SAFETY-CRITICAL COMPUTER SYSTEMS

Storey

ADDISON WESLEY

42787

SAFETY-CRITICAL COMPUTER SYSTEMS

Neil Storey

ADDISON-WESLEY

and therefore, for the period where the failure rate is constant,

$$\text{MTTF} = \int_0^\infty e^{-\lambda t}\,dt = \frac{1}{\lambda} \tag{7.1}$$

This leads to the very simple result that the MTTF is the inverse of the constant failure rate of the system. Thus a system with a constant failure rate of 0.001 failures per hour will have a mean time to failure of 1000 hours. It is important to note that it is *not* reasonable to assume that such a system will operate correctly for 1000 hours. The reliability of a system at a time t is given by

$$R(t) = e^{-\lambda t}$$

and thus at a time $t = 1/\lambda$ (that is, at a time equal to its MTTF), the reliability is

$$R(t) = e^{-\lambda(1/\lambda)} = e^{-1} = 0.37 \tag{7.2}$$

Thus any given system has only a 37% chance of functioning correctly for an amount of time equal to the MTTF, or conversely a 63% chance of failing in this period. However, for a large number of units the MTTF represents the average time for which they would operate before their first failure.

Mean time to repair

The mean time to repair (MTTR) is quite simply the average time taken to repair a system that has failed, and to get it operational. The figure includes the time taken to detect the failure, locate the fault, effect a repair and reconfigure the system. Often the MTTR may be estimated during the design stage to allow system performance to be predicted, but may need to be determined experimentally when the system is operational.

Just as we describe the reliability of a system using its failure rate λ, we can quantify the repairability of a system using its repair rate μ, which is the average time taken to repair the system. It has units of repairs per hour. Just as the MTTF is $1/\lambda$, the MTTR is $1/\mu$.

Mean time between failures

If it can be assumed that once a failed system has been repaired its performance will be equivalent to the original system, then it is possible to predict the mean time between failures (MTBF), which is simply given by

$$\text{MTBF} = \text{MTTF} + \text{MTTR}$$

In most cases the time taken to repair the system will be small compared with the time for which the system operates, so in practice the MTBF will be numerically similar to the MTTF.

Availability

The availability of a system is the probability that the system will be functioning correctly at any given time. In other words, it is the fraction of the time for which it is operational. This can be expressed in terms of previously defined terms as

$$\text{Availability} = \frac{\text{Time system is operational}}{\text{Total time}} = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}}$$

In critical systems the availability will normally be close to unity, and it is sometimes more convenient to describe a system in terms of its **unavailability**, where

$$\text{Unavailability} = 1 - \text{Availability}$$

It was noted in Chapter 2 that high availability is of paramount importance in some applications. Its relevance in terms of safety depends on whether safety can be guaranteed when the system is inoperative. In applications that have failsafe states it may be possible to maintain safety even when the computer-based system is not operating. In such cases availability may not be primary to safety, but will still be of importance to the overall performance of the system.

7.2 Reliability modelling

In the previous section we discussed several ways of describing the reliability of individual components. During the design stage of a project it is essential to be able to predict the final reliability of a complete system containing many parts. In this section we look at the use of reliability modelling to estimate the reliability of complex systems. The two most common methods are the 'combinational modelling' and the 'Markov state modelling' approaches.

Combinational models

Combinational reliability models allow the reliability of a system to be calculated from the reliability of its component parts. The components in question could be subsystems or individual electronic devices.

The model distinguishes between the situation where failure of any one of a number of components will cause system failure, and the case where several components must fail simultaneously to cause a malfunction. These two situations are modelled by the series and parallel models respectively. The symbols within reliability block diagrams are described within the international standard IEC 1078 (IEC, 1991).

Figure 7.2 A series combination of components.

Series systems

Within any module that is not itself fault tolerant, it can be assumed that failure of any of its components may cause a system failure. Such an arrangement is represented in Figure 7.2. Here the components are shown in series, as any failure will prevent input data from correctly reaching the output. It should be noted that this representation is diagrammatic and does not correspond to the physical interconnection of the components. The model simply shows that failure of any of the components implies failure of the complete system.

As failure of any of the components will result in overall failure, the failure rate of a series system is equal to the sum of the failure rates of the individual components. If a system contains N components, and it may be assumed that failures in the various components are independent, then the system's failure rate λ, during its constant failure rate period, is given by

$$\lambda = \lambda_1 + \lambda_2 + \ldots + \lambda_N$$

where λ_i is the constant failure rate of the ith component. This may be rewritten as

$$\lambda = \sum_{i=1}^{N} \lambda_i$$

The reliability of the arrangement may also be expressed in terms of the reliability of the components. If $R_i(t)$ is the reliability of the ith component in the system, then the overall system reliability $R(t)$ is given by the expression

$$R(t) = R_1(t)R_2(t)\ldots R_N(t)$$

which may be written as

$$R(t) = \prod_{i=1}^{N} R_i(t) \qquad (7.3)$$

Example 7.1

A system is composed of 100 components and failure of any component will result in failure of the system. If the failure of the various components is completely independent, and each component has a reliability of 0.999, calculate the overall system reliability.

From Equation 7.3 we know that the reliability of a series arrangement of components is simply the product of their individual reliability. In this case, the reliability is therefore

$$R(t) = 0.999^{100} = 0.905$$

This example illustrates the relationship between component and system reliability. Here we see that the resultant reliability of the overall system is considerably worse than that of the individual components. If we consider the unreliability of the components $Q_c(t)$, we see that

$$Q_c(t) = 1 - R_c(t) = 1 - 0.999 = 0.0001$$

whereas the unreliability of the resultant system is

$$Q(t) = 1 - R(t) = 1 - 0.905 = 0.095$$

This represents an increase in unreliability by a factor of almost 1000.

