APOLLO SPACECRAFT

George M. Low
Manager, Apollo Spacecraft Program
NASA Manned Spacecraft Center
Houston, Texas

Abstract

The flawless performance of the five manned Apollo flights is attributed to reliable hardware; thoroughly planned and executed flight operations; and skilled, superbly trained crews. Major factors contributing to spacecraft reliability are simplicity and redundancy in design; major emphasis on tests; a disciplined system of change control; and closeout of all discrepancies. In the Apollo design, the elimination of complex interfaces between major hardware elements was also an important consideration. The use of man, in flying and operating the spacecraft, evolved during the course of the program, with a tendency to place more reliance on automatic systems; however, the capability for monitoring and manual takeover was always maintained. The spacecraft test effort was increased during the 18 months preceding the first manned flight with emphasis on environmental acceptance testing. This test method screened out a large number of faulty components prior to installation.

Introduction

One year ago today Apollo 7, the first manned flight in Apollo, completed its 11-day earth orbital flight. Three months ago today, Astronauts Armstrong, Collins, and Aldrin were on the return leg from man's first landing on the moon. In the intervening 9 months, there were three additional manned flights — two to the moon and one in earth orbit — all highly successful, and each completing all its assigned mission objectives. Five manned flights in 9 months, with two new, highly complex spacecraft, and each a complete success. How was this possible? What was the key to this string of accomplishments?

There is no simple answer to these questions. In fact, the task of explaining the success of Apollo in 20 or 30 minutes is almost an impossible one. For, in my opinion, the results we have achieved stem primarily from a painstaking attention to detail, by all people, at all levels, in industry and in NASA. And it is hard to tell the story of Apollo without having the time to go into the details that were at the heart of its success.

There are, of course, three basic ingredients: hardware that is most reliable; flight operations that are extremely well planned and executed; and a superbly trained and skilled flight crew.

Others in this session will talk about flight operations and the flight crew. I'll confine my remarks to the spacecraft hardware — to Columbia and Eagle — to the Command and Service Modules and the Lunar Module.

There are four aspects of spacecraft development that stand out: design, test, control of changes, and an understanding of all discrepancies.

Spacecraft Design

The principles of manned spacecraft design involve a combination of aircraft design practice and elements of missile design technology: Build it simple; and then double up on many components or systems so that if one fails the other will take over. There are many examples in our spacecraft: ablative thrust chambers that don't require regenerative cooling; hypergolic propellants that do not require an ignition source; three fuel cells, where one alone could bring the spacecraft back from the moon; series/parallel redundancy in valves, regulators, capacitors, and diodes so that neither an open failure nor a closed failure will be catastrophic. Some of these points are illustrated in the schematic diagram of the Command Module Reaction Control System (Fig. 1). This system not only has many internal redundancies, but also is duplicated in its entirety.

There is another important design rule that we have not discussed as often as we should: Minimize functional interfaces between complex pieces of hardware. In this way, two organizations can work on their own hardware relatively independently of each other. Examples in Apollo are the interfaces between the spacecraft and the launch vehicle; and between the Command Module and the Lunar Module. Between the Saturn and the spacecraft, there are only approximately 100 wires. Most of these have to do with the Emergency Detection System. The reason that this number could not be even smaller is twofold: Redundant circuits are employed; and the electrical power always comes from the module or stage where a function is to be performed. For example, the closing of relays in the launch vehicle could, in an automatic abort mode, fire the spacecraft escape motor. But the electrical power to do this, by design, originates in the spacecraft batteries. The main point is that a single man can fully understand this interface, and can fully cope with all the effects of a change on either side of the interface. If there had been 10 times as many wires, it probably would have taken a hundred or a thousand times as many people to handle the interface.

