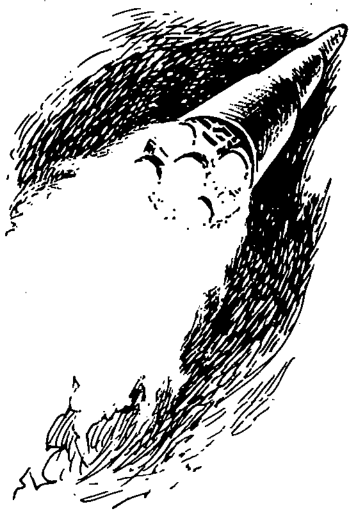

STRIVING FOR RELIABILITY*

Gerald J. Lieberman *Stanford University*

THE APOLLO mission to land men on the moon required the creation of an artificial world to support the needs of three men and complex electronic gear in a hostile environment for approximately two weeks. For a successful mission all of the components of the Apollo system had to function in a satisfactory manner. Failure in any one of innumerable areas could lead to a failure of the mission.

If reliability is defined as the probability that a device performs adequately for a specified period in a given environment, it is clear that the reliability problem was one of the greatest challenges facing the Apollo program. What is the role that probability and statistics play in reliability? Reliability is just a "probability," and hence the mathematical structure of probability theory is important in its evaluation. There are two types of general problems that arise in reliability. First, we must develop a mathematical model to represent the system, and second, we must be able to provide estimates of the numbers that enter in such a model to provide a numerical value for the reliability.

In the context of the Apollo, the command and service modules have approximately two million functional parts, miles of wiring, and thousands of joints. These parts, wiring, and joints may be broken into subsystems in some fashion. Each subsystem may be assumed to have a known reliability associated with it. A mathematical model of the Apollo system can then be abstracted from the physical processes, and the theory of combinatorial probability utilized to predict the reliability of the Apollo system. Simplified models will be discussed subsequently.

The second general problem is the statistical problem associated with reliability. Again in the context of Apollo, the system may be broken into subsystems in some fashion, with each subsystem having an assigned reliability. In practice, estimates of these reliabilities must be obtained from experimental data. Ultimately, an estimate of the overall reliability of the Apollo system based upon these subsystem estimates is desired. Alternatively, performance data may be obtained directly on the complete Apollo system, and these experimental results used to obtain an estimate of the reliability. These are some of the statistical problems associated with reliability.

RELIABILITY MODEL

In order to describe the basic problem of model development encountered in reliability theory, a simplified version of the basic Apollo module is presented in Figure 1. Suppose that the Apollo system works if, and only if, all five components shown operate properly; that is, the system will fail if any component of the system fails. Such a system is called a *series system* and a

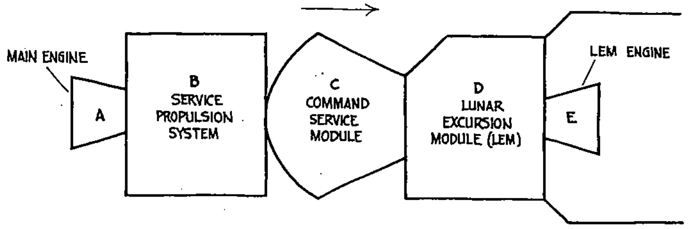

FIGURE 1
Simplified model of the Apollo module

[*] This paper was prepared while the author was visiting Imperial College of Science and Technology. It was supported in part by the Army, Navy, and Air Force under contract Nonr 225(53) (NR-042-002) with the Office of Naval Research.

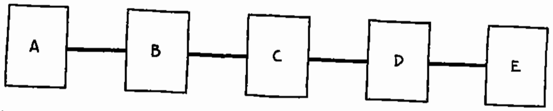

FIGURE 2
Black-box representation of series system

typical example is a string of Christmas tree lights of 1940 vintage (if one bulb fails, the entire string is darkened).

If Figure 1 is a series system it may be replaced by the slightly more abstract (*black box*) description shown in Figure 2, where each black box is in one of two states, operating or failed. The Apollo system, then, will fail if any of the five components fails during its mission. If it can be assumed that the five components of Apollo depicted in Figure 1 operate independently and are connected in series, the system reliability (the probability that Apollo performs adequately for a specified period in a given environment) is easily calculated in terms of the component reliabilities. If the component reliabilities are respectively R_A, R_B, R_C, R_D, and R_E, then probability theory indicates that the system reliability is given by R, where

$$R = R_A R_B R_C R_D R_E.$$

Thus, for example, if each component has reliability 0.999 for a mission. the reliability of the Apollo system is given by

$$R = (0.999)^5 = 0.995,$$

to three decimals.

