1

Introduction to Reliability Engineering

1.1 What is Reliability Engineering?

No one disputes the need for engineered products to be reliable. The average consumer is acutely aware of the problem of less-than-perfect reliability in domestic products such as TV sets, computers, and automobiles. Most products and industries are affected by the costs of unreliability. Manufacturers often suffer high costs of failure under warranty. Arguments and misunderstandings begin when we try to quantify reliability values or try to assign financial or other cost or benefit values to levels of reliability.

The customer, having accepted the product, accepts that it might fail at some future time. This simple approach is often coupled with a warranty, or the customer may have some protection in law, so that he may claim redress for failures occurring within a stated or reasonable time. However, this approach provides no measure of quality or reliability over a period of time, particularly outside a warranty period. Even within a warranty period, the customer usually has no grounds for further action if the product fails once, twice, or several times, provided that the manufacturer repairs or replaces the product as promised each time. If it fails often, the manufacturer will suffer high warranty costs, and the customers will suffer inconvenience. Outside the warranty period, only the customer suffers. In any case, the manufacturer will also probably incur a loss of reputation, possibly affecting future business.

Whether failures occur or not, and their times to occurrence, can seldom be forecast accurately. Reliability is therefore an aspect of engineering uncertainty. Whether an item will work for a particular period is a question that can be answered as a probability.

The most commonly used definition of reliability is: *Reliability is the probability that an item will perform its intended function without failure in specified operating conditions (or environments) for a specified period of time or usage*. The terms in this definition, such as probability, intended function, failure, specified operating conditions, and specified period of time, are all very important and carry a special meaning, which will be addressed and discussed in this book.

This definition also contains several key elements of making a reliable product. It is important to understand that reliability science is a fusion of multiple engineering subjects, including key disciplines, such as reliability statistics and physics of failure (PoF), sometimes referred to as reliability mathematics and reliability physics.

Reliability physics (this term will be used interchangeably with the term “physics of failure”) addresses the definitions of failure and of the stated conditions. It studies failure modes and failure mechanisms, which a product might experience under certain conditions, i.e., stress environments. For example, the failures caused by vibration are often attributed to fatigue, the failures experienced in high humidity environments are often caused by corrosion, mechanical shock often causes

Practical Reliability Engineering, Sixth Edition. Patrick D.T. O’Connor and Andre V. Kleyner.

fracture, and so on. Understanding the physics of failure is critical to identifying, understanding, and correcting product failures to improve the reliability and the overall product design.

Since reliability is expressed as a probability, mathematical and statistical methods are also important for modeling reliability (for prediction, measurement, assessment, etc.) and for analyzing reliability data. Statistical methods are used to define reliability as a function of time, $R = f(t)$ or a function of the appropriate usage equivalent of time, such as distance driven, number of ignition cycles, temperature cycles, mechanical shocks, ON/OFF cycles, and other product usage measures. Obtaining the reliability function, even with a degree of uncertainty, would allow an engineering professional to make an assessment of the expected reliability at the end of the product's mission life or at any time in between. This would allow us to make an assessment if the product is meeting (or not) its engineering requirements. Mathematical and statistical methods are covered in Chapters 2–6 of this book. A reliability professional needs to be knowledgeable in both key areas of reliability engineering—physics and mathematics.

Durability is a particular aspect of reliability, related to the ability of an item to withstand the effects of time or usage on failure mechanisms such as fatigue, wear, creep, and corrosion. Durability is usually expressed as a minimum time before the occurrence of *wear-out* failures. In repairable systems, it often characterizes the ability of the product to function while maintained.

The objectives of reliability engineering, in the order of priority, are:

1) To apply engineering knowledge and specialist techniques to prevent or reduce the likelihood or frequency of failures.
2) To identify and correct the causes of failures that do occur, despite the efforts to prevent them.
3) To determine ways of coping with failures that do occur, if their causes have not been corrected.
4) To apply methods for estimating the likely reliability of new designs and for analyzing reliability data.

The reason for the priority emphasis is that it is by far the most effective way of working, in terms of minimizing costs and generating reliable products. The primary skills that are required, therefore, are the ability to understand and anticipate the possible causes of failures, and knowledge of how to prevent them. It is also necessary to have knowledge of the methods that can be used for analyzing designs and data. The primary skills are nothing more than good engineering knowledge and experience, so reliability engineering is first and foremost the application of good engineering, in the widest sense, during design, development, manufacture, and service.

Overriding all of these aspects, though, is the management of the reliability engineering effort. Since reliability (and very often safety) is such a critical parameter of most modern engineering products, and since failures are directly or indirectly generated by the people involved (designers, test engineers, manufacturing, suppliers, maintainers, users), it can be maximized only by an integrated effort that encompasses training, teamwork, discipline, and application of the most appropriate methods. Reliability engineering specialists cannot make this happen alone. They can provide support, training, and tools, but only managers can organize, motivate, lead, and provide the resources. Reliability engineering is a team effort and, ultimately, effective management of engineering.

1.2 Why Teach Reliability Engineering?

Engineering education is traditionally concerned with teaching how engineering products work. However, the ways in which products fail, the effects of failure, and aspects of design, manufacture, maintenance, and use that affect the likelihood of failure are not usually paid as much attention in engineering schools. The engineer's tasks are to design, make, and maintain the product so that the failed state is deferred. In these tasks, an engineer faces the problems of variability of engineering materials, processes, and applications. Variability and chance play an important role in determining the reliability of most products. Basic parameters like mass, dimensions, friction coefficients, strengths, and stresses are never absolute but are in practice subject to variability due to process and material variations, human factors, and applications. Some parameters may also vary with time. Understanding the laws of chance and the causes and effects of variability is, therefore, necessary for the creation of reliable products and for the solution of problems of unreliability.

