a destructive test for lot acceptance as part of the requirements of the FPS.

3.5.3 Production Process Controls

Strict process controls are used throughout procurement and assembly. The component procurement processes include the identification of critical processes and the establishment of methods for process control. Assembly processes like welding, wirewrapping, and potting are specified and are under tight control. As an example, in the case of welding all lead materials are controlled. The weld setting of the welding machine is specified for every set of materials to be welded, and the in-process inspection procedures are established. Periodic quality control inspections are made on each welding machine to verify that the machine and the operator are producing weld joints that can pass destructive-type tests. The material, size, and shape of electronic component leads are standardized where possible without sacrificing the reliability of the component. The standard lead materials used are kovar, dumet, and nickel. The interconnection wiring is nickel, thus limiting the number of different kinds of weld joints that must be made during assembly. The fact that the process of welded interconnection lends itself to tight control was one of the primary reasons for its use in the APOLLO computer design.

3.5.4 Final Acceptance Tests

Final acceptance procedures were designed to test the functional capability of the computers and DKSYS in addition to subjecting the assemblies to stresses that would excite potential failure mechanisms. These test procedures were used for all testing whether the computer was being sold off, returned for repair, etc. The test conditions were not to be exceeded for any flight computer.

The final assembly was subjected to extreme vibration, temperature, and voltage that were in excess of the maximum mission requirements. The modules are subjected to temperature cycling, operational tests under thermal extreme, and in some cases operational vibration tests to detect design and workmanship defects. Some of the tests that were specified initially were changed to increase their effectiveness as a screen. The history of vibration testing as applied to the detection of component contamination represents an example of how the procedures were changed to increase the effectivity.

Briefly, the history of vibration testing starts with sine vibration that was changed to random. Later the vibration axis of the computer was changed to increase the sensitivity to logic gate contamination, and finally operational vibration of individual logic modules was introduced. The computer long-term aging test is an example of decreasing the requirement, since the test was not contributing significantly to the screening of potential failures. The Block I long-term aging required 200 hours operating time before sale of a computer. In Block II the requirement was reduced to 100 hours, since the experience during the Block I testing and in field operations indicated that no potential failure mechanisms were being detected by the test.

4. PROJECT EXPERIENCE

The preceding sections have been concerned with matters of design and specification of the AGC. This section treats problems with actual components or entire computers after the design and specification stage. The first part deals with problems uncovered in the manufacturing process; the second, with problems uncovered in the field.

4.1 MANUFACTURING PROBLEMS

The manufacturing problems during the development and production phase of the program were primarily concerned with obtaining or maintaining a

component quality level that might be considered beyond the state-of-the-art for even high-reliability components. Some problems were caused by the component design or the manufacturing processes. Other problems were the result of a discrepancy between the component application and its design characteristics. The former were usually detected by means of the FPS; the latter, during computer assembly and test.

4.1.1 Component Defects

The types of component quality problems experienced during production can be illustrated by problems with the switching diode, the two switching transistors, the NOR gate, and the relays used in the DSKY.

4.1.1.1 Diodes

Three major problems with the switching diode were: junction surface instabilities detected by increases in reverse leakage current, intermittent short circuits caused by loose conducting particles entrapped within the package, and variation in forward voltage drop.

4.1.1.2 Transistors

All significant transistor problems were related to the internal leads and lead bonds. They were: "purple plague" which results in open bonds caused by aluminum rich, gold-aluminum intermetallic; a time-dependent failure mode resulting from motion in the aluminum lead wire when the transistor was switched on and off at a relatively slow rate; and occasional die-attach problems that caused difficulty in applications that required low thermal resistance for proper heat conduction.

4.1.1.3 Block II Flatpack Dual NOR Gate

The three major problems with the dual NOR gate were package leaks and leak testing; open bonds caused by a gold rich, aluminum-gold intermetallic; and shorting caused by loose conducting particles.

The problem with loose conducting particles in the logic gate is of special interest. It developed in severity throughout the production cycle. The change in severity of the problem was due in part to an increased awareness of the problem, and in part as a result of corrective action to alleviate some poor die attach problems. The corrective action was a harder die scrub during die attach that resulted in gold "pile up" around the chip. The "pile up" would break loose thus becoming a source of conductive particles within the package. Other sources are pieces of lead material, gold-tin solder from the cover sealing process and chips of silicon.

