SAFETY-CRITICAL COMPUTER SYSTEMS

Increasingly, microcomputers are being used in applications where their correct operation is vital to ensure the safety of the public and the environment: from anti-lock braking systems in automobiles, to fly-by-wire aircraft, to shut-down systems at nuclear power plants. It is, therefore, vital that engineers are aware of the safety implications of the systems they develop.

This book is an introduction to the field of safety-critical computer systems, and is written for any engineer who uses microcomputers within real-time embedded systems. It assumes no prior knowledge of safety, or of any specific computer hardware or programming language.

This book:

- Covers all phases of the life of a safety-critical system from its conception and specification, through to its certification, installation, service and decommissioning
- Provides information on how to assess the safety implications of projects, and determine the measures necessary to develop systems to meet safety needs
- Gives a thorough grounding in the techniques available to investigate the safety aspects of computer-based systems and the methods that may be used to enhance their dependability
- Uses case studies and worked examples from a wide range of industrial sectors including the nuclear, aircraft, automotive and consumer products industries

Audience:

This text is intended for both engineering and computer science students, and for practising engineers within computer-related industries. The approach taken is equally suited to engineers who consider computers from a hardware, software or systems viewpoint.

Neil Storey is in the Department of Engineering of the University of Warwick, UK, and has published a large number of journal and conference papers in the area of safety-critical systems. He has many years of experience in teaching undergraduate, postgraduate and professional engineers, and is the author of a widely-used electronics textbook.

ISBN 0-201-42787-7

9 780201 427875 >

Addison Wesley Longman Limited

SAFETY-CRITICAL COMPUTER SYSTEMS

Storey

ADDISON WESLEY

42787

SAFETY-CRITICAL COMPUTER SYSTEMS

ADDISON-WESLEY

4 Risk Analysis

CHAPTER CONTENTS

4.1 Introduction

In the previous chapter we noted that hazards are present in all aspects of our everyday life. Associated with these hazards are certain 'risks', which determine their relative importance and enable us to judge their acceptability. The *possibility* of being struck by a meteorite represents a hazard, as does the *possibility* of being stung by a mosquito or electrocuted by a wrongly wired appliance. However, the risks associated with these potential dangers may be very different. A large part of our lives is spent in assessing risks and our survival is closely linked to our abilities at this task.

In order to understand the nature of risk it is useful to consider the relationship between hazards and events that result in harm to an individual or the environment. It should be remembered that hazards represent situations of *potential* danger. When a hazard results in an event that causes actual harm, this occurrence is commonly termed an **accident**. In fact, safety engineers formally define an accident as *an unwanted and unexpected release of energy*, a definition that excludes some harmful occurrences such as exposure to dangerous chemicals. This has led some engineers to use the term **mishap** to include both accidents and other causes of harm (Leveson, 1986). In this text we will adopt the more common use of the term 'accident' that is synonymous with the term 'mishap':

> *An **accident** is an unintended event or sequence of events that causes death injury, environmental or material damage.*

Much of the process of hazard analysis is aimed at identifying sequences of events that may lead to an accident. These trains of events are often termed **accident sequences** or **accident scenarios**. Only by investigating these sequences can the true significance of individual hazards be determined.

A great deal of useful information can also be gained by looking at events that might be considered as 'near misses'. These are often referred to as incidents:

> *An **incident** (or **near miss**) is an unintended event or sequence of events that does not result in loss, but, under different circumstances, has the potential to do so.*

The importance of a hazard is related to the accidents that may result from it. Two factors are of significance here, namely:

- the potential consequences of any accident that might result from the hazard;
- the frequency (or probability) of such an accident occurring.

The risk associated with a hazard is determined by these two factors and may be defined as follows:

> ***Risk** is a combination of the frequency or probability of a specified hazardous event, and its consequence.*

This definition allows risk to be treated in either a qualitative or a quantitative manner. If the latter approach is adopted, numerical estimates of both frequency and severity may be combined to produce a single measure of risk. This is shown in Example 4.1.

Example 4.1

Failure of a particular component is likely to result in an explosion that could kill 100 people. It is estimated that this component will fail once in every 10 000 years. What is the risk associated with this component?