Competition, the pressure of schedules and deadlines, the cost of failures, the rapid evolution of new materials, methods and complex systems, the need to reduce product costs, and safety considerations all increase the risks of product development. Figure 1.1 shows the pressures that lead to the overall perception of risk. However, in today's world reliability is almost taken for granted, i.e., a consumer expects the product to be reliable, whether an automobile, mobile phone, appliance, or any other device, although it takes effort and a lot of work behind the scenes to achieve the expected level of product reliability.

Later chapters will show how reliability engineering methods can be applied to design, development, manufacturing, and maintenance to control the level of risk. The extent to which the methods are applicable must be decided for each project and each design area. They should be used to supplement good engineering practice. However, there are times when new risks are being taken, and the normal rules and guidelines are inadequate or do not apply. Sometimes we take risks unwittingly when we assume that we can extrapolate safely from our present knowledge. Designers and managers are often over-optimistic or are reluctant to point out risks about which they are unsure.

It is for these reasons that an understanding of reliability engineering principles and methods is now an essential ingredient of modern engineering. Despite its obvious importance, quality and reliability education, for some reason, is insufficient in today's engineering curricula. Few engineering schools offer degree programs or even an adequate number of courses in quality or reliability methods. This causes the demand to exceed the supply of reliability professionals; therefore, a majority of quality and reliability practitioners receive their professional training from colleagues, professional seminars, publications, and technical books, like this one. The lack of formal education opportunities in this field greatly emphasizes the importance of technical publications for professional development. Reliability texts, including this book, serve the purpose of reliability education and further contribute to the art, science, and beauty of the reliability engineering discipline.

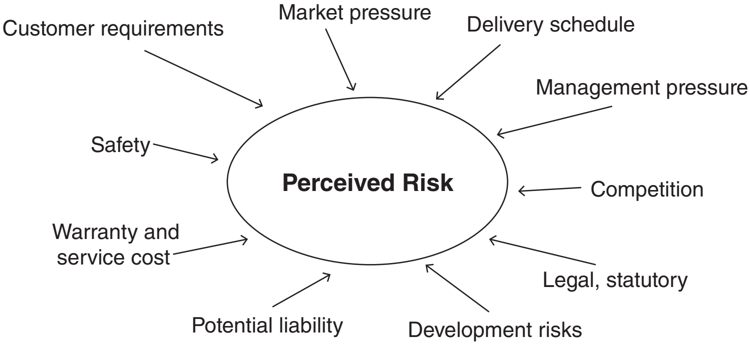

Figure 1.1 Perception of risk.

1.3 Why Do Engineering Products Fail?

It is important to define what constitutes a product's failure. The reliability definition in the previous section mentions; therefore, it is important to understand what a product's intended function is. A failure occurs when one or more intended functions of a product are no longer fulfilled to the customer's satisfaction. Failures like a total loss of product function are obvious, but there are situations when the definition of product failure is not that clear. For example, it can be a partial loss of function or some form of a "soft failure" when a system can be reset after a restart or reboot. In some cases, a system can still be functional, but some of its technical parameters can be outside of the required limits, like a radio producing a high level of noise or slow internet speed. More on the topic of a failure definition will be covered later in this book.

There are many reasons why a product might fail. Knowing and understanding the potential causes of failures is fundamental to preventing them. The reliability engineering effort, during design, development, and in manufacture and service should address all of the anticipated and possibly unanticipated causes of failure, to ensure that their occurrence is prevented or minimized.

The main reasons why failures occur are:

1) The design might be *inherently incapable*. It might be too weak, consume too much power, suffer a vibration resonance at the wrong frequency, etc. The list of possible reasons is endless, and every design problem presents the potential for errors, omissions, and oversights. The more complex the design is, the greater this potential.
2) The item might be *overstressed* in some way. If the stress applied exceeds the strength, then failure will occur. An electronic component will fail if the applied electrical stress (voltage, current) exceeds the ability to withstand it, and a mechanical strut will buckle if the compression stress applied exceeds the buckling strength. Overstress failures such as these do happen, but fortunately not very often, since designers provide margins of safety. Electronic component specifications state the maximum rated conditions of application, and circuit designers take care that these rated values are not exceeded in service. Mechanical designers work in the same way: they know the properties of the materials being used, e.g., the yield stress and they ensure that there is an adequate margin between the strength of the component and the maximum applied stress. However, it might not be possible to provide protection against every possible stress application.
3) Failures might be caused by *variation*. In the situations described above the values of strength and load are fixed and known. If the known strength always exceeds the known load, as shown in Figure 1.2, then failure will not occur. However, in most cases, there will be some uncertainty about both. The actual strength values of any population of components will vary: there will be some that are relatively strong, others that are relatively weak, but most will be of nearly average strength. Also, the loads applied will vary. Figure 1.3 shows this type of situation. As before, failure will not occur so long as the applied load does not exceed the strength. However, if there is an overlap between the distributions of load and strength, i.e., a load value in the high tail of the load distribution is applied to an item in the low tail of the strength distribution so that there is overlap or *interference* between the distributions (Figure 1.4), then failure will occur. We will discuss load and strength interference in more detail in Chapter 6.
4) Failures can be caused by *wear-out*. We will use this term to include any mechanism or process that causes an item that is sufficiently strong at the start of its life to become weaker with age. Well-known examples of such processes are material fatigue, wear between surfaces in moving contact, corrosion, insulation deterioration, and the wear-out mechanisms of light bulbs

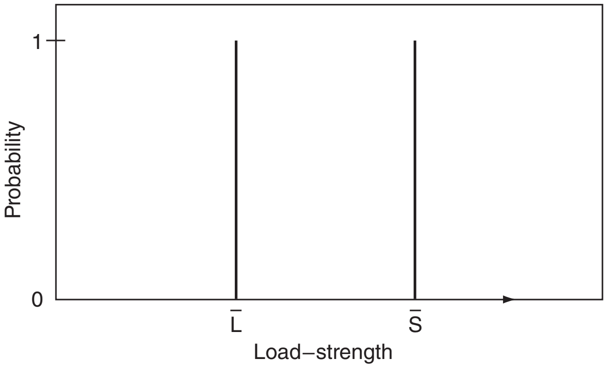

Figure 1.2 Load–strength–discrete values.