Another design question for manned flight concerns the use of man himself. Here again, there is no simple rule as to how man should interface with his machine. Generally, tedious, repetitive tasks are best performed automatically; selection of the best data source to use, or selection of control modes, or switching between redundant systems are tasks that are best performed by the pilot. In Apollo, the trend has been to rely more and more on automatic modes as systems experience was gained. For example, computer programs for rendezvous were reworked to require far less operator input than had originally been planned; but the entire rendezvous sequence was designed so that the pilot could always monitor the automatic system's performance and apply a backup solution if deviations were

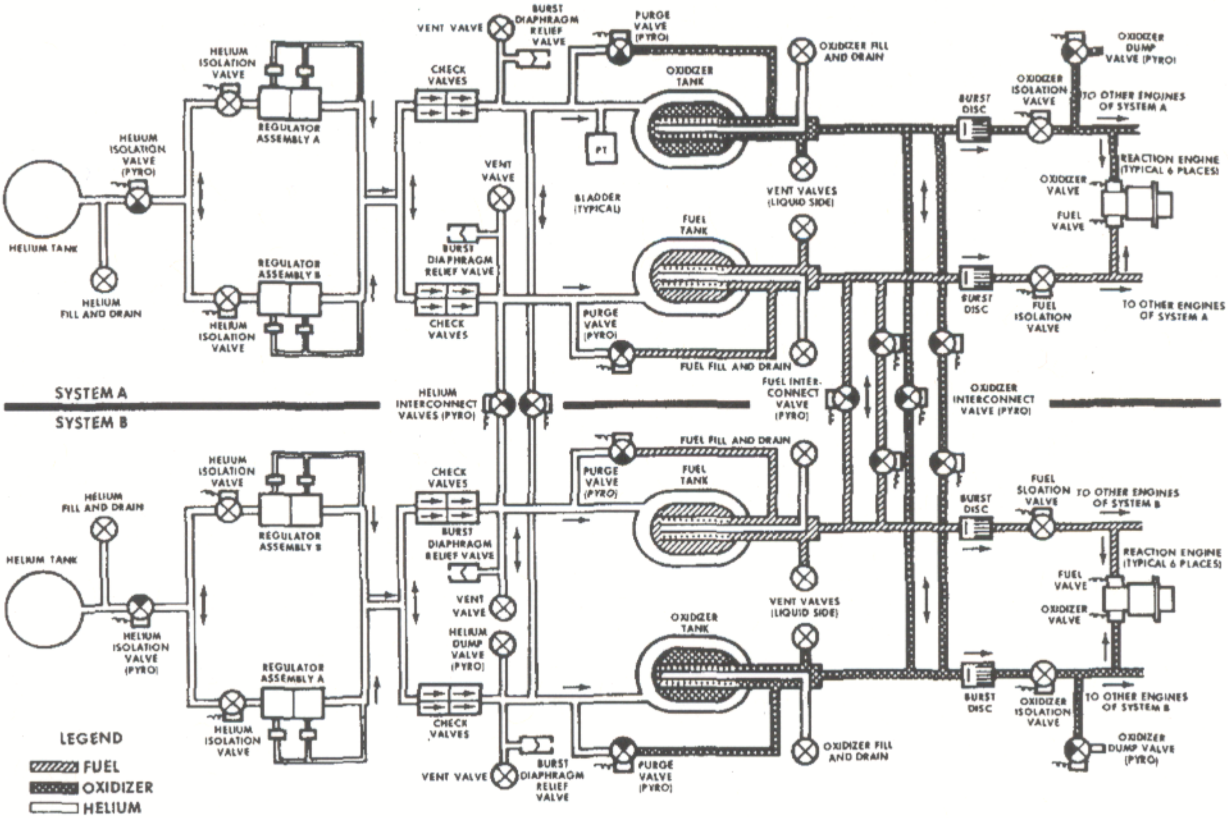

Figure 1. Command Module Reaction Control System.

noted. A tremendous amount of time and effort was spent to determine how the crew could best decide which data source to use and which of many redundant systems to rely on. This was always a basic mission design consideration.

The concept of inflight maintenance was discarded entirely as being impractical for flights of Apollo's specific purpose and duration. In its place, more telemetry was added and full advantage was taken of the ground's capability to assess systems performance, predict trends, and compare data with preflight test experience.

Apollo Test Activities

The single most important factor leading to the high degree of reliability of the Apollo spacecraft was the tremendous depth and breadth of the test activity.