One failure in every 10 000 years represents a failure rate of 0.0001 failures per year. Therefore, the risk is given by:

$$\begin{aligned}\text{Risk} &= \text{severity} \times \text{frequency} \\ &= 100 \times 0.0001 \\ &= 0.01 \text{ deaths per year}\end{aligned}$$

Risk may also be expressed in terms of its effects on individuals. This is of particular relevance when society as a whole is at risk. This might be the case when considering the global implications of environmental hazards, or simply when considering hazards that affect a great many people. This approach is shown in Example 4.2.

Example 4.2

In a country with a population of 50 000 000 approximately 25 people are killed each year by lightning. What is the risk associated with death from this cause?

The fraction of the population killed per year is simply $25/50\,000\,000 = 5 \times 10^{-7}$. The associated risk may be expressed by saying that each individual has a probability of 5×10^{-7} of being killed by lightning in any given year. Alternatively, we could say that the population as a whole is exposed to a risk of 5×10^{-7} deaths per person-year.

In some cases measures of severity and frequency are divided into classes to simplify the use of guidelines and standards. Recommendations, or requirements, are then based on these classes rather than on numeric values for these quantities. The classifications used vary greatly between industries, and the definitions of the various classes are very closely linked to the areas in which they are used. In the following sections we look at some of the classifications used within specific industrial sectors.

4.2 Consequences of malfunction – severity

All safety-related industries classify hazards in terms of their severity. Inevitably, the classifications used tend to be closely related to the nature of the relevant industry, thereby making comparisons difficult. In this section we look at a few examples of the classes used within particular areas.

Avionics

Civil aviation standards within Europe and the US categorize hazard severity as shown in Table 4.1 (RTCA/EUROCAE, 1992). It can be seen that the severity is defined in terms of the effects on the aircraft and its crew, as well as the likely influence on the safety of the occupants. Military avionics standards follow similar lines.

Table 4.1 Hazard severity categories for civil aircraft.

Category	*Definition*
Catastrophic	Failure condition which would prevent continued safe flight and landing
Hazardous	Failure conditions which would reduce the capability of the aircraft or the ability of the crew to cope with adverse operating conditions, to the extent that there would be: (1) a large reduction in safety margins or functional capabilities (2) physical distress or higher workload such that the flight crew could not be relied on to perform their tasks accurately or completely (3) adverse effects on occupants, including serious or potentially fatal injuries to a small number of those occupants
Major	Failure conditions which would reduce the capability of the aircraft or the ability of the crew to cope with adverse operating conditions to the extent that there would be, for example, a significant reduction in safety margins or functional capabilities, a significant increase in crew workload or in conditions impairing crew efficiency, or discomfort to occupants, possibly including injuries
Minor	Failure conditions which would not significantly reduce aircraft safety, and which would involve crew actions that are well within their capabilities. Minor failure conditions may include, for example, a slight reduction in safety margins or functional capabilities, a slight increase in crew workload, such as routine flight plan changes, or some inconvenience to occupants
No effect	Failure conditions which do not affect the operational capability of the aircraft or increase crew workload

Military systems

The UK Ministry of Defence has defined the techniques to be used for hazard analysis and safety classification for military computer-based systems in Interim Defence Standard 00-56 (MoD, 1995). This defines four levels of hazard severity, termed 'accident severity categories', as outlined in Table 4.2. It can be seen that here the definitions are more generally applicable, with the emphasis on death and injuries to individuals.

Proposed international standard – IEC 1508

In 1995 the International Electrotechnical Commission (IEC) issued a draft of a proposed international standard that is intended to form a generic basis for standards in all industrial sectors (IEC, 1995). Although this standard is only in draft form at the time of writing, its importance is such that its contents are

Table 4.2 Accident severity categories for military systems.

Category	*Definition*
Catastrophic	Multiple deaths
Critical	A single death, and/or multiple severe injuries or severe occupational illnesses
Marginal	A single severe injury or occupational illness, and/or multiple minor injuries or minor occupational illnesses
Negligible	At most a single minor injury or minor occupational illness

discussed in several sections of the book. The standard is also discussed in more detail in Chapter 14.