Figure 1.3 Load–strength–distributed values.

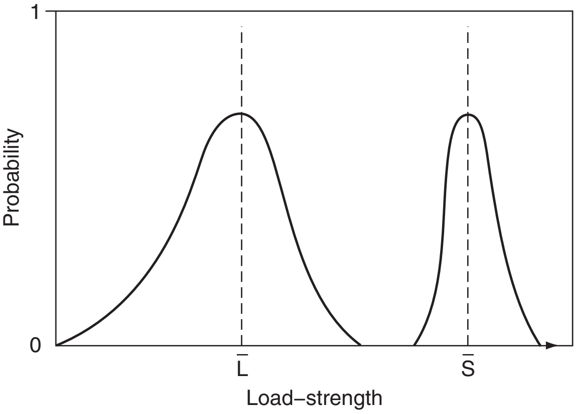

Figure 1.4 Load–strength–interfering distributions.

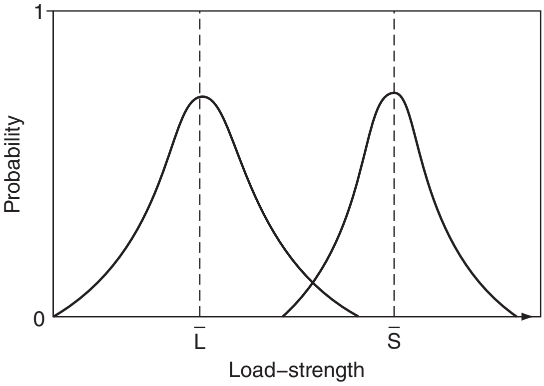

and fluorescent tubes. Initially, the strength is adequate to withstand the applied loads, but as weakening occurs over time the strength decreases. In other words, the relationship between load and strength shown in Figures 1.2 and 1.3 becomes that shown in Figure 1.4. More on this phenomenon will also be covered in Chapter 6.

5) Failures can be caused by other time-dependent mechanisms. Battery run-down, creep caused by simultaneous high temperature and mechanical stress, as in turbine discs and fine solder

joints, and progressive drift of electronic component parameter values are examples of such phenomena.

6) Failures can be caused by *sneaks*. A sneak is a condition in which the system does not work properly even though every part does. For example, an electronic system might be designed in such a way that under certain conditions incorrect operation occurs. The fatal fire in the Apollo spacecraft crew capsule was caused in this way: the circuit was inadvertently designed the way that an electrical short circuit would occur when a particular sequence was performed by the crew. Sneaks can also occur in software designs.
7) Failures can be caused by *errors*, such as incorrect specifications, designs, or software coding, by faulty assembly or test, inadequate or incorrect maintenance, or incorrect use. The actual failure mechanisms that result might include most of that listed above.
8) Failures can be caused by misuse or abuse. Not following the user guide or operating a product in conditions it was not designed to operate can lead to failure. For example, filling a fuel tank with the wrong fuel may result in engine failure.
9) There are many other potential causes of failure. Gears might be noisy, oil seals might leak, display screens might flicker, operating instructions might be wrong or ambiguous, electronic systems might suffer from electromagnetic interference, etc.

Failures have many different causes and effects, and there are also different perceptions of what kinds of events might be classified as failures. The burning O-ring seals on the Space Shuttle booster rockets were not classed as failures, until the ill-fated launch of Challenger. We also know that all failures, in principle and almost always in practice, can be prevented.

1.4 Probabilistic Reliability

The definition of reliability as a probability means that any attempt to quantify it must involve the use of statistical methods. In engineering, we try to ensure 100% reliability, but our experience tells us that we do not always succeed. Therefore, reliability statistics are usually concerned with probability values, and quantifying such numbers brings increased uncertainty since we need correspondingly more information. Other sources of uncertainty are introduced because reliability is often about people who make and people who use the product, and because of the widely varying environments in which typical products might operate.

Reliability is quantified in different ways. Besides the probability of survival, we can specify reliability as the mean number of failures in a given time (failure rate), or as the *mean time between failures* (MTBF) for items that are repaired and returned to use, or as the *mean time to failure* (MTTF) for items which are not repaired, or as the proportion of the total population of items failing during the mission life (unreliability or fraction failing). Reliability can also be quantified in terms of the product life where we expect no more than x% of the population to fail—Bx-life. Reliability metrics will be covered in detail in Chapters 2 and 3.

The application and interpretation of statistics to deal with the effects of variability on reliability are less straightforward than in, say, public opinion polls or measurement of human variations such as IQ or height. In these applications, most interest is centered around the behavior of the larger part of the population or sample, variation is not very large, and data are plentiful. In reliability we are concerned with the behavior in the extreme tails of distributions and possibly unlikely combinations of load and strength (see Figure 1.4), where variability is often hard to quantify and data are limited and often expensive or difficult to obtain.

Further difficulties arise in the application of statistical theory to reliability engineering, owing to the fact that variation is often a function of time or of time-related factors such as operating cycles, diurnal or seasonal cycles, or maintenance periods. Engineering is often concerned with change, hopefully (but not always) for the better. Therefore, the reliability data from any past situation cannot be used to make credible forecasts of future behavior, without taking into account non-statistical factors such as design changes, maintainer training, and even imponderables such as unforeseeable production or service problems. A statistician working in reliability engineering needs to be aware of these realities.

Chapter 2 provides the statistical basis of reliability engineering, but it must always be remembered that quality and reliability data contain many sources of uncertainty and variability that cannot be rigorously quantified. It is also important to appreciate that failures and their causes are often open to interpretation and argument. They also differ in terms of importance (cost, safety, other effects). Therefore, we must be careful when applying conventional scientific, deterministic thinking to the interpretation of failures. For example, a mere count of the total reported failures of a product will not provide a complete picture. It tells us nothing about causes or consequences, and therefore nothing about how to improve the situation. Nevertheless, it is necessary to derive values for decision-making as well as to understand the root cause of the problem, so both mathematics and physics are essential. The important point is that the reliability engineer or manager is not like an insurance actuary, a powerless observer of his statistics, but rather an active participant in the design improvement and decision-making process. Statistical derivations of reliability are not a guarantee of results, and these results can be significantly affected by actions taken by quality and reliability engineers and managers.

