1 Basic principles of reliability, human error, and other general issues

1.1 Introduction

A number of basic qualities or conditions are of value whenever reliability is an issue. These include: (a) simplicity, (b) redundancy (providing duplicate or backup components or systems), (c) margins of safety, (d) modularity (dividing complicated things into simple components), and (e) conservatism (using conservative technology). These factors, and others, are considered in the following chapter.

Human error is, of course, a very important cause of problems in all activities. It might be thought that little can be done to prevent such errors, but this is far from the case. For example, numerous investigations have been carried out (mostly in the aviation and nuclear industries), which show that errors are generally not completely random and unpredictable events, but usually follow regular patterns. These results, which are discussed below, suggest ways of avoiding errors, or at least mitigating their consequences.

Other sections of the chapter discuss record keeping in the laboratory (the lack of which is a common cause of problems), the maintenance and calibration of equipment, and general strategies for troubleshooting apparatus and software.

1.2 Central points

The following are very general principles of reliability that recur repeatedly in all activities in research.

(a) *Simplicity* The imperative to keep things simple is usually well understood, but not always practised. It is especially important in a university environment where inexperienced research workers are involved.

A frequent cause of the erosion of simplicity is the desire to make something (experimental apparatus, computer software, calculations) as general-purpose as possible, rather than tailoring it to a specific task. Also, workers sometimes feel inclined to add extra unnecessary features to apparatus or software that is under development, in the belief that, although they are not needed immediately, these features might be useful in the future. This tendency often leads to difficulties. Another cause of troublesome complexity is the desire (particularly common among

beginners in research) to demonstrate one's ability through the mastery of complicated experimental, theoretical, or computational "machinery."

It is possible to take the principle of simplicity too far. For example, in the case of electronic circuit design, the need to provide proper circuit protection (e.g. over-voltage protection on sensitive input circuits), adequate margins of safety, and where necessary redundancy, is normally more important than reducing the number of components [1]. Furthermore, in systems in which human error during use is an important source of reliability problems, an increase in system complexity for the purpose of automating tasks can lead to an overall improvement in reliability. An example of this would be the automation of valve operations in a high-vacuum system.

An important benefit of simplicity, which is perhaps not always appreciated, is that it often simplifies troubleshooting in the event of a failure.

(b) *Redundancy* The implementation of redundancy can range from very elementary measures to relatively sophisticated ones. At the most basic level one may, for example, have a backup piece of experimental apparatus, which can be swapped with a malfunctioning unit when necessary. To take another example, it is usually feasible to provide extra wires in cryogenic instruments, as insurance against possible losses of working wires due to breakages or short circuits. Also, for instance, containers that might be subjected to gas overpressures can be supplied with two pressure relief valves, so that if one fails to open when required, the other should do so.

At the other end of the scale of sophistication, there exist multiply redundant computer systems, involving several computers that work separately on the same calculation, and then automatically compare their results. If the results of one computer are found to be different from those of the others, it is ignored, and that computer is then taken out of the system.

The use of redundancy can be a very effective method of improving reliability, and it is heavily used in areas where high levels of reliability are needed (such as on spacecraft). However, since its use normally involves additional cost, complexity, and bulk, it should not be employed as a substitute for sound design or technique. In general, redundancy should be applied only after all other methods for improving reliability have been tried, unless the highest levels of reliability are essential. (The use of redundant pressure relief devices on cryogenic vessels would be an example of the latter situation.)

There are several other situations in which redundancy is regularly used. For instance, computer hard drives can be unreliable, and in order to prevent the loss of important information in the event of failure, redundant arrays of these devices are often employed. This is not difficult to do – see page 492. Power supplies are also frequently a source of trouble (see pages 395 and 494). In circumstances in which high reliability is very important, power can be provided by redundant power-supply systems, which comprise two or more independent power supplies. If one such unit fails, the others will automatically and immediately compensate. These systems are often used in server computers. Finally, in some measuring systems, redundant sensors are employed to guard against errors due to sensor damage or loss of calibration.

It is possible for redundancy to lead to a reduction of reliability, if sufficient thought has not been given to its implementation. A classic example of this involves a twin-engine airplane with a single-engine ceiling of 1220 m above sea level [2]. The two engines are redundant only under some conditions. If the airplane is flying over Denver (with a ground height above sea level of 1610 m), the presence of two engines *doubles* the chances of crashing because of engine failure.