The proposed standard adopts a four-level system for categorizing the severity of hazards using classes with names identical to those given in Table 4.2. The document does not explicitly define the four classes, but it is likely that definitions similar to those given above are appropriate. When generating industry-specific standards, different industries will undoubtedly adopt class definitions to match their needs.

4.3 Probability of malfunction – frequency

The probability, or frequency, of a hazardous event may be expressed in a number of ways, and may be given quantitatively or qualitatively. Sometimes the rate is expressed in terms of the number of events per hour or per year of operation. Alternatively, it may be given as the number of likely events during the lifetime of the unit. In systems that are used intermittently, such as emergency shutdown systems, it is common to express the information in terms of the number of **failures on demand**, that is, the number of failures expressed as a fraction of the total number of times it is called upon to operate. It is perhaps worth noting that it is common to see failure probabilities given as a single probability figure *without units*. This practice can be very confusing, as the significance of the figure is clearly very different if it refers to 'failures per year', 'failures per hour' or 'failures on demand'. Unfortunately, even standards are not immune from this weakness.

Avionics

Within standards for aircraft systems the frequency, or probability, associated with hazards is normally expressed in terms of the number of occurrences that might be expected per hour of flight. Similar definitions are used within Europe (JAR, 1994) and in the US (FAR, 1993). The range of probabilities is divided

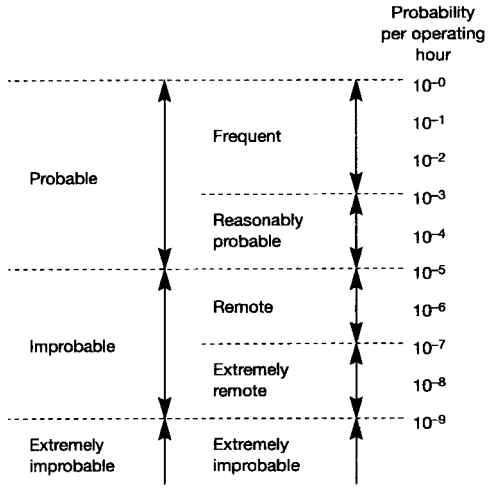

Figure 4.1 Hazard probability classes for aircraft systems.

into three main classifications: probable, improbable and extremely improbable. These three basic ranges are then subdivided into five classes. The definitions of these classes are shown in Figure 4.1.

Military systems

Interim Defence Standard 00-56 defines six categories of hazard probability, as shown in Table 4.3 (MoD, 1995).

Table 4.3 Accident probability ranges for military systems.

Accident frequency	*Occurrences during operational life considering all instances of the system*
Frequent	Likely to be continually experienced
Probable	Likely to occur often
Occasional	Likely to occur several times
Remote	Likely to occur some time
Improbable	Unlikely, but may exceptionally occur
Incredible	Extremely unlikely that the event will occur at all

Proposed international standard – IEC 1508

The draft of the proposed IEC standard adopts a six-level system for classifying the frequency of a hazard, using names for the classes identical to those given in Table 4.3. The classes are not defined within the standard, but it is likely that definitions similar to those used above would be appropriate in many cases. The definitions used for any particular applications are likely to reflect the nature of the hazards associated with the relevant industry.

4.4 Risk classification

We have seen that where quantitative methods are used to describe the severity and frequency of a hazard, it is possible to produce a numerical value for the associated risk by combining these two values. Examples 4.1 and 4.2 demonstrated this process. However, when qualitative measures are used to describe severity or frequency, as, for example, in Tables 4.1 and 4.3, direct calculation of risk is not possible. In such cases severity and frequency data must be combined in some less mechanical way (Figure 4.2). The result of this process is a classification of the risk associated with a particular hazard. This classification is sometimes called a **risk class**, a **risk level** or a **risk factor**. In fact, the use of risk categories is common even where numerical values are used for the severity and frequency of hazards, as their adoption simplifies the use of guidelines. Most standards define a number of risk classes and then set out development and design techniques appropriate for each category of risk. Unfortunately, there is little consensus on the naming or definition of the various classes. In this section we shall look at a few examples.