1.5 Repairable and Non-repairable Items

It is important to distinguish between repairable and non-repairable items when predicting or measuring reliability. For a non-repairable item such as a light bulb, a transistor, a rocket motor, or an unmanned spacecraft, reliability is the survival probability over the item's expected life, or for a period during its life, *when only one failure can occur*. During the item's life, the instantaneous probability of the first and only failure is called the *hazard rate*. Life values such as the mean life or *mean time to failure* (MTTF), or the expected life by which a certain percentage might have failed (say 10%), which would be expressed as B_{10}-life as mentioned in the previous section, are other reliability characteristics that can be used. Note that non-repairable items may be individual parts (light bulbs, transistors, microprocessors, fasteners) or systems comprising many parts (spacecraft, mobile phones, etc.)

For items that are repaired when they fail, reliability is the probability that failure will not occur in the period of interest *when more than one failure can occur*. It can also be expressed as the *rate of occurrence of failures* (ROCOF), or as the *failure rate* (repairable systems will be covered in Chapter 4). However, in the industry, the term failure rate became generic and has been often applied to both repairable and non-repairable systems expressing the number of failures per unit of time. Reliability metrics will be discussed in more detail in Chapters 2 and 3.

We are also concerned with the *availability* of repairable items since repair takes time. Availability is defined as the probability that a component or a system is performing its required function at a given point in time when used under stated operating conditions. Availability can be interpreted as the percentage of time the system is operating over a specified time interval, and it is affected by the failure rate and by the maintenance time. Maintenance can be *corrective*, i.e., repair of a failed system or *preventive* intended to reduce the likelihood of failure, e.g., lubrication or a scheduled part replacement. We therefore need to understand the relationship between reliability and maintenance, and how both reliability and maintainability can affect availability.

It is important to understand that repairable and non-repairable items are modeled using different statistical methods. While non-repairable items can be analyzed using statistical distributions, such as Normal, Lognormal, Weibull, and others (see Chapters 2 and 3), repairable items are modeled by a stochastic process (Chapter 4), where the same unit can fail multiple times and the number of failures in the population can potentially exceed the total number of operating units.

Sometimes an item may be considered as both repairable and non-repairable. For example, a missile is a repairable system while it is in store and subjected to scheduled tests, but it becomes non-repairable when it is launched. Reliability analysis of such systems must take account of these separate states. Repairability might also be determined by other considerations. For example, whether an electronic circuit board is treated as a repairable item or not will depend upon the cost of repair. An engine or a vehicle might be treated as repairable only up to a certain age.

Repairable system reliability data analysis is covered in Chapter 4 and availability and maintainability in Chapter 17.

1.6 The Pattern of Failures With Time (Bathtub Curve)

Life cycle product reliability is often represented by a model called the *Bathtub curve*. It is an idealized plot representing the rate of product failures over its lifetime. It is often applied to electronic products produced in large quantities but can be easily expanded to other types of devices. The bathtub curve, Figure 1.5, consists of three distinct regions representing different failure patterns of the product.

The first region is referred to as *infant mortality* or *early failures*. The failures in this region are mostly attributed to manufacturing defects, poor quality, installation issues, software bugs, etc. Many of these problems are eventually resolved and fixed, thus making the rate of failures decrease with time.

The second region is referred to as *useful life*, where breakdowns are mostly due to random events. This region is characteristic of failures which are often caused by the application of loads in excess of the design strength. For example, overstress failures due to accidental or transient circuit

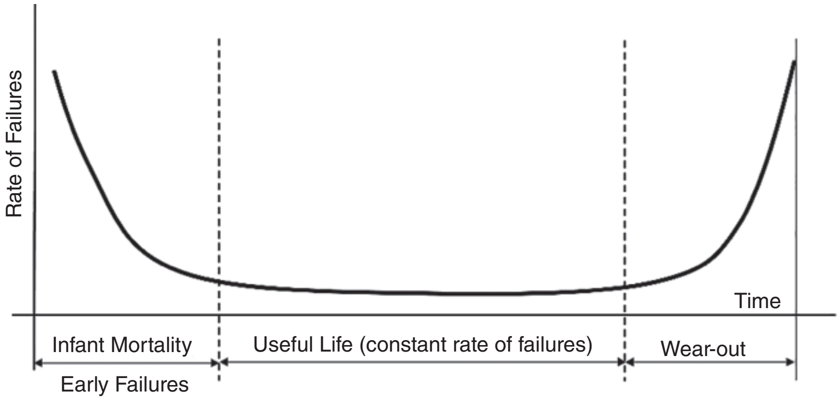

Figure 1.5 The "Bathtub" curve.

overload, or maintenance-induced equipment failures. These typically occur randomly and at an approximately constant rate. A constant failure rate is also typical of complex systems subject to repair and overhaul, where different parts exhibit different patterns of failure with time and parts have different ages since repair or replacement.

The third and probably the least studied part of the bathtub curve is the *wear-out* region. Engineering systems, especially those with mechanical parts, naturally deteriorate over time. The number of failure occurrences that a system experiences would increase after a certain age or a certain amount of usage, e.g., a number of button pushes, miles driven, or number of thermal cycles. For example, material fatigue brought about by strength deterioration due to cyclic loading is a failure mode that does not usually occur for a finite time and then exhibits an increasing probability of occurrence. Therefore, the wear-out region in the bathtub curve is characterized by an increasing rate of failures. Eventually, this leads to total failure unless the system is maintained or overhauled.

The reason why the product wear-out region is less studied is that some products do not survive long enough to experience wear-out or that field failures occurring outside of the contractual warranty period are not as well documented as failures happening during manufacturing or those coming back as warranty claims.