The corrective actions to solve the contamination problems started by introducing vendor internal visual inspection changes in December 1966. By August 1967 MIT/DL had completed a study on the use of X-Rays as a screen and had attempted to change the FPS to provide for a 100-percent X-Ray screen. The change was not processed until August 1968 because of many debates about the effectiveness of the screen. To illustrate this lack of an agreement, the following is a quote from one published memo: "to perform 100-percent X-Ray examination of several thousand flatpacks, looking for slight anomalous conditions indicated by white or greyish spots on the film, is not conducive to good efficiency". This attitude prevailed in management, until it became obvious that the time consumed in debugging computers with intermittent failures during vibration was not conducive to good efficiency either. When this became obvious, it was almost too late to X-Ray screen because most of the lots were in module assembly. However, the few remaining lots were processed through X-Ray, and the FPS was changed to specify the procedure.

The only remaining corrective action possible was the introduction of a module vibration test with the capability of detecting transient failures induced by mobile conducting particles. This module vibration procedure that was introduced in the early fall 1968 was effective, since no more failures occurred during computer vibration, but it was also costly and time consuming. The gross failure rate during module vibration was lower for those modules using a high percentage of X-Rayed lots, however an analysis which should determine the effectiveness of X-Ray screen has not been completed.

4.1.2 Design Defects

This section deals with manufacturing problems that were the result of marginal design or component application, in particular, the type of design problem that wasn't detected during the engineering or qualification tests of preproduction hardware. Although there were relatively few of these problems, they were of interest because they illustrate where engineering analysis or testing to worst case conditions did not excite the latent failure mechanism. The randomness of the variables that trigger the failure masked the failure mode during all the preproduction and qualification tests.

4.1.2.1 E-Memory

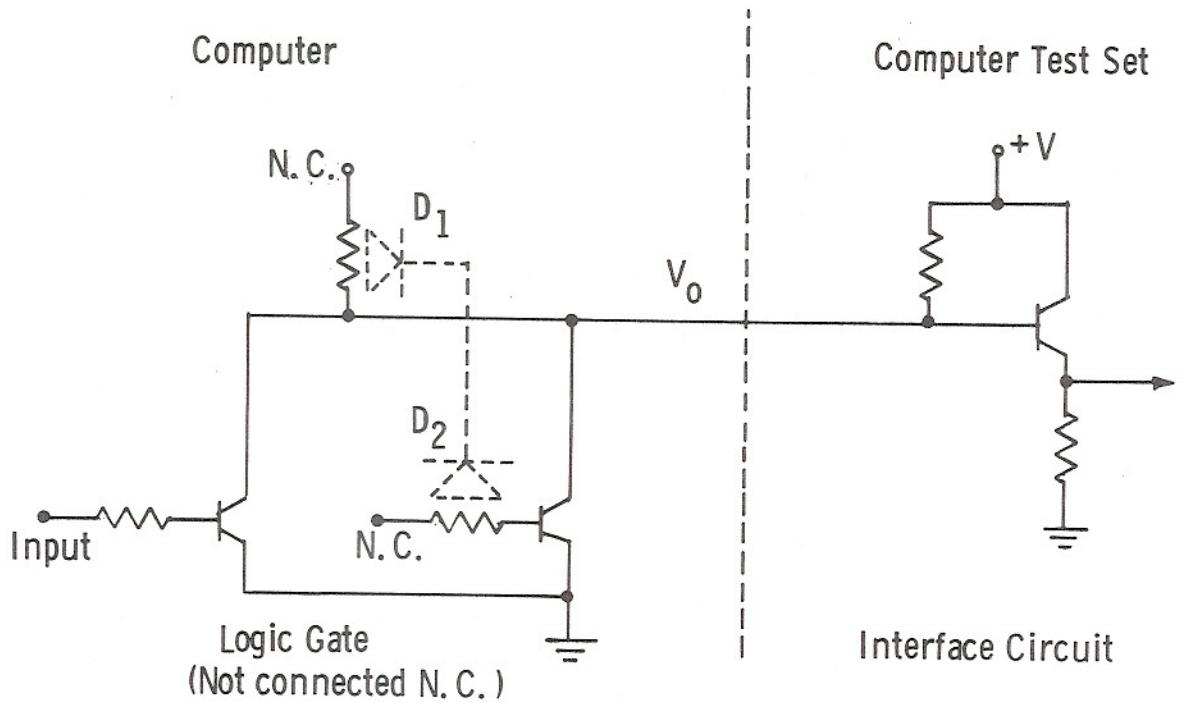
A complicated problem developed when there were several failures of the erasable memory modules due to breaks in the #38 copper wire used for internal wiring of the core stacks and from the core stack to module pins. Analysis of the breaks concluded that they occurred when the wire was subjected to tensile or fatigue stresses caused by excessive motion of the core stack and module pins within the potting material during vibration testing.