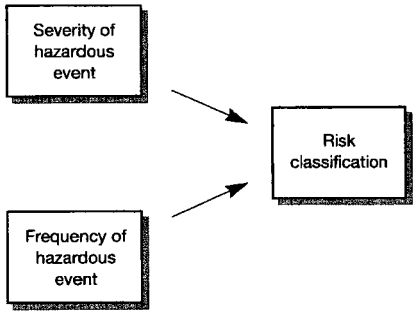

Figure 4.2 Determination of risk classification.

operate in a 'continuous mode'. Failure rates for this group are expressed in **failures per year**, although discussion documents produced during the development of this standard gave rates in failures per hour. The difference between these two sets of units represents a factor of approximately 10^4, and so a knowledge of the units being used is vital. The second class covers systems that operate in what is termed a 'demand mode'. This classification includes applications such as shutdown systems, which are called upon only when needed. Failure rates for this group are expressed in terms of **failures on demand**, referring to the probability that the system will fail to operate when called upon to do so. Discussion documents for the standard referred to the two classes as 'continuous control systems' and 'protection systems', but later versions have adopted more precise definitions. Target failure rates for the four safety integrity levels are given in Table 4.10. It can be seen that the chosen sets of units have the characteristic that the numerical values of the ranges for the two classes are identical. This fact should not be allowed to mask the distinction between these two figures.

In Chapter 1 we noted that the production of safety-critical systems may be considered as a problem of fault management. We also saw that faults may be classified as random, as in the case of random hardware failures, or systematic, as in hardware or software design faults. When determining the safety integrity of a system it is necessary to consider its performance in respect of both these areas. We may therefore identify two elements that contribute to the overall safety integrity of the system:

> ***Hardware integrity*** *is that part of the safety integrity relating to dangerous random hardware failures.*
>
> ***Systematic integrity*** *is that part of the safety integrity relating to dangerous systematic failures.*

Systematic integrity covers all aspects of the design of the system and therefore includes considerations of hardware and software. However, in some cases it is appropriate to look at the software in isolation, leading to a third classification:

> ***Software integrity*** *is that part of the safety integrity relating to dangerous software failures.*

Table 4.10 Target failure rates for the safety integrity levels of draft IEC 1508.

Safety integrity level	*Continuous mode of operation (probability of a dangerous failure per year)*	*Demand mode of operation (probability of failure to perform its designed function on demand)*
4	$\geqslant 10^{-5}$ to $< 10^{-4}$	$\geqslant 10^{-5}$ to $< 10^{-4}$
3	$\geqslant 10^{-4}$ to $< 10^{-3}$	$\geqslant 10^{-4}$ to $< 10^{-3}$
2	$\geqslant 10^{-3}$ to $< 10^{-2}$	$\geqslant 10^{-3}$ to $< 10^{-2}$
1	$\geqslant 10^{-2}$ to $< 10^{-1}$	$\geqslant 10^{-2}$ to $< 10^{-1}$

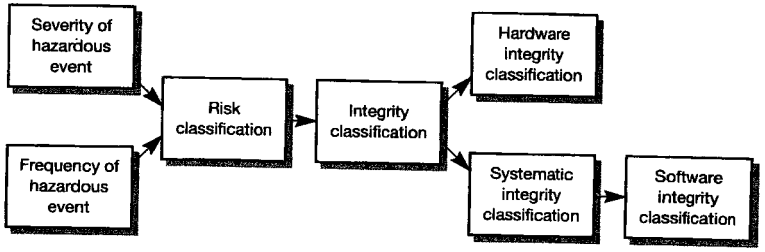

Figure 4.5 Assignment of integrity levels.

The assignment of an integrity level to a system, or part of a system, is based on the classification of the risks associated with it. Following the assignment of an overall system safety integrity level to a system or subsystem, individual integrity levels may be assigned to different aspects of the unit. For example, a software integrity level might be assigned to define the importance of the correctness of the software elements. This in turn will determine the development methods used and the level of testing performed. The process of assigning integrity levels is illustrated in Figure 4.5.

Clearly, the assignment of safety integrity levels is related to the risk class associated with an application. However, it is important to remember the distinction between these two classifications. *Risk* is a measure of the likelihood, and consequences, of a hazardous event. *Safety integrity* is a measure of the likelihood of the safety system correctly performing its tasks.