It is interesting to know that human life actuarial statistics follow a similar pattern, with an infant mortality phase, a steady life period, and then aging resulting at the end of life. The statistical treatment of failure data is covered in Chapters 2, 3, 4, and 13.

The bathtub curve sets expectations for how a system typically performs over its life cycle. For repairable systems, it helps to develop and plan the best repair and preventive maintenance schedule and for non-repairable systems, it helps to set the product’s target mission life. Ideal design is when the mission life requirement matches the duration of the useful life period.

Obviously, not all the products would follow the bathtub curve pattern, but this model helps to understand the general effect of failures on the lifetime reliability of engineering systems. We will continue to revisit the bathtub curve model throughout this book.

1.7 The Development of Reliability Engineering

Reliability engineering, as a separate engineering discipline, originated in the United States during the 1950s. The increasing complexity of military electronic systems was generating failure rates which resulted in greatly reduced availability and increased costs. Solid-state electronics technology offered long-term hope, but conversely, miniaturization was to lead to proportionately greater complexity, which offset the reliability improvements. The gathering pace of electronic device technology drove an increasing use of large numbers of new component types, involving new manufacturing processes, with the inevitable consequences of low reliability. The users of such equipment were also finding that the problems of diagnosing and repairing the new complex equipment were seriously affecting its availability for use, and the costs of spares, training, and other logistics support were becoming excessive. Against this background, the US Department of Defense (DoD) and the electronics industry jointly set up the Advisory Group on Reliability of Electronic Equipment (AGREE) in 1952. The AGREE report concluded that to break out of the spiral of increasing development and ownership costs due to low reliability, disciplines such as reliability and quality engineering should become an integral part of the development cycle for electronic equipment. The report laid particular stress on the need for new equipment to be tested for several thousand hours in high-stress cyclical environments including high and low temperatures, vibration, and ON/OFF cycles in order to discover the majority of weak areas in a design at an early enough stage to enable them to be corrected before production begins. Until that time, relatively short environmental tests had been considered adequate to prove the suitability of a design. The report also recommended that formal demonstrations of reliability, in terms of statistical confidence that a specified MTBF had been exceeded, be set as a condition for acceptance of the equipment by the procuring agency. A large part of the report was devoted to providing detailed test plans for various levels of statistical confidence and environmental conditions.

The AGREE report was accepted by the DoD, and AGREE testing quickly became a standard procedure. Furthermore, it was necessary to make the product speak for itself, by causing it to fail, and then to eliminate the weaknesses that caused the failures. The DoD reissued the AGREE report on testing as US Military Standard (MIL-STD) 781, *Reliability Qualification and Production Approval Tests*.

Meanwhile, the revolution in electronic device technology continued, led by integrated microcircuitry. Screening techniques, in which all production devices are subjected to elevated thermal, electrical, and other stresses, were introduced in place of the traditional sampling techniques. These screening techniques were formalized in military standards covering the full range of electronic components. Specifications and test systems for electronic components, based on the US Military Standards, were developed in the United Kingdom and in continental Europe, and internationally through the International Electrotechnical Commission (IEC).

Improved quality standards in military electronics resulted in dramatic improvements in the reliability of commercial components. As a result, during the 1980s the US Military began switching from military-grade electronic components to commercial off-the-shelf (COTS) parts to reduce development time and costs.

Engineering reliability efforts in the United States developed quickly, and the AGREE and reliability program concepts were adopted by NASA, and many other major suppliers and purchasers of high-technology equipment. In 1965 the DoD issued MIL-STD-785–*Reliability Programs for Systems and Equipment*. This document made mandatory the integration of a program of reliability engineering activities with the traditional engineering activities of design, development, and production. It was by then realized that such an integrated program was the only way to ensure that potential reliability problems would be detected and eliminated at the earliest (and therefore the cheapest) stage in the development cycle. Much written work appeared on the cost–benefit of higher reliability showing that efforts and resources expended during early development testing, plus reliability demonstrations to MIL-STD-781 led to reductions in service costs which more than repaid the reliability program expenditure. The concept of life cycle costs (LCC), or whole life costs, was introduced.

In the 1990s the series of European Reliability/Dependability[1] standards began to be developed and became integrated into the International Standards Organization (ISO). For example, IEC 60300 describes the concepts and principles of dependability management systems. It identifies the generic processes for planning, resource allocation, and control, necessary to meet dependability objectives.

Starting in the early 1980s, the quality and reliability of new Japanese industrial and commercial products took Western competitors by surprise. Products such as automobiles, electronic

---

[1] In this context dependability is defined as including reliability, maintainability, availability, and safety.

components and systems, and machine tools achieved levels of reliability far in excess of previous experience. These products were also less expensive and often boasted superior features and performance. The "Japanese quality revolution" had been driven by the lessons taught by American teachers brought in to help Japan's post-war recovery. The two that stand out were J. R. Juran and W. Edwards Deming, who taught the principles of "total quality management" (TQM) and continuous improvement. Japanese pioneers, particularly K. Ishikawa, also contributed. These ideas were all firmly rooted in the teaching of the American writer on management, Peter Drucker (Drucker, 1995), that people work most effectively when they are given the knowledge and authority to identify and implement improvements, rather than being expected to work to methods dictated by management. These ideas led to great increases in productivity and quality, and thus in reliability and market penetration, as Drucker had predicted. Many Western companies followed this new path and also made great improvements.

A large number of reliability-related standards were issued in the 1990s and through the 2000s. ISO standards regulating the ways to achieve high reliability in design including environmental testing include ISO 16750 and IEC 60068 series guiding climatic, mechanical, electrical, and other loads applied during testing and validation of electrical and electronic equipment. Functional Safety standards, such as an Automotive ISO 26262 and its parent ISO 61508, affected product design and development practices in the automotive and other industries. Many design testing and validation activities have been regulated for safety critical systems in airplanes, cars, medical and other devices. There are also a number of industry standards, and even some large companies issue their own reliability standards/requirements. For example, all the large automakers, such as Ford, Toyota, General Motors, Stellantis, BMW, Volvo, and others, have their own reliability standards and expect their suppliers to follow their procedures.