4.1.2.2 Diode Switching

Another problem was that of diode turn-on time in the rope modules caused by the fact that static matching of the forward voltage drop was insufficient and dynamic matching was required to reduce the variation in turn-on time between matched diodes.

4.1.2.3 Logic Gate

The "Blue Nose" problem is a component design problem of special interest. It occurs because a fundamental characteristic of the component was not considered in its applications. The characteristics of the isolation regions of the integrated circuit NOR gate caused the problem because: (1) the behavior of the isolation regions was not understood during the design, and (2) the engineering evaluations were not detailed enough to expose the existence of marginal conditions. The problem developed late in Block I production in the interface between the computer and computer test set. Figure 2 shows the circuit schematic, and the parasitic elements that caused the problem are shown as dotted lines. When V_o rises to about 2 volts, the diode-capacitor coupling occurs through the resistor substrate, diodes D_1 and D_2 , to the unused transistor. This coupling is a feedback path that slows the pulse rise time as indicated. The rise time will be a function of the gain of the unused transistor as well as a function of the repetition rate of the driver. Diode D_2 behaves as a capacitor that charges rapidly but discharges slowly, since the reverse impedance of D_1 is in series. The first pulse of a pulse train will be slow, and all succeeding ones faster, if the period between the pulses is small compared to the discharge time of D_2 . Since the magnitude of the effect is also dependent upon the gain of the unused transistor, it can be seen why engineering tests may not detect the problem. The condition required to detect the slow rise time is one where the transistors are high gain, and the rise time of the pulse is critical yet the data rate is low. Late in production a shift in the distribution of the transistor gain to a higher average gain caused



Integrated Circuit Gate
Illustrating "Blue Nose" Problem

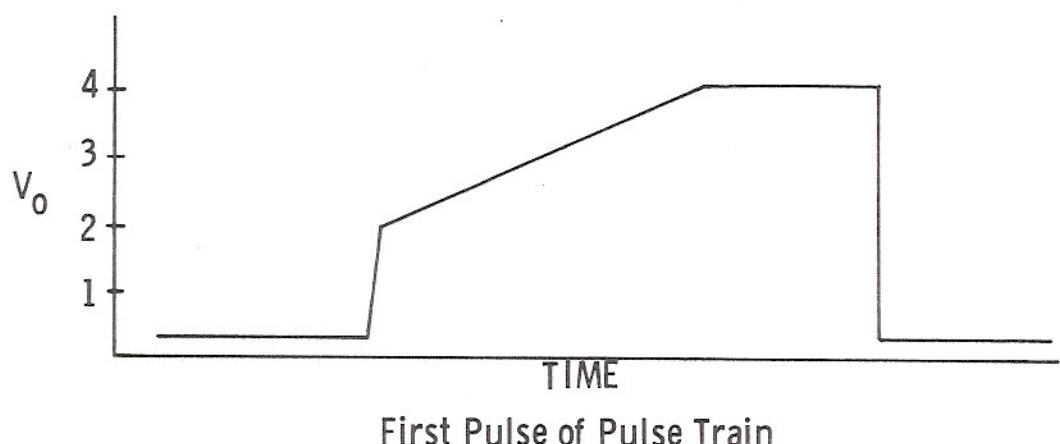


Figure 2 Blue Nose Circuit

this problem to be detected and become very troublesome. The most expeditious solution at that point in production was to select the low gain components for use in the critical locations. Another possible solution, that could not be as easily phased into production, was to ground the unused inputs of the gate.