When using redundancy, one should always be on guard for "common mode failures," in which the benefits of having redundant elements are negated by the occurrence of a fault that affects all the elements. For example, in the case of pressure relief valves, the advantage of having two valves would be lost if they shared a common passage to the container, and this passage became blocked. In the case of redundant computer systems, the use of redundancy would be to no avail if all the computers used the same algorithm to perform the calculation, and an error occurred due to a problem with the algorithm. The use of redundancy is discussed in depth in Ref. [1].

(c) *Margins of safety* The use of a margin of safety is often applied in those cases where some operating parameter cannot pass beyond certain limits without causing failure, and one would like to take into account uncertainties or unforeseen conditions. Examples of physical parameters include power, electric current, pressure, and number of operating cycles. A specific case involves using a pressure vessel only at some fraction of its bursting pressure, in order to allow for ignorance about the material properties, uncertainties in calculations of the bursting pressure, errors in the pressure measurement, and mechanical fatigue.

The notion of a "margin of safety" is used very generally when one would like to allow for uncertainty, even when there is no element of actual physical danger. A nonphysical example is the use of extended-precision arithmetic in order to take account of round-off errors in numerical calculations. The reduction of the magnitude of a physical operating parameter in order to increase reliability is often referred to as "derating," and is discussed in more detail on page 58.

The use of a margin of safety usually involves making a tradeoff with performance (operating pressure, in the above example). This means that there is often a temptation to reduce the safety margin in an arbitrary way in order to gain performance. Needless to say, in general this should be resisted.

It has been said that, during the development of a device or system, $\frac{2}{3}$ of the difficulties and $\frac{1}{3}$ of the costs are incurred while attaining the last 10% of the desired performance [3]. While this rule-of-thumb was coined with regards to aircraft development, the sentiment is equally applicable to research. If the performance of an instrument or a technique is being pushed close to the very edge of what is possible, margins of safety will inevitably have to be compromised, and reliability problems are bound to appear more frequently.

(d) *Modularity* The management of complicated things is normally done by dividing them up into a number of simpler independent ones ("modules"). If necessary, these can be further subdivided, until the stage is reached where further subdivision is no longer necessary. This technique is a very general way of dealing with complexity.

In the case of experimental apparatus, the use of modules makes it possible to:

(i) more easily understand the workings of the apparatus by hiding irrelevant complexity within the modules,
(ii) create reliable complicated-apparatus by permitting the assembly of a number of simple units (possibly of a standard design with well characterized behavior), which can be designed and debugged easily in isolation,
(iii) readily diagnose faults by allowing suspect portions of the apparatus (one or more modules) to be easily swapped with known working ones,
(iv) quickly repair faults by making it straightforward to replace defective parts of the apparatus, and
(v) more easily implement redundancy strategies.

The use of modularity is also invaluable in writing computer software (in which the modules are called "routines"), and performing mathematical calculations (where complicated operations are broken up into a number of simpler ones, which can then be handled separately).

Although the use of modularity normally results in improving the reliability of a complex system, it should not be thought that just because the modules work reliably, the system as a whole will necessarily do so. There is a class of faults called "sneaks," in which the system as a whole fails even though the modules that comprise it work correctly. In such cases, the overall design (e.g. interconnections between components in an electronic circuit) is incorrect. Computer software bugs are often a form of sneak. Hence, it is always necessary to think about the possibility of sneaks during the design of a system. Furthermore, one must test the system as a whole once it has been assembled, and not just assume *a priori* that satisfactory behavior is guaranteed by the proper functioning of its subunits. The issue of sneaks is dealt with in more detail in Ref. [1].

(e) *The advantage of small incremental improvements* Making large changes in some parameter (sensitivity, power level, etc.) brings with it the risk of unanticipated changes in some other quality (or qualities) that could otherwise have been predicted, and accounted for, if the change were small. Small incremental improvements, made one at a time, have the advantage that since one is close to the starting point, the relationship between the change and any negative consequences is fairly simple, so that it is fairly easy to tell what needs to be done in order to make corrections. When everything is under control, another incremental change can be made. In this way, one carefully and controllably alters the parameter until the desired improvement is achieved.

(f) *Using conservative technology* If reliability is an issue, it is seldom a good idea to employ a new technology (e.g. a novel type of commercially made instrument) without giving it time to be tested by other users, and for improvements to be made as a result of their experiences. First-generation things are often best avoided. It may take years for the inevitable problems to be sorted out.