A closer examination of these components must be made before calculating system reliability. Do the performance characteristics of these components interact with each other? Does a successful main engine performance portend that the LEM engine will also behave properly? Does degradation of one subsystem put a high load on other subsystems? If the answers to these and similar questions are positive, then the performances of the subsystems may not be assumed to be independent of each other. Unfortunately, if the components are not independent, a more complex expression is required for calculating the system reliability.

Another simple type of configuration for components is a *parallel configuration.* Suppose that two (possibly smaller) engines A_1 and A_2 replace the main engine A. Suppose further that they both are turned on during the mission and that the engine function is satisfactorily performed if either A_1 or A_2 is operative, i.e., a failure occurs if and only if both engines A_1 and A_2 fail. Such a parallel configuration can be represented as black boxes as shown in Figure 3.

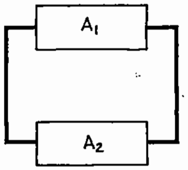

FIGURE 3
Black-box representation of parallel components

A representation of the entire system is shown in Figure 4. An actual example of a parallel system is the second stage of the Saturn rocket used in the Apollo program. This has five J.2 rocket motors, and even though, on the Apollo 13 launch, the center engine failed early, a satisfactory earth orbit was achieved by extending the burn of the remaining engines.

Again, if components connected in parallel can be assumed to be independent, the system reliability can easily be obtained from probability theory. If the reliability of the engine function is denoted by R_A and the engines have reliabilities R_{A_1} and R_{A_2} respectively, then

$$R_A = 1 - (1 - R_{A_1})(1 - R_{A_2})$$

that is, the probability that the engine power function fails, $(1 - R_A)$, is just the probability that both engines A_1 and A_2 fail. Suppose that both engines A_1 and A_2 have reliability 0.980 (considerably less than the reliability of the single main engine described earlier); then

$$R_A = 1 - (1 - 0.980)^2 = 1 - (0.020)^2 = 0.9996,$$

to four decimals, which is higher than the reliability of the single main engine described earlier. In fact, the black box representation of the Apollo system shown in Figure 4 can be replaced by that shown in Figure 2, provided the reliability of the main engine is equal to 0.9996.

A third type of configuration is a standby parallel system. If engines A_1 and A_2 were connected as in Figure 4 but so that engine A_2 was switched

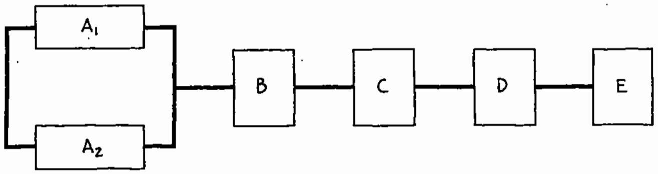

FIGURE 4
Black-box representation of Apollo system with parallel main engine units and other components connected in series

on only after engine A_1 failed, engines A_1 and A_2 would be considered as a parallel standby system. An example of a parallel standby system appeared in the ill-fated Apollo 13 mission when the astronauts actually used the oxygen supply in the LEM after the oxygen supply in the command service module failed. In both the regular parallel configuration and the standby parallel configuration, the additional units are called *redundant units,* and their existence tends to increase the reliability of the overall system (as shown in the parallel engine example) usually at the expense of cost and/or weight. In fact, some of the most interesting and important reliability design problems are those which seek to maximize the reliability of a system subject to constraints on weight or cost or both.

If all complex systems could be decomposed into independently behaving series and parallel systems, reliability calculations would become relatively simple. This is seldom the case, unfortunately, and systems such as Apollo are so complex that high-speed digital electronic computing machines are required to calculate the system reliability (assuming the reliability of individual components is known). The existence of such computing programs is also useful to assess the effect of design changes on the overall reliability of the system. In some situations, systems are so complex that detailed reliability computations are too cumbersome, and other techniques are required to assess the reliability of the system and to determine the effect of design changes.

STATISTICAL ANALYSIS OF RELIABILITY

The previous description of the assessment of the reliability of a system assumed that the reliability of the individual components was known, but nothing was said about how these numerical values were to be determined. As indicated previously this is the statistical aspect of the reliability problem. There are two ways in which estimates of the system reliability can be obtained. The simplest from a statistical point of view, but generally the most costly (and therefore most impractical), is to test by performing the mission using the final complex system. This testing is often destructive in that the system cannot be used again (e.g., launching an Apollo system) so that clearly this mode of testing is used infrequently. The other method is to test the subsystems in an environment, possibly simulated, but presumably similar to that encountered during the mission. Based upon test data, estimates of the subsystem reliabilities can be obtained, and a subsequent estimate of the system reliability then can be found.