"Blue Nose" is an expression in the parlance of the MIT logic designers indicating a logic gate used without power applied, such as a gate used to increase the fan-in. It takes its name from the graphical symbol used to denote it.

4.2 SYSTEM INTEGRATION EXPERIENCE

The system integration problems, that were experienced during GN&C and spacecraft checkout, were the most troublesome during computer development. As operator experience developed, and as the software and hardware anomalies were eliminated, checkout ran quite smoothly. Since transient or non-repeating type anomalies were the most common, it was extremely difficult to analyze the symptoms and satisfactorily explain the anomaly. Although there were many failures, and all had to be explained, there were only a few that were indications of design faults or software bugs. In general, many of the faults were the result of electrical transients of many types. Power-line transients and transient behavior of subsystems during power up and power down were the most common. The interference on signal lines, induced by operation of various switch contacts, was the result of marginal shielding and grounding. In some cases these transient signals were due to coupling within the computer between signal interface and other logic signals. All of these electrical interference problems indicated that the early computers and interface cabling were more sensitive to interference than desirable, even though the system would pass the standard EMI susceptibility specifications. A series of design changes, related to shielding and grounding, eliminated electrical interference problems except those induced by temporary power failures that would cause a V-fail alarm and a software restart.

4.2.1 Example — Software Problem

A problem, characterized by a TC Trap alarm during spacecraft testing, is typical of the type that is extremely difficult to analyze. When the actual cause of the alarm was determined, it was concluded that it was a software problem, even though the initial symptoms misled the investigators into suspecting noise as the cause. In fact, it was erroneously concluded after a brief analysis that there was no software bug. Later, after all possible hardware noise conditions had been eliminated, a software interaction was detected between test programs loaded into erasable memory and the executive activity which was located in the fixed memory.

4.2.2 Example — Hardware Problems

There was a class of integration problems that resulted from the lack of understanding about how the computer and other subsystem interfaces operated during the power-up sequences. For example:

1. When the uplink equipment was turned on, or in some cases when turned off, the equipment would emit one or more pulses. These pulses would remain in the AGC register and would cause the first data transmission to be in error, unless the register was cleared before transmission.
2. When the computer was turned on, it would indicate a warning alarm for as long as 20 seconds and would trigger the spacecraft master caution and warning.
3. When the computer was switched between standby and operate, a power transient internal to the computer would modulate the clock sync signals to the spacecraft. Sometimes the modulation would cause the down telemetry to drop out of sync for approximately one second.

These problems were relatively minor in terms of corrective action required but were troublesome to analyze. The corrective action taken was to modify the operating procedures and update the ICD to identify the signal behavior during the transient conditions.

4.2.3 Example — Mission Problems

4.2.3.1 Uplink Problem — APOLLO 6 Mission

There was one interference type problem that occurred during the APOLLO 6 mission. The AGC generated frequent uplink alarms both during and in the absence of ground initiated uplink data. Interference conditions made the process of loading data into the computer very difficult. The alarms were determined to be the result of noise on the uplink interface wiring that the computer would interpret as signal, since the noise amplitude was equal to or greater than signal.

The occurrence of noise during the mission initiated an intensive investigation that not only located the source of the noise in the spacecraft but also the sensitivity of the routing and shielding of the spacecraft cabling used on this interface. The umbilical input lines, used during prelaunch checkout and connected in parallel with the uplink input to the AGC, were determined to be the lines that were susceptible to the interfering noise. After launch these unterminated lines remained connected to the umbilical and also passed through several connectors within the spacecraft.

4.2.3.2 APOLLO 11 And 12 Examples

Both APOLLO 11 and APOLLO 12 missions had anomalies that are of interest. During the lunar landing phase of APOLLO 11, the computer in the LM signaled an alarm condition several times. These alarms were an indication to the astronauts that the computer was eliminating low priority tasks because it

was carrying a computational load in excess of its capacity. The computer was designed and programmed with the capability of performing the high priority tasks first and causing low priority tasks to wait for periods of reduced activity. Several times during the landing the computer had to eliminate low priority tasks and signaled the astronauts of this fact via the alarms.