1.8 Courses, Conferences, and Literature

Reliability engineering and management courses are taught at some universities and engineering colleges although not in sufficient quantities to satisfy all the reliability engineering demands of the industry.

Conferences on general and specific reliability engineering and management topics have been held regularly in the United States since the 1960s and in Europe and Asia since the 1970s. One of the best known is the annual Reliability and Maintainability Symposium (RAMS), held in the US and sponsored by most of the important engineering associations and institutions in the United States. It is held every year and its conference proceedings are published by IEEE and contain much useful information and are often cited. The European Safety and Reliability Conference (ESREL) is also held annually and publishes proceedings on a variety of reliability topics, and more and more conferences take place in other countries and regions.

Technical journals and other periodical publications on reliability also play an important role in the continuous education of reliability professionals; some are listed at the end of this chapter. Much of the reliability literature in the past has tended to emphasize the mathematical and analytical aspects of the subject, making them difficult to apply in everyday engineering practice. However, some later publications including this book cover more practical aspects and integrate reliability work into the overall engineering and management process. These aspects are covered in later chapters.

1.9 Organizations Involved in Reliability Work

Several organizations have been created to develop policies and methods in reliability engineering and to undertake research and training. Among these organizations, it is important to mention ASQ (American Society for Quality), which has become a truly international organization with members in almost every country in the world. ASQ has many internal sections including the Reliability & Risk Division which is a worldwide professional group with a focus on reliability-specific training, education, networking, and best practices.

Also, there are other organizations involved in reliability engineering activities, such as SAE (Society of Automotive Engineers), ISO (International Organization for Standardization), SRE (Society of Reliability Engineers), PHM (Prognostics and Health Management) Society, AIAG (Automotive Industry Action Group), IEEE (Institute of Electrical and Electronic Engineers) Reliability Society, VDA (Verband der Automobilindustrie)—German Association of the Automotive Industry. These and a few other organizations are involved in promoting, assisting, and regulating engineering activities directed at achieving high product reliability.

1.10 Reliability as an Effectiveness Parameter

With the increasing cost and complexity of many modern systems, the importance of reliability as an effectiveness parameter, which should be specified and paid for, has become apparent. For example, a radar station, a process plant, or an airliner must be available when required, and the cost of non-availability, particularly if it is unscheduled, can be very high. Therefore, in order to improve a system's reliability or availability additional systems or sub-systems can be employed adding operational redundancy to the original system. For example, the Apollo project second-stage rocket was powered by six rocket motors; any five would have provided sufficient impulse, but an additional motor was specified to cater for a possible failure of one. As it happened there were no failures, and every launch utilized an "unnecessary" motor. These considerations apply equally to less complex systems, such as vending and copying machines, even if the failure costs are less dramatic in absolute terms.

As an effectiveness parameter, reliability can be "traded off" against other parameters. Reliability generally affects availability, and in this context, maintainability is also relevant. For example, automatic built-in test equipment can greatly reduce diagnostic times for electronic equipment, at a cost of a slight reduction in overall reliability and an increase in unit costs. Many other parameters can be considered in trade-offs, such as weight, redundancy, cost of materials, parts and processes, or reduction in performance.

1.11 Reliability Program Activities

A formal reliability program is necessary for virtually any product development process. We have already seen how reliability engineering developed as a result of the high cost of unreliability of military equipment, and later in commercial applications. The reliability program must begin at the earliest phase of the project. In many cases a reliability professional needs to be involved as early as the business quoting process. The reason is that engineering requirements often contain reliability targets including the methods of how these targets can be achieved, such as the required test and

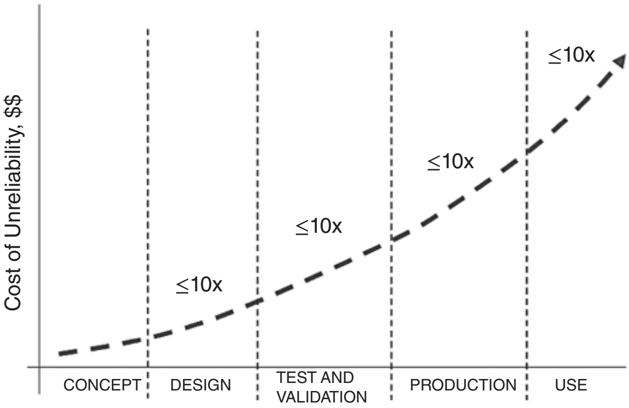

Figure 1.6 Cost of Unreliability.

validation procedures. The input of a reliability professional will help to determine the cost involved in achieving such targets and therefore contribute to the assessment of the total product's cost.

From the very beginning of the product development process reliability engineering must be integrated with the design team and help them to make the right decisions and guide the *Design for Reliability* (DfR) process. It is important to identify and eliminate the potential causes of failure as early in the design process as possible. It is widely considered that the cost of a design error grows up to 10 times from one phase of product development to the next (Figure 1.6). Obviously, this is an approximate number that would vary from industry to industry and even from product to product; however, it is important to understand what triggers this increase in cost. Figure 1.6 shows the curve of unreliability as we move from the concept phase to design then further to testing and validation of the product, and finally to production and use.

As can be seen from Figure 1.6, with each new phase the cost of finding and correcting a design error grows significantly from the simple cost of sketches, schematics, and engineering time (concept phase) to the costs of retesting, failure analysis, increased warranty cost, damaged brand and other serious financial consequences (production/operation). Individual contributors to this cost will be discussed in detail in Chapter 8.

Therefore, engineering activities at the early stages of the design process can significantly affect reliability and the overall cost of the program. These are decisions related to the risks involved in the specifications (performance, complexity, cost, producibility, etc.), development timeline, resources applied to evaluation and test, skills available, and other factors. The shorter the project's development time, the more important are these decisions, particularly if there will be few opportunities for an iterative approach. The activities appropriate to this phase are the generation of reliability objectives and the assessment of the trade-offs associated with them.