Allocation of integrity levels

The various standards give useful guidance on the process of assigning integrity levels to systems, or subsystems, that are safety critical. Draft IEC 1508 gives target failure rates for the various integrity levels but does not specify how these should be allocated to specific applications. This is because the IEC standard is intended to be generic in nature, rather than tied to any particular industry. The standard does, however, give guidance on the design and development techniques appropriate for each integrity level.

The process of assigning integrity levels to particular applications is considered within industry-specific standards. Being concerned with only a single application area, these can give detailed information on the relevant classifications for typical system components. This process was illustrated in Table 4.8, which shows how integrity classifications are assigned to various systems within nuclear power stations in Germany. Once an integrity level has been assigned to a system, its design and method of development must ensure that it meets the needs of that level. In many cases the standards define the requirements of the various integrity levels in a manner appropriate to the

particular industry (as in Table 4.9), and set out development methods that must be used for each level.

In some cases standards specify an integrity requirement directly in terms of a maximum probability of failure, rather than in terms of an integrity level. This practice is adopted in civil aircraft standards (FAR, 1993; JAR, 1994), as shown in Table 4.11. Comparison of this table with Table 4.1 and Figure 4.1 shows that the failure requirements represent the specification of an allowable hazard probability class for each of the hazard severity categories. The requirements could therefore be expressed as 'any failure that could result in a catastrophic condition must be extremely improbable; any failure that could result in a hazardous condition must be extremely remote;' and so on.

The requirements set out in Table 4.11 define the performance that a civil aircraft equipment manufacturer must achieve if his products are to be acceptable. From this it can be seen that if failure of a piece of equipment could result in a 'major' incident, it must be expected to fail at most once in every 100 000 hours of service. If, on the other hand, its effects could be 'catastrophic', resulting in loss of the aircraft, it must be expected to fail not more than once in every 1000 million hours. This latter figure represents less than one failure in every 100 000 years of operation. Note that even in this most critical situation the standard does not say that the equipment must not fail. No system can guarantee freedom from failure. The standard simply defines acceptable levels for the dependability of the system.

Table 4.11 Relationship between the severity of an effect and its allowable probability for civil aircraft systems.

Category	*Severity of effect*	*Maximum probability per operating hour*
Normal		10^{0}
		10^{-1}
Nuisance		10^{-2}
Minor	Operating limitation; emergency procedures	10^{-3}
		10^{-4}
Major	Significant reduction in safety margins; difficult for crew to cope with adverse conditions; passenger injuries	10^{-5}
		10^{-6}
Hazardous	Large reductions in safety margins; crew extended because of workload or environmental conditions. Serious injury or death of a small number of occupants	10^{-7}
		10^{-8}
Catastrophic	Multiple deaths, usually with loss of aircraft	10^{-9}

Although the maximum probability figures for civil aircraft are normally expressed in failures per hour of operation, as in Table 4.11, it is interesting to compare the various ranges of probabilities with the integrity levels defined in the draft of the IEC 1508 standard. If allowance is made for the different units used (failures per hour as against failures per year), it can be seen that events that could result in a 'major' incident must have a probability corresponding to the failure of a system of category 1, those resulting in a 'hazardous' incident must have a probability corresponding to category 3, and those that could be 'catastrophic' must fall within the high range of category 4.

Achievable levels of integrity

We have seen that the allocation of an integrity level to a system or subsystem imposes requirements upon it in terms of its dependability. In later chapters we shall look at several methods that may be used to increase the dependability of a system in order to meet these targets.

In developing a safety-critical system it is necessary not only to achieve a high level of integrity, but also to be able to demonstrate that this has been done. Unfortunately, the latter often proves to be difficult, or perhaps even impossible, in highly critical systems. We have seen, for example, that aircraft systems are often required to have a probability of failure corresponding to less than one failure in every 100 000 years of operation. At present we know of no method of testing a system to demonstrate this level of performance. It is therefore surprising that some engineers, such as those responsible for components of the Paris Métro system, have claimed to have achieved failure rates 1000 times better than this value (Guiho, 1990).

As the highest level of integrity of IEC 1508 is itself beyond our current abilities to demonstrate dependability, it can be argued that any system that requires a system of greater performance is simply too dangerous. It is therefore becoming generally accepted that it is inappropriate to specify systems requiring an integrity above this level. We shall return to look at methods of investigating the dependability of a system in Chapter 12.