The overload condition resulted from the fact that the rendezvous radar was on but was not in the GN&C mode. In this mode the radar angle data was being sent to the GN&C with a phase different than during normal operation. The analog to digital converters in the GN&C system could not lock onto the angle signals. The resulting hunt or dither caused a maximum data rate into the AGC counters that consumed more than 15% of the computational time. The loss of computational time was sufficient to overload the computer several times during the landing.

The APOLLO 12 anomaly was attributed to lightning striking the vehicle during the first few seconds of launch. The lightning induced temporary power failures in the fuel cell system. The transfer to the backup battery power resulted in a power transient and a condition of V-Fail in the AGC. Subsequent tests on the computer indicated no damage or loss of E-memory contents during the lightning or power transients.

4.3 FIELD FAILURE HISTORY

In addition to the problems discussed in the last section which were solved without modifying the computer hardware, there was a class of failures, the solution of which required modifications to the computer itself. Both design changes and computer repair situations are included.

In all, there were 16 computer failures and 36 DSKY failures of equipment on flight status which are of primary interest. The period of time implied by "on flight status" is defined as that part of the computer's life cycle which

begins with the date of acceptance by NASA as determined by the Material Inspection and Receiving Report (DD-250) and ends for the following reasons:

1. End of period of compilation 31 Dec. 1970.
2. Completion of flight mission.
3. Removal from flight status for other reasons (exposure to qualification environment, allocation to ground function not under quality control surveillance, etc.).

During this period of flight status and during the acceptance testing prior to acceptance by NASA, quality control surveillance was maintained, failure reports were written on all indications of anomalous behavior, and a record of operating time was accumulated. The failure experience during the factory acceptance testing was summarized in the previous section. The failures of primary interest for this section of the report are those with a "Cause" classification of "Part" in the failure reporting system. Failures with a "Cause" classification such as "Secondary", "Induced", "Procedure Error", "Test Error", "Handling", etc. are not considered here. Table I is a breakdown of the total number of failure reports written into these classifications. The DSKY failures are less interesting and are not covered in detail since DSKY components are of a largely obsolete technology (pushbutton switches, indicator panels, and relays).

There were 42 computers manufactured and delivered for flight status. Failure history has been accumulated in these systems. The first of these was delivered in the Fall of 1966 and the last one in the Spring of 1969. See Table II for the history of time on flight status for each of these computers.

Of the 16 failures, 4 are of particular interest since they are of the type for which no corrective action was taken. A complete breakdown of the failures is presented in Table III. The first four are the failures counted in the determination of an MTBF for the computer or for the prediction of a mission

TABLE I
AFR CAUSE CLASSIFICATION

FAILURE "CAUSE" CLASSIFICATION	AGC	DSKY
Development type dated before 1967	252	67
Procedure and testing errors	199	32
Induced by GSE and Cabling	150	28
Handling and Workmanship	336	42
Electrical Part	182	237
Factory acceptance testing	166	201
On flight status	16	36

success probability. The other 12 include 10 failures due to contamination in the flatpacks which were detected when a flight status computer was returned to the factory and subjected to a vibration screen more severe than the acceptance level and an order of magnitude higher than flight levels. These 10 are not counted since failures indicated during those factory test environments which are more severe than normal mission environments are not counted against the computer for purposes of reliability prediction unless they corroborate field failures. The 11th failure (also not counted) was the result of the diode design problem mentioned in the previous section. All flight hardware which is sensitive to this design problem has been purged of the defect. The 12th failure (also not counted) was a transistor bond failure at the post. This was an aluminum wire interconnect bonded to a gold plated post (not the transistor chip) which was open. Analysis indicated there was no evidence of a bond ever having been made between the wire and the post. None of the previous testing had caused the contact to open. The computer had been on flight status for over a year without indication of this defect and had been returned to the factory as part of a retrofit program to make an unrelated design change. After this retrofit, the failure was first detected when the computer was operating at the upper temperature limit of the thermal cycle. The failure was not repeatable, but after further diagnostic vibration and thermal cycling, it was again detected and located.