As the project proceeds from the initial study to detailed design, the reliability risks are controlled by a formal, documented approach to the review of design and to the imposition of design rules relating to the selection of components, materials, and processes, stress protection, tolerancing, etc. The objectives at this stage are to ensure that known good practices are applied, that deviations are detected and corrected, and that areas of uncertainty are highlighted for further action. The program continues through the initial hardware manufacturing and test stages, by planning and executing tests to show design weaknesses and to demonstrate achievement of specified requirements and by collecting, analyzing, and acting upon test and failure data. During production, quality assurance (QA) activities ensure that the product is manufactured as designed, but further testing may be applied to eliminate weak items and to maintain confidence. The data collection, analysis, and action process continues through the production and in-use phases. Throughout the product life cycle, therefore, the reliability is assessed, first by initial predictions based upon past experience in order to determine feasibility and to set objectives, then by refining the predictions as detail design proceeds, and subsequently by recording performance during the test, production, and in-use phases. This performance is fed back to generate corrective action and to provide data and guidelines for future products. All these activities are part of the Design for Reliability (DfR) process which will be covered in detail in Chapter 8.

1.12 Reliability Economics and Management

Obviously the reliability program activities described can be expensive. Figure 1.7 is a commonly described representation of the theoretical cost–benefit relationship of effort expended on reliability (and production quality) activities. It shows a U-shaped total cost curve with the minimum cost occurring at a reliability level somewhat lower than 100%. This would be the optimal reliability target, from the total cost point of view.

Unfortunately, it is not easy to quantify the effects of given reliability program activities, such as additional design analysis or testing on achieved reliability. The costs (including those related to the effects on project schedules) of the activities are known, and they arise in the short term, but the benefits arise later and are often much less certain. However, achieving levels of reliability close to 100% is often not realistic for complex products. Recent research on reliability cost modeling (Kleyner, 2010) showed that in practical applications the total cost curve is highly skewed to the right due to the increasing cost and diminishing return on further reliability improvements, as shown in Figure 1.8. The tight timescales and budgets of modern product development can also impact the amount of effort that can be applied. On the other hand, there is often strong market pressure to achieve near-perfect reliability. See more on the cost of reliability in Chapters 14 and 18.

It is important to remember that while achieving 100% quality in manufacturing operations, or 100% reliability in service is extremely rare in real-life applications, especially in high-volume production, it should nevertheless be considered as an ultimate goal for any product development and production program.

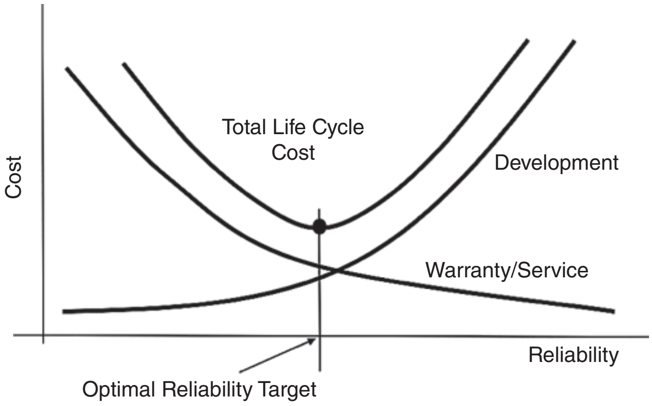

Figure 1.7 Reliability and life cycle costs (traditional view).

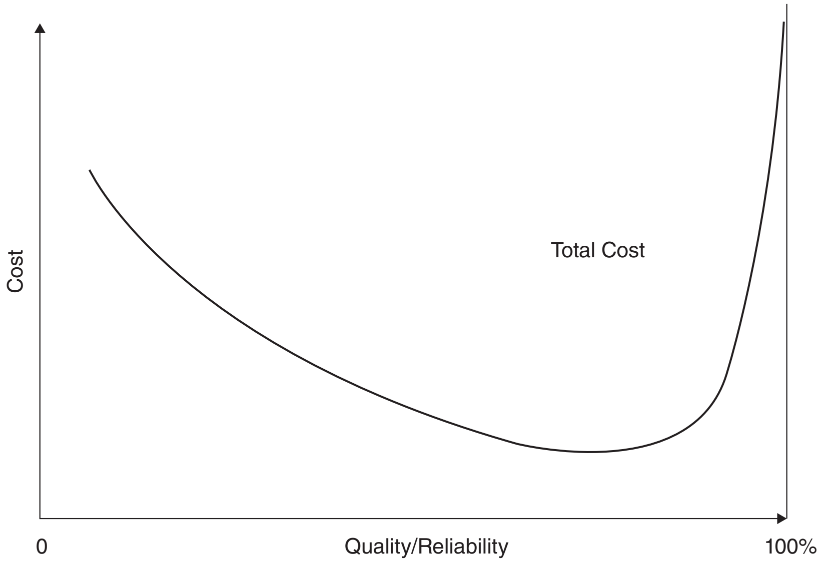

Figure 1.8 Reliability and life cycle costs (practical applications).

Achieving reliable designs and products requires a totally integrated approach, including design, testing, production, as well as reliability program activities. The integrated engineering approach places high requirements for judgment and engineering knowledge on project managers and team members. Reliability specialists must play their parts as members of the design team.

The task of design engineers is to ensure that all components are correctly applied, that margins are adequate (particularly in relation to the possible extreme values of strength and stress, which are often variable), that wear-out failure modes are prevented during the expected life (by safe life design, maintenance, etc.), and that system interfaces will not lead to failure (due to interactions, tolerance mismatches, etc.). Because achieving all this on any modern engineering product is a task that challenges the capabilities of the very best engineering teams, it is very likely that aspects of the initial design will fall short of meeting all the reliability requirements. Therefore we must submit the design to analyses and tests in order to show not only that it works, but also to show up the features that might lead to failures. When we find out what these are, we must redesign and re-test, until the final design is considered to meet the requirements. With all the modeling, simulation, and DfR techniques the expectation is that the amount of product testing will be drastically reduced compared to the old TAAF (test, analyze, and fix) approach. However, it is unrealistic to expect that in the future we can avoid testing the product altogether. With the complexity products, it is impossible to model all the possible scenarios also taking into account the unit-to-unit variation; therefore, testing will remain an important part of the product development. However, the goal in this case would be to minimize the amount of testing, thus making the product level testing a confirmation of reliable design, rather than a method to find the design weakness. More on the product testing will be covered in Chapters 12 and 13.