There are two general categories of tests: Tests performed on a single prototype device (or on a few devices) to demonstrate that the design is proper and will perform properly in all environments; and tests performed on each flight item to assure that there are no manufacturing errors and that the item will function as intended. Both categories apply to individual parts, components, subsystems, systems, and entire spacecraft. The first category is generally called development testing or qualification testing. The second category is called acceptance testing.

Instead of reviewing the entire development and qualification test program, let us look at only those tests involving complete spacecraft or boilerplates. These are listed in Table I.

TABLE I. DEVELOPMENT AND QUALIFICATION TESTS

[Full-Scale Spacecraft Testing]

Escape motor flight tests	7 flights
Parachute drop tests	40 drops
Command Module land impact tests	48 tests
Command Module water impact tests	52 tests
Lunar Module structural drop tests	16 drops
Lunar Module complete drop tests	5 drops
Command and Service Module acoustic/vibration tests	15.5 hr
Lunar Module acoustic/vibration tests	3.5 hr
Command and Service Module modal survey testing	277.6 hr
Lunar Module modal survey testing	351.4 hr
Command and Service Modules thermal vacuum tests	773 hr
Lunar Module thermal vacuum tests	2652 hr
Service Module propulsion systems tests	1474.5 min
Ascent stage propulsion systems tests	153 min
Descent stage propulsion systems tests	220 min

Each of these tests taught us more about our spacecraft — their strengths and their weaknesses. As a result of the thermal vacuum tests, we were able to withstand the translunar and lunar environment without a single thermal problem; passive thermal control modes were developed that required minimum crew inputs and gave a perfect thermal balance. The land impact tests demonstrated that the Command Module could survive an emergency land landing, provided that the wind velocity was within certain limits. These tests also led to the design of a new couch impact attentuation strut which allowed us to increase the permissible launch wind speed and thereby gave us more flexibility in an otherwise constrained launch window. Other tests led to other significant results.

But most important of all, these tests gave us a tremendous amount of time and experience on the spacecraft and their systems. An example of this is shown in Figures 2 and 3 for the Service Propulsion System. In Figure 2, the total running time of this system is shown for both ground and flight tests prior to Apollo 11. Each unit represents one mission duty cycle — the total running time for this system during Apollo 11. Figure 3 is a similar representation for the number of times the engine has been started in terms of the five starts of Apollo 11.

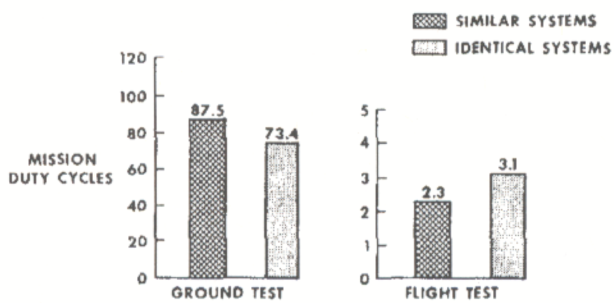

Figure 2. Service Propulsion System running time. One mission duty cycle is 532 seconds long.

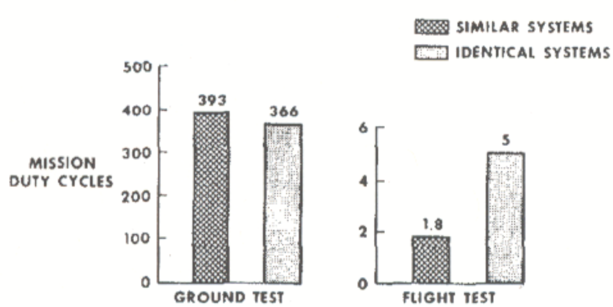

Figure 3. Service Propulsion System start experience. One mission duty cycle is five starts.

It was this kind of experience, together with a detailed analysis of all previous failures, discrepancies, or anomalies, that led us to the conclusion that we were ready to fly to lunar orbit last December on Apollo 8 and that we were ready to make a lunar landing in July of this year.