(g) *Testing versus sound design and construction* While there is no doubt about the importance of testing in gaining assurance that a given item of equipment or software operates correctly and reliably, testing is not a substitute for sound design and construction. One cannot expect that one will be able to expose by testing, and subsequently

repair, all possible potential problems. Reliability can come only from a correct design and its proper implementation. This is especially important when intermittent faults are a possibility, since the presence of these may not be detected by testing (see the discussion on page 60).

1.3 Human factors

1.3.1 General methods and habits

1.3.1.1 Introduction

In general, human error is responsible for a great many of the reliability problems that can occur. Strictly speaking, it is responsible for virtually all reliability problems. However, here we will concern ourselves only with errors taking place within the research environment, and not those occurring (for example) at a factory where some apparatus may not be designed correctly. Therefore, suitable general approaches to research, habits and abilities are very important in averting such problems.

The habit of being careful is obviously a desirable attribute, but is insufficient by itself. Experience and imagination are also invaluable in foreseeing and averting potential difficulties. Patience and attention to detail (see page 7) are yet other useful characteristics.

1.3.1.2 Finding out what is known

It has been said that six months of work in the laboratory may be saved by six hours spent in the library [4]. This is not an exaggeration, and indeed the importance of reviewing previously published work before beginning a research project is hard to overemphasize. It is not uncommon for new investigators in a field to reinvent the same techniques, and make the same blunders, as others who have already described their experiences in print. This redundant knowledge is often gained at the cost of considerable effort and expense. Ignorance of well-known pitfalls in experimental method can, and sometimes does, also lead to the publication of erroneous data.

It sometimes happens that a "new" scientific phenomenon or theory is observed or developed, and announced, only to be subsequently revealed as something that is already known. The history of science abounds with instances of such rediscoveries. An example and brief discussion of this is provided in Ref. [5].

With the availability of a huge variety of information sources on the Internet, and some excellent search engines, it is hard to justify not making the effort to find out what is known. Strategies for carrying out literature searches are discussed in detail in Ref. [4].

1.3.1.3 A digression on sources of information

Much useful information about research instruments and techniques can be found in scientific instrumentation journals and related periodicals. These include: *Review of Scientific Instruments*, *Measurement Science and Technology*, *Cryogenics*, *Journal of Vacuum Science and Technology*, *Nuclear Instruments and Methods*, *Applied Optics*, and *SIAM Journal on Numerical Analysis*.

There are numerous books on various topics related to these subjects (some of which are listed in the following chapters). If a book or journal is not present at one's own research establishment, one should consider the possibility of making use of the "interlibrary loan" facilities that are often available. The use of these is usually very straightforward. (Keep in mind that not everything is on the Internet.)

Doctoral dissertations can be useful sources of detailed technical information that does not get published. The solutions to technical problems provided by these are often inelegant and extempore. However, they can provide information about the kinds of problem that can occur and their character, and are frequently based on firsthand experience. Copies of dissertations can be obtained online from national libraries (e.g. the British Library in the UK) or commercial sources (e.g. Ref. [6]).

Company websites are often a useful source of information on the correct use of a generic type of equipment, potential problems with these, and possible solutions. Nor should one neglect printed catalogs. These can contain information that is complimentary to that provided by the websites, and are often in some ways easier to use.

Searching for journal articles containing needed information has long been possible with computer databases of journal titles and abstracts. Using appropriate keywords, it is often possible to find very helpful information amidst the vast number of articles that have been printed. Some useful databases are INSPEC, Web of Science, and Metadex. Google provides an online facility that makes it possible to search the entire contents of a very large number of journals (and not just the titles and abstracts) [7].

Using the appropriate keywords in the correct combination is very important, and the acquisition of skill in this activity can be highly valuable. One useful way of obtaining a variety of suitable keywords is to search the Internet in the normal way. Even if the results of this search are themselves not relevant, the web pages may contain words that can be used as keywords in database searches. A thesaurus (or *synonym dictionary*) can be a useful source of keywords.

Google also provides a very useful website that allows one to do an online keyword search of a large fraction of the books that have ever been printed, and to display parts or all of the book pages containing these keywords[1] [8]. This facility does for books what the computer databases do for journals. In addition, it can be used to augment the index of a printed book that one already has in one's possession. It allows searches of a particular book for words that are not necessarily contained in the book's index, and can also search

[1] As of the time of writing there are plans to allow all the pages in any book to be viewable, thereby turning this website into an online library of great comprehensiveness. However, these plans are mired in legal controversy.

for combinations of words and phrases. The latter is, of course, generally not possible using an ordinary printed index.