Much of the remaining discussion will be concerned with estimating the reliability of a component because a subsystem (and even a system) can be thought of as a single component. Two important ways of testing a component directly are as follows: the first method, called *testing by attributes,* places the unit on test in the appropriate environment and for the appropriate

time and determines whether the unit fails or operates successfully; the second method, called *life testing,* places the unit on test in the appropriate environment, and the time to failure is recorded.

The statistical model for testing by attributes is relatively simple to describe. A random sample of n (independent) units are placed on test in a specified environment for a predetermined time t. It is assumed that each unit has the same reliability R, which is now assumed to be unknown. The total number of "successes" is recorded. Based upon these data, the reliability is to be estimated. In technical terms, the number of successes in the n trials is said to have a binomial distribution with parameters n and R, and standard techniques can be used to estimate the reliability. A general comment is in order. Usually, a large number of items must be tested in order that the estimate of the reliability be "good." For example, 230 consecutive units must be tested successfully in order to demonstrate (with probability 0.90) that the reliability of a system exceeds 0.99.

We will now discuss the statistical problems in life testing. Recall that units are placed on test in the appropriate environment, and the time to failure is recorded. In order to use these data to estimate the component reliability, some assumptions are usually required about the underlying failure mechanism. A simple characterization of the failure mechanism is in terms of the instantaneous failure rate, commonly called the hazard function (i.e., the probability of the unit failing in a small additional increment of time given that it has survived to the present). This is analogous to those mortality tables in life insurance, which are concerned with the probability of an individual surviving an additional year given that he has survived to the present. A reasonable hazard function that a component may possess is referred to as the bathtub instantaneous-failure function and is shown in Figure 5.

Initially, the hazard function tends to decrease, reflecting the "break-in" of marginal parts which, though operative, are operating improperly and hence may cause premature failures. The longer these parts function, the better

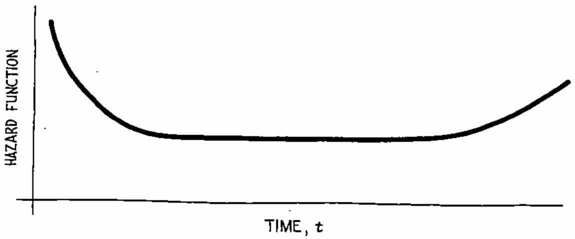

FIGURE 5
Bathtub hazard function

they perform. After this break-in period, the hazard function tends to remain constant for what is believed the normal operating period. This constant failure rate implies that no wear and deterioration are present; that is, during this period, the probability of a still working unit failing in a small additional time interval is independent of the age of the item. (It is the same for a unit that has lasted a thousand hours as for a unit that has lasted one hour.) Finally, the last part of the hazard function reflects an increasing segment which implies that wear or deterioration is present. Unfortunately, the locations of these break points are not easily determined, so erroneous assumptions about the location of the normal operating interval frequently occur.

OTHER USES OF RELIABILITY CONCEPTS

In most of the previous discussion, the Apollo space vehicle was used as an example to illustrate the importance of reliability concepts. Problems in reliability are not confined to esoteric systems, however, but they are encountered in items in everyday usage. Reliability is an important feature of household appliances, automobiles, telephones, power supplies, and so on, whether viewed from the vantage of the producer or the consumer. Important decisions are based upon the reliability of the product. For example, the five-year guarantee given by automobile manufacturers resulted from their determination of the reliability of the components falling under this guarantee. Consumers often choose to purchase the brand of items whose failure rate is low. High-reliability "consumer" products may be important from other than economic considerations. An unreliable pacemaker inserted in a heart patient could result in his death.

To summarize then, obtaining systems that perform adequately for a specified period of time in a given environment is an important goal for both government and industry. It has been frequently stated that the cost of maintenance and repair for such items as electronic equipment in their first year of operation often exceeds the purchase cost. Hence study and use of the theory of reliability, which can be applied in the research, development, and production phases of a system to enable the user to evaluate and improve performance, is a worthwhile venture. If reliability theory is to be useful, it must be quantitative in nature, because reliability must be demonstrable. Hence, probability and statistics play an important role in its development.