After the design phase, the product has to be manufactured. In principle, every unit should be identical and correctly made. Of course, this is not achievable, because of the inherent variability of all manufacturing processes, whether performed by humans or by machines. It is the task of the manufacturing people to understand and control variation and to implement inspections and tests that will identify non-conforming products. For many engineering products, the quality of operation and maintenance also influence reliability.

The essential points that arise from this brief and obvious discussion of failures are that:

1) Failures are mostly caused (directly or indirectly) by people (designers, suppliers, assemblers, users, maintainers). Therefore the achievement of reliability, besides being a team effort, is also a management task, to ensure that the right people, skills, teams, and other resources are applied to prevent the creation of failures.
2) Reliability (and quality) specialists cannot by themselves effectively ensure the prevention of failures. High reliability and quality can be achieved only by effective team(s) working together.
3) There is no fundamental limit to the extent to which failures can be prevented; therefore, in principle, we can design and build for ever-increasing reliability.

Deming explained how, in the context of manufacturing quality, there is no point at which further improvement leads to higher costs. This is, of course, even more powerfully true when considered over the whole product life cycle, so that efforts to ensure that designs are intrinsically reliable, by good design, thorough analysis and effective development testing, can generate even higher pay-offs than improvements in production quality. The creation of reliable products is also an important management task. Guidance on reliability program management and costs is covered in Chapter 18.

Questions

1 Define (a) failure rate, and (b) hazard rate. Explain their application to the reliability of components and repairable systems. Discuss the plausibility of the "bathtub curve" in both contexts.

2 **a** Explain the theory of component failures derived from the interaction of stress (or load) and strength distributions. Explain how this theory relates to the behavior of the component hazard function.
   b Discuss the validity of the "bathtub curve" when used to describe the failure characteristics of non-repairable components.

3 What are the main objectives of a reliability engineering team working on an engineering development project? Describe the important skills and experience that should be available within the team.

4 Briefly list the most common basic causes of failures of engineering products.

5 It is sometimes claimed that increasing quality and reliability beyond levels that have been achieved in the past is likely to be uneconomic, due to the costs of the actions that would be necessary. Present the argument against this belief. Illustrate it with an example from your own experience.

6 Describe the difference between repairable and non-repairable items. What kind of effect might this difference have on reliability? List examples of repairable and non-repairable items in your everyday life.

7 Explain the difference between reliability and durability and how they can be specified in a product development program.

8 List the potential economic outcomes of poor reliability. What could be the intangible losses from product failures? Based on the cost curve Figure 1.7 determine:
  a What would be the major contributors to the quality/reliability program cost? Discuss the ways to reduce these cost factors.
  b What would be the major contributors to the failure costs? Discuss the ways to reduce these cost factors.

9 After processing the existing program cost data and running a regression model on the previous projects, the cost of product development and manufacturing (CDM) has been estimated to follow the equation: $\text{CDM} = \$0.8 \text{ million} + \$3.83 \text{ million} \times \text{R}^2$ (R is the achieved product reliability at service life and is expected to be above 90%). The cost of failure (CF) has been estimated as the sum of the fixed cost of \$40 000 plus a variable cost of \$150 per failure. The total number of the expected failures is $n \times (1 - R)$, where n is the total number of produced units. Considering that the production volume is expected to be 50 000 units, estimate the optimal target reliability and the total cost of the program.

10 Select everyday items (coffee maker, lawnmower, bicycle, mobile phone, CD player, computer, refrigerator, microwave oven, cooking stove, etc.)
  a Discuss the ways this item can potentially fail. What can be done to prevent those failures?
  b Based on Figures 1.3 and 1.4, what would be an example of the load and strength for this item? Do you expect load and strength for this item to be time-dependent?

11 What is the difference between MTBF and MTTF?

12 What is increasing failure or hazard rate an indicator of? What kind of system behavior does it represent?

Selected Bibliography

British Standard, BS 4778-1 *Quality Vocabulary Part 1: International Terms*, British Standards Institution, London.
British Standard, BS 5760 *Reliability of Systems, Equipments and Components*, British Standards Institution, London.
Deming, W. (1986) *Out of the Crisis*, MIT University Press, (originally published under the title *Quality, Productivity and Competitive Position*).
Drucker, P. (1995) *The Practice of Management*, Heinemann.
Kleyner, A. (2010) *Determining Optimal Reliability Targets*, Lambert Academic Publishing.
UK Defence Standard 00–40 *The Management of Reliability and Maintainability*, HMSO.
US MIL-STD-721 *Definitions of Effectiveness Terms for Reliability, Maintainability, Human Factors and Safety*, National Technical Information Service, Springfield, Virginia.

Periodic Publications on Reliability

Quality and Reliability Engineering International. John Wiley (published quarterly).
Reliability Engineering and Systems Safety. Elsevier (published monthly)

Microelectronics Reliability. Elsevier (published monthly)
IEEE Transactions on Reliability. IEEE (published quarterly).
Proceedings of the US Reliability and Maintainability Symposium (RAMS). IEEE (published annually).
Quality Engineering. Taylor & Francis on behalf of the American Society for Quality (published quarterly)
International Journal of Performability Engineering (online journal)
International Journal of Quality & Reliability Management. Emerald Publishing (10 issues per year).
International Journal of Reliability and Safety. INDERSCIENCE Publishers (published quarterly)
Reliability: Theory & Applications. Online Journal by Gnedenko Forum www.gnedenko.net (published quarterly)