1.3.1.4 Periodically review the state of the art

It is a good idea to review the state of the art and capabilities in any given area from time to time – things change. One should not make assumptions about the best methods or technologies to use on the basis of what was done many years ago.

Instrumentation and methods are often in a state of flux. An entirely new and superior technology or method may suddenly appear. Or perhaps an alternative, previously inferior, technology or method may improve to the point where it becomes the one of choice.

While we are normally accustomed to things improving, they sometimes get worse. Excellent products can disappear if their manufacturers go out of business, or the product is discontinued either because it wasn't selling or because an important part that went into its construction is no longer available. Alternatively, the quality of a product may be degraded in order to cut costs, or because technical people with unique skills have left the firm that makes it.

1.3.1.5 Paying attention to detail

In technical matters, the neglect of small details (even of an apparently trivial nature) can have great consequences. Reliability problems are very often caused by simple things that have not been given sufficient attention [1]. For example, a tiny screw that has not been properly secured in an instrument might work its way loose with vibration and cause short circuits, thereby producing erratic behavior, or even damaging the device. One of the common characteristics of successful experimenters is a knack of paying attention to the right details [4]. The potential importance of such things is vividly illustrated by the following historical cases.

In 1962, the Mariner 1 space probe, which was on its way to Venus, suffered a failure that led to its destruction. The cause of this failure was a missing hyphen[2] in the spacecraft's computer guidance software [9]. Mariner 1 cost about 19 million US dollars.

In 2008, the Large Hadron Collider (LHC) particle accelerator underwent a failure that started a chain of damaging events. These led to the complete shutdown of the facility. It has been projected that it will take a total of about a year[3] and some 21 million US dollars to repair the damage and restart the accelerator. The cause of the failure was a single bad solder joint that connected two superconducting magnets being used to control the particle beam [10]. Although a very large number of such joints are present in the accelerator, a

[2] Some reports have described the error as being a missing overbar (i.e. $\overline{\dot{R}_n}$) in the transcription of the original mathematical description of the guidance algorithm. The overbar indicates an averaging operation.
[3] This section was written in early 2009.

large-scale and highly professional effort was made to ensure that they were satisfactory [11]. However, apparently this was not enough.

1.3.1.6 Difficulties caused by improvisation

In modern scientific research, there is often tremendous pressure to come up with results as quickly as possible. In some ways this is commendable. However, such pressure often leads to improvisation: in choosing and devising an experiment, designing and building apparatus, and taking measurements and interpreting the data. Often the result of such an approach is an ill-conceived research project, apparatus that is difficult to use and unreliable, data that are noisy and untrustworthy, and conclusions that are groundless. Furthermore, extempore solutions to problems can sometimes hide fundamental defects in the various aspects of a research project, and it may be difficult to tell at a later stage precisely where things have gone wrong.

Improvisation is sometimes justified in the early developmental stages of a research project, when it is desirable to validate an experimental concept (e.g. a method of measurement or a type of apparatus) that is not amenable to exact analysis. Sometimes the rapid pace of a particular field of research precludes long-term planning and construction. In such situations, improvisation at some level may be unavoidable. A common problem is that arrangements that were intended to be temporary end up becoming permanent. The subject of improvisation versus planning in the construction of apparatus is discussed in Ref. [4].

1.3.2 Data on human error

1.3.2.1 Frequency of problems caused by human error

It has been reported [12] that during human test or maintenance activity in general, the probability of a fault being put on an item lies in the range between 10^{-4} and 10^{-2} *per operation*, depending on the complexity of the task. Therefore, for a maintenance routine comprising a large number of operations, the chances of introducing a failure may be very significant. (Many things done in a research laboratory, e.g. changing samples in an instrument, are similar in character to maintenance activities.)

A study of 180 "significant event reports" at nuclear power plants between 1983 and 1984 [13] suggests that of the 387 root causes that were identified, "human performance" comprised 52% of causes, "design deficiencies" were given as 33%, while "manufacturing, etc." and "other/unknown" external causes roughly equally comprised the rest.

(NB: The precise values of the statistical data provided in this and the following sections are not important, since the statistical uncertainties are large. The numbers merely provide a means of identifying the most salient problems and their causes, and indicating the relative qualitative significance of such problems and causes.)