Acceptance testing played an equally important role. It starts with piece parts. Although Apollo was late in applying this rule, I believe that screened and burned-in electronic parts are a firm requirement. Next, each component, or black box, is tested before it is delivered, and again before it is installed in the spacecraft. Then factory testing of the complete spacecraft begins: First, the wiring is wrung out; then, individual subsystems are tested as installed; next, groups of systems are jointly tested; and finally, the complete spacecraft, with all its systems functioning, is run in an integrated test. All normal, emergency, and redundant modes are verified.

After delivery to the launch site, similar (when practical, identical) tests are performed. A major test at the Cape is a manned altitude chamber run of each spacecraft. The final acceptance test, of course, is the countdown itself.

A most important facet of acceptance testing is environmental acceptance testing. The primary purpose of acceptance vibration testing and acceptance thermal testing is to find workmanship errors. To do this, the environment has to be severe enough to find the fault (e.g., a cold solder joint), and yet not so severe as to weaken or fatigue the component. The levels selected for these tests in Apollo are shown in Figures 4 and 5. These were picked on the basis of experience in Gemini and other programs. Each component type, of course, had to pass qualification tests under even more severe environments. Even in spite of this rule, our environmental acceptance tests sometimes found design faults (as opposed to workmanship faults) that had been missed in the qualification tests. The reason for this is that a single qualification test may have missed a marginal condition which was later found as a result of a large number of acceptance tests.

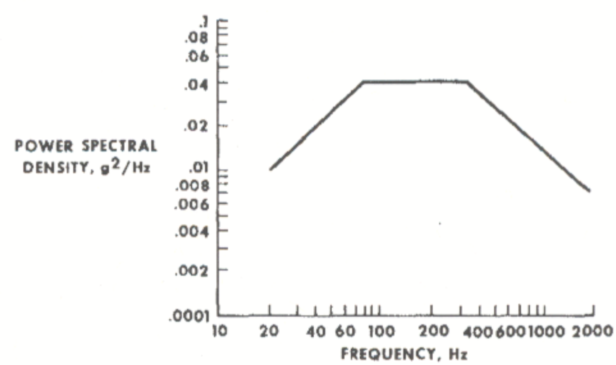

Figure 4. Acceptance vibration test level.

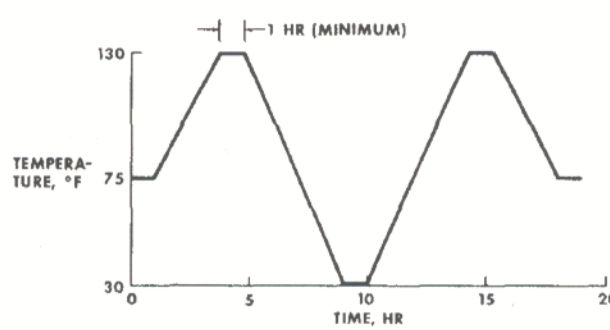

Figure 5. Acceptance thermal test level.

We also considered environmental acceptance tests of complete spacecraft, but decided against this because the environment on most components, when mounted in their spacecraft, is not severe enough to find workmanship faults. The vibration levels on many components are one or two orders of magnitude less than those given in Figure 4. (This conclusion would not be true for smaller, more compact spacecraft.) Temperatures in the spacecraft are generally constant because most electronic components are mounted on cold plates.

The results of the Apollo environmental acceptance test program are summarized in Figures 6 and 7. Note that 5 percent of all components tests failed under vibration, and 10.3 percent of all components did not pass the thermal test. Remember that these were components that were otherwise ready for installation in the spacecraft. A categorization of failure types is given in Table II. If these tests had not been performed and if these failures had occurred in flight, it might be fair to conclude that the lunar landing would not yet have been accomplished.

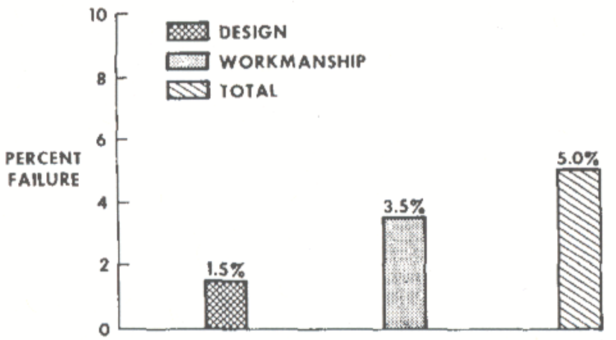

Figure 6.- Results of acceptance vibration tests for 11 447 tests of 166 different components.