The population of DSKYs considered on flight status was 64 with 36 failures as noted previously. The most interesting class of failures in the DSKY is that which resulted from contamination in the relays. During the manufacturing cycle special vibration screens were developed for the component level during FPS processing, for the module level, and finally for the DSKY level of assembly. The experience of continued contamination failures during vibration testing at each level of assembly is a positive indication that the screens were not 100 percent effective. In addition, there was an indication of contamination in the main panel DSKY of the APOLLO 12 command module just before launch. Contamination of any one of 108 relays that operate the

electroluminescent panel can cause the panel to read all eights while the relay contacts are shorted by the contamination. The APOLLO 12 DSKY experienced this condition. During the mission there was no further indication of failure. Since that experience, a small test program has been developed which will cycle all relays and hopefully clear a failure if it were to occur during flight.

In summary, the contamination in flatpacks and DSKY relays has continued to plague the APOLLO program. As discussed under the Section on Manufacturing Problems, the methods for screening components were modified during the production cycle in order to increase screening effectiveness. In the case of the flatpacks, the computers at the end of the production run had the most effective screens which included 100-percent X-Ray of the components, monitored vibration at the module level, and operating vibration at the computer level. Earlier computers had various combinations of these tests but most of them had only operating vibration at the computer level. Even this test was changed to increase the effectiveness at about the mid-point of the production cycle. Experience has shown both for the DSKY and the AGC that a field return which is subjected to the latest methods of module vibration will very likely have failures due to contamination. One of the computers, after successfully flying a mission, had a contamination failure when it was returned to the factory and subjected to the vibration test. Notice that there is no evidence of contamination failures in flight.

The total history of the computers indicates there have been 58 APOLLO Failure Reports (AFR) resulting from contamination in flatpacks. Most of these occurred when the computer was being sold off initially. The 10 failures discussed previously occurred when computers were returned to the factory and were subjected to the latest vibration screens. These 10 were not indicative of any field failures. Only AFR 17275 (listed in Table III) was related to a failure during operation in the field and was verified by subsequent factory testing.

TABLE II
AGC CENSUS

S/N	DD 250 DATE	END DATE	OP TIME HOURS
16 (C-1)	7/25/66	8/23/67	1176.8
18 (C-4)	10/20/66	5/16/67	274.5
19 (C-5)	11/19/66	11/27/68	711.7
20 (C-6)	11/26/66	2/22/69	722.4
22 (C-2)	8/15/66	7/31/67	122.3
23 (C-7)	12/7/67	4/26/68	107.5
24 (C-8)	2/7/67	12/31/70	862.0
25 (C-10)	6/27/67	11/20/69	412.8
26 (C-12)	6/24/67	12/31/70	951.9
27 (C-13)	8/4/67	10/22/68	1545.8
28 (C-14)	8/23/67	12/31/70	713.9
29 (C-9)	4/5/67	12/31/70	831.8
30 (C-11)	6/10/67	1/22/68	987.7
31 (C-15)	10/12/67	5/23/69	1322.2
32 (C-16)	9/1/67	3/7/69	1613.0
33 (C-17)	10/2/67	12/27/68	1471.4
34 (C-18)	10/11/67	11/24/69	1530.7
35 (C-19)	9/6/68	12/31/70	450.4
36 (C-20)	4/30/68	12/31/70	760.4
37 (C-21)	2/8/68	3/13/69	1159.5
38 (C-22)	3/29/68	12/31/70	890.9
39 (C-23)	1/17/69	12/31/70	234.8
40 (C-24)	1/19/68	5/26/69	1206.5
41 (C-25)	12/15/67	12/31/70	771.2
42 (C-26)	1/16/68	7/21/69	1314.4
43 (C-27)	2/12/68	12/31/70	591.6
44 (C-28)	3/25/68	7/24/69	1144.9
45 (C-29)	2/26/68	12/31/70	1245.9
46 (C-30)	8/6/68	4/17/70	971.3
47 (C-31)	1/16/69	12/31/70	205.6
48 (C-32)	4/10/68	12/31/70	312.1
49 (C-33)	8/6/68	12/31/70	1064.8
50 (C-34)	7/25/68	12/31/70	367.2
51 (C-35)	4/29/69	12/31/70	207.0
52 (C-36)	3/31/69	12/31/70	302.8
53 (C-37)	9/25/68	4/17/70	524.8
54 (C-38)	2/10/69	12/31/70	377.1
55 (C-39)	3/26/69	12/31/70	0.0
56 (C-40)	5/6/69	12/31/70	217.8
57 (C-41)	9/10/69	12/31/70	254.2
58 (C-42)	5/13/69	12/31/70	91.2
59 (C-43)	5/15/69	12/31/70	154.9