Table 1.1 Dominant types of human error, from a survey of error types in 200 nuclear power-plant incidents (see Ref. [13]). The actual number of cases of each type is indicated in parentheses

Omission of functionally isolated acts	34%	(68)
Latent conditions not considered	10%	(20)
Other types of error, unclassifiable	10%	(20)
Other types of omission	9%	(17)
Side effect(s) not considered	8%	(15)
Simple mistakes among alternatives	5%	(11)
Alertness low	5%	(10)
Manual variability (i.e. clumsiness)	5%	(10)
Spatial orientation weak	5%	(10)
Strong expectation	5%	(10)
Familiar association	3%	(6)
Absent-mindedness	1%	(3)

1.3.2.2 Dominant types of human error – related problems

In another study, of 200 "significant events" at nuclear power plants [13], a breakdown of the errors types was made (see Table 1.1). The terms in Table 1.1 are defined as follows:

(a) omission: the failure to perform one or more of the actions needed to accomplish some goal [13],
(b) functionally isolated: isolated from the main purpose of the task (e.g. switching equipment from "test" or "standby" mode to the normal operating mode) [14],
(c) latent conditions: conditions whose adverse consequences do not immediately manifest themselves, but lie dormant within a system (e.g. experimental apparatus) for an extended period (> 1–2 days), until they combine with other factors to create a problem [13],
(d) mistakes among alternatives: for example, mistakes in setting switches, such as "up/down," etc.,
(e) strong expectation: make an assumption, rather than observe the actual situation [14],
(f) familiar association: inadequate application of rules for interpreting phenomena.

Latent errors (which result in latent conditions) often occur when a potential problem is dismissed as being unimportant, because it is not realized how latent conditions can multiply over time and eventually combine to create an actual problem. Some major accidents in the public domain, such as the Challenger space-shuttle accident and the Chernobyl disaster, have been attributed to latent errors [13].

It was found that omissions were the dominant type of error [13]. These could include such things as forgetting to set a valve to the correct position, or omitting some steps in a procedure. Omissions comprised 43% of all errors. Other studies have arrived at similar conclusions. The omission errors in the above survey were most closely associated with

testing, calibration, and maintenance operations. Omissions in general are responsible for an immense amount of wasted time in scientific research [4].

Some facts regarding the chances of making an omission error are as follows [13], [15].

(a) Such errors occur more often if one is carrying out routine tasks while distracted or preoccupied.
(b) The presence of a large number of discrete steps in the action sequence comprising the task increase the chances that at least one will be omitted.
(c) If the amount of information needed to carry out a step is large, the chances are high that items in that step will be omitted.
(d) Steps that are not clearly cued by previous ones, or do not succeed them in a direct linear sequence, stand a good chance of being omitted.
(e) If instructions have been given verbally and there are more than five simple steps, those in the middle of the list are more likely to be omitted than those at the beginning or end.
(f) In the case of written instructions, isolated steps (i.e. those not clearly associated with the others) at or near the end of the list are likely to be omitted.
(g) Steps in an action sequence involving reassembly (e.g. of apparatus) are more likely to be omitted than those of the original disassembly.
(h) If certain steps must be performed on some occasions, but not on others, then these steps stand a higher chance of being omitted. This is especially true if such steps are needed relatively infrequently.
(i) In a highly automatic task that has been well practised, unexpected interruptions are likely to lead to omissions.
(j) If the person who finishes a task is not the same as the one who started it, it is more probable that omissions will occur.

Highly automated, routine tasks are also vulnerable to premature exits (omitting some final steps in the task), especially if there is time pressure or another job waiting to be done.

Omission errors are a regular problem during apparatus construction, maintenance or reassembly activities. Forgetting to install or properly tighten fasteners (bolts, screws, etc.) is very common, particularly if multiple fasteners must be installed [15]. In electronic work, it is not uncommon to forget to apply solder to a connection, if a number of connections that require soldering are present. (This can lead to troublesome intermittent electrical contacts.) Hardware items in general are frequently not connected or left loose, or are missing altogether. Naturally, the risk of making such errors increases if the items involved (e.g. fasteners) are subsequently covered up by other components during further assembly work. Also, the removal of foreign objects and tools from a work area (e.g. the inside of a vacuum chamber or an electronic instrument) at the end of a job is frequently omitted.

1.3.2.3 Dominant causes of human error – related problems

In the study of "significant event reports" at nuclear power plants discussed on page 8, the category of "human performance" problems was broken down by cause of problem (see Table 1.2).