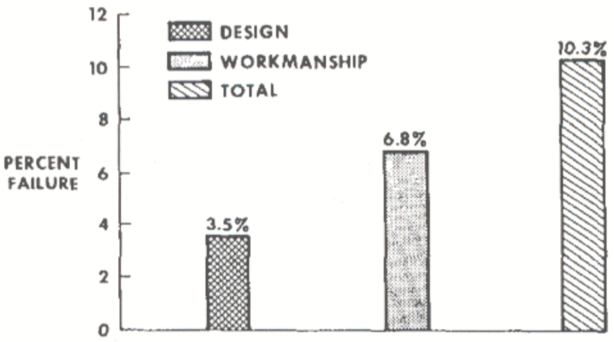

Figure 7.- Results of acceptance thermal tests for 3685 tests of 127 different components.

TABLE II. HISTORY OF ENVIRONMENTAL ACCEPTANCE TEST FAILURES

	Percent
Electrical	57.3
Mechanical	27.4
Contamination	11.5
Other	3.8
	100

Control of Changes

If the design has been verified, and if a thorough test program has been completed, it should not be necessary to make any changes. Of course, this idealized situation does not exist in any program like Apollo where design, test, and flight often overlap and must be carried out at the same time. Changes may be required as a result of test failures, or another look at the design may identify a situation that could lead to a failure, or to the inability to react to failure. Sometimes a more detailed definition of the flight missions or the operational use of the hardware itself leads to a requirement for change.

Since it is not possible to eliminate all changes, we have to start with the premise that any change is undesirable: All previous test and flight experience is voided; a change, no matter how simple, may have ramifications far beyond those that are identified by the initial engineering analysis.

Because changes must be made, it becomes important to understand them and to control all changes, no matter how small. In Apollo, we handled all changes through a series of Configuration Control Panels and a Configuration Control Board. The panels considered minor hardware changes early in the development cycle, as well as crew procedures and all computer programs. The board considered more significant hardware changes, all hardware changes after spacecraft delivery, as well as procedures or software changes that could have schedule or mission impact.

The Apollo Spacecraft Configuration Control Board met 90 times between June 1967 and July 1969. During this time, 1697 changes were considered; 1341 were approved and 356 were rejected. (The low rejection rate resulted because proposed changes were reviewed before they came to the board, and only those that were deemed to be mandatory for flight safety were brought before the board.) The board is chaired by the Program Manager who also makes the final decision on all changes. Its members are the Directors of all major technical elements of the NASA Manned Spacecraft Center. The contractor's program managers also serve as board members.

We considered changes large and small. One example of a large change would be the new spacecraft hatch that was incorporated after the fire. But we reviewed in equal technical detail a relatively small change like the need for a small piece of plastic material inside the Astronaut's ballpoint pen.

Our board was established to discipline the control of changes. But it was found to serve a much bigger purpose: It provided a decisionmaking forum for spacecraft developer and user. In reaching our decisions, we had the combined inputs of the hardware developer, flight operations, flight crew, safety, medicine, and science.

I have recently reviewed the results of the 90 board meetings that preceded Apollo 11. Even with hindsight, there are few, if any, of the board's decisions that I would make differently today than they were made at the time.

Closeout of Failures

Throughout the program, a large number of discrepancies or failures occurred on a daily basis. The relationship may have been a close one: the failure actually took place during a test of the next spacecraft to fly. Or it might have been remote: a component identical to one used on Apollo failed on another program. In both cases, the result was the same: the failure had to be understood; and, if applicable, some corrective action would be taken. Corrective action might involve a design change, or a reinspection, or perhaps a procedural change.

I will confine my remarks to the anomalies that occurred during the five manned Apollo flights. A numerical listing of these is given in Table III.