4.7 The view of society and ethical considerations

Whereas the process of hazard analysis is a logical, mechanistic procedure, the study of risks inevitably involves the use of judgement and opinion. When one assigns an integrity level to a system, or defines a tolerable failure rate, one is defining how much effort should be expended in improving the safety of the arrangement. In so doing, one is placing a value, either monetary or otherwise,

Military systems

Interim Defence Standard 00-56 (MoD, 1995) defines four accident risk classes which are given the symbols A, B, C and D, where A represents the most serious event and D the least serious. Table 4.4 shows the relationships between the various risk classes and the severity and frequency of the hazard. Table 4.5 gives the definitions of the four classes.

Proposed international standard – IEC 1508

The draft of the proposed IEC standard defines four risk classes which are given the symbols I, II, III and IV, where I corresponds to the most serious accident and IV the least serious. The suggested relationship between the risk classes and the severity and frequency of the hazard is shown in Table 4.6. However, the standard states that the actual form of this relationship will be sector dependent and will depend on the definitions used for the various frequency and consequence classes. The definitions of the risk classes are given in Table 4.7.

Table 4.4 Accident risk classes for military systems.

	Consequences			
Frequency	*Catastrophic*	*Critical*	*Marginal*	*Negligible*
Frequent	A	A	A	B
Probable	A	A	B	C
Occasional	A	B	C	C
Remote	B	C	C	D
Improbable	C	C	D	D
Incredible	D	D	D	D

Table 4.5 Interpretation of risk classes for military systems.

Risk class	*Interpretation*
A	Intolerable
B	Undesirable, and will only be accepted when risk reduction is impracticable
C	Tolerable with the endorsement of the Project Safety Review Committee
D	Tolerable with the endorsement of the normal project reviews

Table 4.6 Risk classifications from draft IEC 1508.

	Consequences			
Frequency	*Catastrophic*	*Critical*	*Marginal*	*Negligible*
Frequent	I	I	I	II
Probable	I	I	II	III
Occasional	I	II	III	III
Remote	II	III	III	IV
Improbable	III	III	IV	IV
Incredible	IV	IV	IV	IV

Table 4.7 Interpretation of risk classes from draft IEC 1508.

Risk class	*Interpretation*
I	Intolerable risk
II	Undesirable risk, and tolerable only if risk reduction is impracticable or if the costs are grossly disproportionate to the improvement gained
III	Tolerable risk if the cost of risk reduction would exceed the improvement gained
IV	Negligible risk

It can be seen that the assignment and definition of risk classes in military systems is directly equivalent to that found in the proposed IEC standard. This similarity reflects the great efforts being made to harmonize standards and guidelines throughout the world.

Tables 4.4 and 4.6 clearly illustrate the advantages associated with the use of risk classes. They allow standards and guidelines to deal with a small number of risk classes rather than to attempt to provide guidance for each combination of severity and frequency.

4.5 The acceptability of risk

The definitions of risk classes given in the previous section show that some levels of risk are simply not acceptable. It is clearly not satisfactory to have a hazard that could have catastrophic consequences and that could occur frequently. However, it may be acceptable to have a situation where an accident with negligible consequences could occur frequently, or where a potentially critical or catastrophic accident is improbable or even remote. The acceptability of a given level of risk is determined by the benefits associated with that risk, and by the amount of effort that would be required to reduce it.

The draft of IEC 1508 divides levels of risk into three ranges, as shown in Figure 4.3. The uppermost band of this diagram represents hazards where the risk is so great that it is deemed to be intolerable and cannot be justified on any grounds. In contrast, the lowermost band represents hazards where the risk is so small that it can generally be neglected. Between these two bands lies a third classification where a risk, though not insignificant, *may* be acceptable under certain circumstances. The criterion for acceptance of a particular risk is based on a decision as to whether it is **as low as is reasonably practicable (ALARP)**, bearing in mind the benefits of the system and the costs of any further risk reduction. A risk level satisfying this criterion is termed the **tolerable risk** for the given application. A risk within the ALARP band is never acceptable if it can be easily reduced. Therefore a proposed system that poses even a very small risk may be judged unacceptable if that risk is unjustifiable. Conversely, a system that has a significant risk may satisfy the requirement if it offers sufficient benefits, and if further reduction of the risk is considered impracticable.