TABLE III

FAILURE CLASSIFICATION
(COUNTED FOR MTBF)

AFR	DATE	LOCATION	COMPUTER	PART TYPE	FAILURE MODE
17275	2/6/69	NR	50 (C-34)	Nor Gate	Shorted interconnects Conductive contamination
17272	1/29/69	NR	43 (C-27)	Nor Gate	Open bond Gold rich, aluminum-gold intermetallic
6202	11/20/70	KSC	51 (C-35)	Nor Gate	Open aluminum interconnect corrosion
17291	5/9/69	NR	47 (C-31)	Transformer	Open primary winding

TABLE III (CONT.)

FAILURE CLASSIFICATIONS
(NOT COUNTED FOR MTBF)

AFR	DATE	LOCATION	COMPUTER	PART TYPE	FAILURE MODE
21973	11/15/68	Raytheon	25 (C-10)	Nor Gate	Conducting Contamination
20843	6/13/68	Raytheon	26 (C-12)	Nor Gate	Detected during factory
20845	6/15/68	Raytheon	26 (C-12)	Nor Gate	test environment more
21560	10/10/68	Raytheon	26 (C-12)	Nor Gate	severe than normal mission
20281	12/29/67	Raytheon	34 (C-18)	Nor Gate	(vibration).
19623	2/26/68	Raytheon	34 (C-18)	Nor Gate	
20491	4/27/68	Raytheon	34 (C-18)	Nor Gate	
20925	7/20/68	Raytheon	36 (C-20)	Nor Gate	
21919	9/19/69	Raytheon	45 (C-29)	Nor Gate	
22842	11/11/69	Raytheon	45 (C-29)	Nor Gate	
21812	11/19/68	Raytheon	25 (C-10)	Transistor	Post Bond Missing Detected during factory test environment more severe than normal mission (thermal).
19845	7/5/68	KSC	32 (C-16)	Diode	Design Problem Sensitive to forward switching time; all flight hardware purged.

5. RELIABILITY STATISTICS

In general the life cycle of the computer includes assembly and test as part of the manufacturing cycle, followed by GN&C system assembly and test (which is completed when the system is sold to NASA by means of DD250), a period of storage which includes testing to insure operability as a ready spare, installation into the spacecraft followed by a lengthy cycle of prelaunch checkout, and finally a mission. The life cycle is completed for the Command Module system at splash down. In case of the Lunar Module, the cycle is completed when the operation of the ascent stage of the LM is terminated. In the previous section this cycle was divided into two major periods: first prior to DD250, and second the remaining period defined as flight status. This latter period for each production computer is tabulated in Table II and is used for determining the reliability statistics which are summarized in Table IV. The column labeled Flight is that portion of Column D which computers have spent in flight.

This table classifies the time computers have spent in each environment and identifies each failure with the environment which induces the failure. The failure environments include: a. aging time, which is the total time since sell-off to NASA; b. vibration, which results from shipment, handling and flight; c. thermal cycle, which results from the normal turning power off and on; d. operation, which is the accumulated time the computer was operated. The aging time and operating time are derived from Table II. Vibration time is estimated from the records for shipment, handling, etc. The number of thermal cycles is estimated from operating history recorded in each computer's data package.