TABLE III. APOLLO FLIGHT ANOMALIES

	CSM	LM
Apollo 7	22	0
Apollo 8	8	0
Apollo 9	14	12
Apollo 10	23	15
Apollo 11	9	13

Note that even though each of the flights was completely successful and accomplished all of its objectives, the number of anomalies was quite large. Perhaps this is the best proof of the validity of the Apollo design concept: The spacecraft were designed for mission success.

The closeout of these flight failures had to be accomplished in the time available between the completion of one flight and the start of the next; a period that was generally only about 6 weeks. Yet even these 6 weeks weren't fully available to us because hypergolic propellants were loaded into the spacecraft 1 month before launch. After propellant loading, the ability to make spacecraft changes and to perform the necessary retest is severely limited. Nevertheless, each of the failures listed in Table II was satisfactorily closed out before the next flight.

Let us look at just one example: On Apollo 10, during several of the lunar orbits, a critical fuel cell temperature started to oscillate significantly (Fig. 8). Normally, this temperature is steady, between 155° and 165° F. The oscillations encountered on Apollo 10 triggered the spacecraft alarm system, but otherwise were not detrimental. Yet, unless we understood their cause, we could not be sure that they would always be limited as they were in Apollo 10 and, hence, might lead to a fuel cell failure. Our investigation revealed that small, isolated disturbances in this temperature were often present (Fig. 9). Pratt & Whitney, North American, and NASA then performed a detailed stability analysis of the fuel cell system; transfer functions were experimentally determined; and finally a complete fuel cell was run to verify the results of the analysis. The results were that the disturbances shown in Figure 9 could trigger an instability if the power loading were sufficiently high and the temperature were sufficiently low. The analysis also showed that the amplitude of the oscillations would always be limited as it was in Apollo 10. With this information, it was possible to devise procedures that would eliminate the oscillations, should they occur.

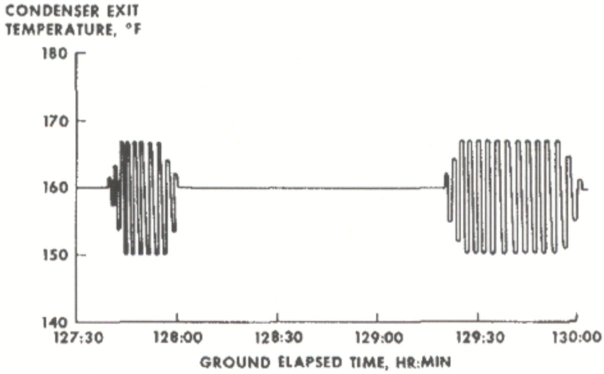

Figure 8. Apollo 10 fuel cell temperature oscillations.

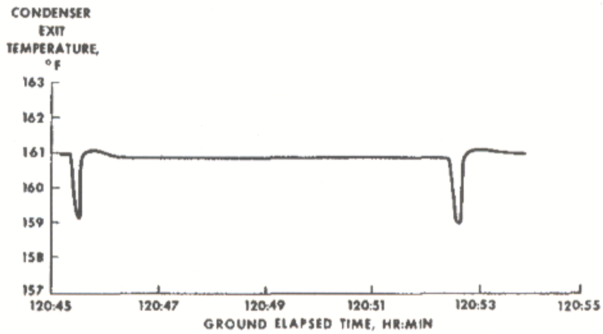

Figure 9. Apollo 10 fuel cell temperature disturbance.

The solution, as described here, probably sounds simple. Yet, a similar task, if undertaken as a research assignment, might have taken a year or more to complete. Here it was accomplished in weeks.

This was only one example of a discrepancy. The total task — that of handling all flight anomalies — was enormous; yet, it was completed prior to each flight.

Concluding Remarks

Design, test, control of changes, closeout of failures — each was important to the accomplishment of the lunar landing last July. (I realize that I left out manufacturing, partly because of time limitations, and partly because our industrial partners — North American Rockwell and Grumman — are much better qualified than I to address this subject.) But, above all, in each of these areas, what counts most is a meticulous and painstaking attention to detail by industry and NASA alike. No change was too small to consider in detail, no anomaly too small to understand. This is the story of the success of Apollo 7, 8, 9, 10, and 11. And this is what is needed to give Apollo 12, and all future manned flights in space, the highest probability to succeed.