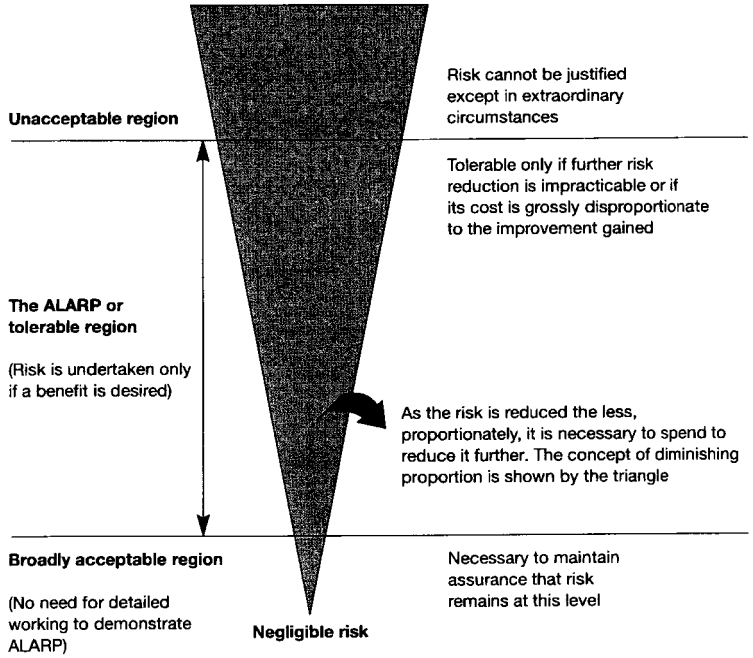

Figure 4.3 Levels of risk from draft IEC 1508.

The ALARP region of Figure 4.3 may be thought of as a zone where a system may be acceptable provided that it is possible to show that the risk is ALARP. Inevitably, a decision as to whether it is practicable to further reduce the risk associated with a particular application is a matter of professional judgement. In highly critical applications, it is the certifying authority that ultimately determines the acceptability of a given system. In reaching their decision they are likely to take into account the standards used within their industry, as well as the arguments put forward by the design team. In applications where the risks are high, those responsible will be expected to expend great effort, and possibly money, in reducing that risk. In lower-risk applications the effort that is expected from the developers is proportionately less. This concept is represented in Figure 4.3 by the width of the triangle.

If Tables 4.5 and 4.7 are considered together with Figure 4.3 it is clear that in each case the most serious risk class corresponds to the region of intolerable risk; the least serious risk class corresponds to the region of negligible risk; and the two remaining classes correspond to the upper and lower bounds of the ALARP region.

The task of producing a safety-critical system can be seen as a process of **risk management** or **risk reduction** (Bell and Reinert, 1993). The equipment under control can be considered to represent a certain level of risk in the absence of any protective measures taken by the system designer. If this level of risk is so low that it may be considered to be negligible, then no further action is required to reduce it further. If this is not the case, safety features must be incorporated within the design to reduce the risk to a level at or below that judged to be tolerable for the application. This process is illustrated in Figure 4.4.

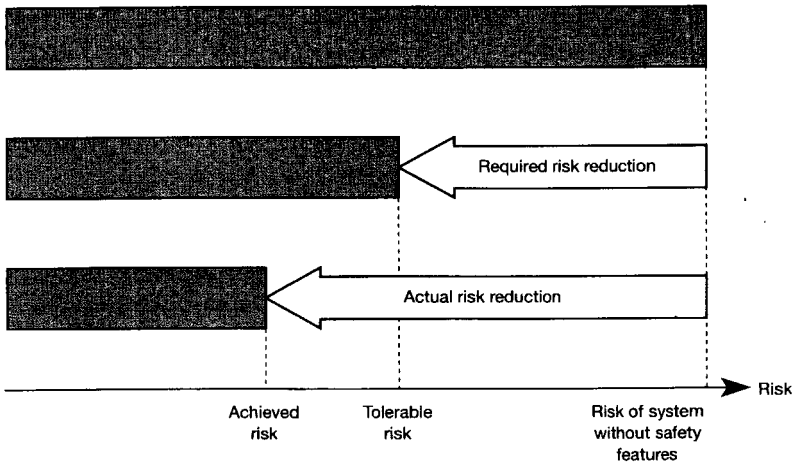

Figure 4.4 The process of risk reduction.