The failure modes listed in Table III are categorized in Table IV according to the type of environment which induces that type of failure. The two logic

TABLE IV
AGC RELIABILITY STATISTICS

FAILURE ENVIRONMENT	A AGING TIME	B VIBRATION	C THERMAL CYCLES	D NORMAL OPERATION	E FLIGHT
SAMPLE SIZE (NUMBER OF AGC'S)	42				17
AGGREGATE TIME IN ENVIRONMENT	670,000 HR.	6,500 HR.	4200 CYCLES	31,000 HR.	1,400 HR.
NUMBER OF FAILURES	2	1	1	0	0
MEAN TIME BETWEEN FAILURES	335,000 HR.	6,500 HR.	4200 CYCLES	52,000 HR.*	
APOLLO 14 MISSION	MISSION TIME IN ENVIRONMENT PROBABILITY OF SUCCESS	CM 200 HR. LM 100 HR. CM = 0.9994 LM = 0.9997	0.75 HR. 0.9999	1 CYCLE 0.9998	CM 200 HR. LM 48 HR. CM = 0.996 LM = 0.9991

$$P_S = e^{-t/MTBF}$$

$$CM \ P_S = (0.9994)(0.9999)(0.9998)(0.996) = 0.995$$

$$LM \ P_S = (0.9997)(0.9999)(0.9998)(0.9991) = 0.9985$$

*Assume 0.6 failures for MTBF computation

gate failure modes are time dependent but reasonably independent of temperature for the range of normal operation; therefore, these are assigned to the aging time column. The contamination failure is assigned to vibration. The transformer failure was an open winding which, due to the potted construction, is stressed by temperature cycling. The failure was intermittent under the conditions of computer warm up. As indicated there are no failures which are classified under operation since the failure rates associated with these four failure modes are not accelerated by the additional environments of temperatures, current, voltage, etc. which are imposed by operation.

The MTBF and success probabilities are calculated as indicated in Table IV for both CM and LM computers of the APOLLO 14 mission. For each computer, the probability of success (P_s) of the mission is the joint probability that both computers survive all environments.

6. SUMMARY AND CONCLUSIONS

From the information in Table IV and the parts count of Table V, the failure rate of various components can be calculated. The resulting numbers may be of interest, but of more interest are some conclusions that can be derived from the APOLLO experience.

1. The composite MTBF for the computer, when operating in the mission environments for an Apollo Command Module flight of 200 hours, can be computed from the results of Table IV ($P_s = 0.995$). This MTBF is 40,000 hours. If computed in the more conventional fashion by charging the four failures against the total computer hours (670,000 hours), the result is 180,000 hours. Total clock time is used in this calculation of computer hours since none of the failure modes experienced are accelerated by computer operation.

TABLE V
AGC PARTS COUNT

NAME	TOTAL	GENERIC TYPE	SUB-TOTAL
Capacitors	221	Solid Tantalum Ceramic Glass Dielectric	200 11 10
Resistors	2918	Wire Wound Tin Oxide Film	111 2807
Transistors 550		NPN Switching PNP Switching Power	443 94 13
Diodes	3325	Switching Zener	3300 25
Transformers	123	Pulse Signal	120 3
Inductors	108		
Thermistors	4		
Cores, Magnetic	35840	Ferrite Tape Wound	32768 3072
Integrated Circuits	2826	Dual Nor Gate Dual Expander Sense Amplifier	2460 334 32
Connectors - Pins	19,957		

2. It can be concluded from the material presented that the computer failure rate is independent of whether the computer is operating or not. This conclusion is based on an understanding of the physics of the failure modes experienced to date. It is also a result of a very careful thermal and electrical design which constrains operating conditions of the components to very reasonable limits.
3. A fairly reasonable development period and a reasonably large number of flight computers were necessary in order to shake down the problems and develop confidence in the reliability statistics.
4. Considerable effort was expended to make the various methods of testing and screening used in the APOLLO program as effective as possible. Even so, they were not 100-percent effective for many of the prevalent failure modes (bonds and contamination) in components being produced.
5. Contamination material in electronic components (flatpacks and relays) has shown a tendency to move around under fairly severe vibration, but has shown no tendency to float freely when at zero gravity.
6. There are long life type failure modes which are hard to predict initially and even harder to screen out of the hardware. Therefore long-term missions which require a reasonably high probability of success must depend upon techniques of redundancy and reconfiguration.