4.6 Levels of integrity

From the discussions in this and earlier chapters, it is clear that the importance of safe operation differs widely between applications. Although safety is always significant, the relative importance of safety in different situations is clearly related to the risks involved. Both an electric toaster and a nuclear reactor protection system should be adequately safe, but the definition of 'adequate' would be different in these two cases.

One could view the differing safety requirements between projects in terms of the level of risk reduction required. In projects associated with potentially high-risk situations a great deal of risk reduction is required in order to achieve a tolerable level of risk. In order to provide a high degree of risk reduction, the risk management mechanisms must themselves be highly dependable. In situations where a relatively low degree of risk reduction is needed, the requirements of the safety systems are less stringent. In such a situation financial considerations will inevitably preclude many techniques that might be appropriate for a more demanding application.

Differing requirements for safety systems lead to the concept of **levels of integrity** for safety-critical systems. In this context the word 'integrity' is concerned with 'safety integrity', which may be defined as follows:

> ***Safety integrity** is the likelihood of a safety-related system satisfactorily performing the required safety functions under all the stated conditions within a stated period of time.*

Although safety integrity could be expressed quantitatively, it is more common to allocate a system to one of a number of safety integrity levels. These in turn may be defined either quantitatively, in terms of measures of performance, or qualitatively, in terms of system characteristics. Table 4.8 shows the proposed integrity classifications for computer systems in nuclear power stations in Germany. The various classes also have associated numerical requirements for such factors as 'failures per year' and 'failure probability during an accident'. These requirements are illustrated in Table 4.9.

Each of the various standards classifies safety-critical systems into a number of integrity levels, but unfortunately there is a great variation in the number and definition of these levels. The number of classifications ranges from a single level to an eight-level system, and the numbering or naming of these levels also varies in form and direction. The international community is gradually converging on a four-level classification, with level 1 being the least critical and level 4 the most critical. This is the convention used in the draft of IEC 1508, and will be used throughout the remainder of this text.

The IEC standard sets out target failure rates for each of the safety integrity levels. In other words, it defines the maximum number of times that a system built to a particular integrity level would be expected to fail in a given period of time. In doing so the standard differentiates between two classes of system that are used in different ways. The first class covers systems that

Table 4.8 Proposed integrity classifications of computer systems in nuclear power stations in Germany.

Class	*Functional requirement*	*Systems involved*
I	Highest requirements of nuclear safety	Systems that release automatic actions for protection of human life and environment
II	High requirements of nuclear safety	Systems that act for the protection of human life and environment by • guiding safety variables under abnormal conditions • causing operator actions deterministically
III	Normal requirements of nuclear safety	Systems that • limit plant variables to specific values • avoid scrams • report disturbances in systems of class I and II
IV	High requirements of plant safety	Systems that • release actions automatically for protection of persons in the plant • protect important parts of the plant
V	High requirements of plant availability	Systems that • increase plant availability • protect normal parts of the plant
VI	High functional requirements	Systems that serve for optimal plant operation, e.g. with respect to efficiency or manœuvrability
VII	Component-related control	Systems for simple requirements

Table 4.9 Proposed requirements for the integrity classes of nuclear systems in Germany.

Class	*I*	*II*	*III*	*IV*	*V*	*VI*	*VII*
Maximum unavailability	10^{-5}	10^{-4}	10^{-2}	10^{-4}	10^{-2}	10^{-1}	10^{-1}
Maximum unavailability for safety-related systems	10^{-7}	10^{-5}	10^{-4}	10^{-5}			
Maximum failure probability per year	10^{-4}	10^{-2}		10^{-2}			
Maximum probability of dangerous failure per year	10^{-7}	10^{-5}	10^{-4}	10^{-5}			
Maximum probability of failure during an accident	10^{-6}	10^{-4}	10^{-3}	10^{-4}	10^{-2}		