Example 7.2

A series system containing 100 components is required to have a reliability of at least 0.999. Assuming that each of the components is equally reliable, what minimum reliability would they require to achieve the specified system performance?

If the reliability of the individual components is $R_c(t)$, then from Equation 7.3 we have

$$R(t) = R_c(t)^{100} \geqslant 0.999$$

This gives

$$R_c(t) \geqslant \sqrt[100]{0.999} \geqslant 0.999\,99$$

It can be seen that in order to achieve very high reliability using a series arrangement of many components, the individual components must themselves be extremely reliable.

Parallel systems

In systems that contain redundancy, failure of one component or subsystem need not result in failure of the complete system. Such an arrangement is described as a parallel system and is shown diagrammatically in Figure 7.3. In this arrangement it is assumed that the system will remain operational provided that at least one of the parallel elements is functioning correctly.

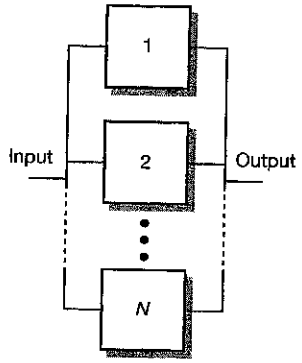

Figure 7.3 A parallel combination of components.

To determine the reliability of a parallel system we start by considering the probability of failure, first of an individual module, and then of the complete system. As the reliability of a component $R(t)$ is the probability of that component functioning correctly for a period of time t, then $[1 - R(t)]$ must be the probability of it failing within that time. You will recall that the quantity $[1 - R(t)]$ is termed the unreliability of the component and is given the symbol $Q(t)$. If a system contains N parallel modules, then the probability of all the units failing independently will be the product of the probabilities of each unit failing individually. Thus, the probability of failure of the system is given by

$$Q(t) = [1 - R_1(t)][1 - R_2(t)] \ldots [1 - R_N(t)]$$

where $R_i(t)$ is the reliability of the ith module. The reliability of the system is therefore

$$R(t) = 1 - Q(t) = 1 - [1 - R_1(t)][1 - R_2(t)] \ldots [1 - R_N(t)]$$

or simply

$$R(t) = 1 - \prod_{i=1}^{N} [1 - R_i(t)] \tag{7.4}$$

If, as is often the case, the parallel modules are identical, each with a reliability of $R_m(t)$, this expression may be simplified. The system reliability then becomes

$$R(t) = 1 - [1 - R_m(t)]^N \tag{7.5}$$

Example 7.3

A system consists of three identical modules and will operate correctly provided that at least one module is operational. If the reliability of each of the modules is 0.999, what will be the reliability of the complete system, assuming that the modules fail independently?

From Equation 7.5 the reliability is

$$\begin{aligned} R(t) &= 1 - [1 - R_m(t)]^N \\ &= 1 - [1 - 0.999]^3 \\ &= 0.999\,999\,999 \end{aligned}$$

When dealing with systems that have a reliability that is very close to unity it is often more convenient to use a measure of unreliability. In this case the unreliability $Q(t)$ is $1 - 0.999\,999\,999$, or 1.0×10^{-9}. It can be seen that a parallel combination of reliable modules produces a system of very high reliability, even when the number of modules is low.

Example 7.4

A system requires a minimum reliability of 0.999. A module designed to fulfil the requirements of the system is found to have a reliability of only 0.85. If a parallel combination of these modules is used to implement the system, what is the minimum number of modules needed to achieve the required reliability?

From Equation 7.5 we have

$$R(t) = 1 - [1 - R_m(t)]^N$$

and therefore

$$R(t) = 1 - [1 - 0.85]^N \geqslant 0.999$$

Rearranging gives

$$\begin{aligned} 0.001 &\geqslant [1 - 0.85]^N \\ N &\geqslant 3.64 \end{aligned}$$

Because N must take an integer value, the minimum number of modules required is four.

Example 7.5

A system is required to have a minimum reliability of 0.999 and is to be constructed using a parallel combination of modules each having a reliability of 0.65. What is the minimum number of modules required?

Repeating the analysis of the previous example we have

$$R(t) = 1 - [1 - 0.65]^N \geqslant 0.999$$

This gives

$$0.001 \geqslant [1 - 0.65]^N$$
$$N \geqslant 6.58$$

Because N must take an integer value, the minimum number of modules required is seven. It can be seen that producing high reliability using unreliable components implies large numbers of modules. In practice such an arrangement is rarely attractive, as the added complexity leads to a high implementation cost and a low MTTF.

Series–parallel combinations

In practice, real systems are often more complicated than the simple series and parallel combinations described above. However, all systems may be reduced to some combination of these two forms, which can then be reduced systematically to produce a single equivalent element. This process is illustrated in Figure 7.4. Figure 7.4(a) shows a system consisting of a series combination of parallel modules. The reliability of this arrangement can be calculated by first combining the parallel elements into a single module of equivalent reliability and then combining the resulting series elements into a single module. Figure 7.4(b) shows the first of these operations. Here module 10 has an equivalent reliability to the parallel combination of modules 1, 2 and 3. Similarly, modules 11 and 12 represent the effective reliability of the parallel combinations of modules 4, 5 and 6, and 7, 8 and 9 respectively. The overall reliability of the system is represented by that of module 13. This is determined by evaluating the series combination of modules 10, 11 and 12.

Parallel combinations of series elements may be analysed in a similar manner by first combining the series modules and then combining resulting parallel elements. Using these techniques, any combination of series and parallel units may be simplified and its reliability assessed.

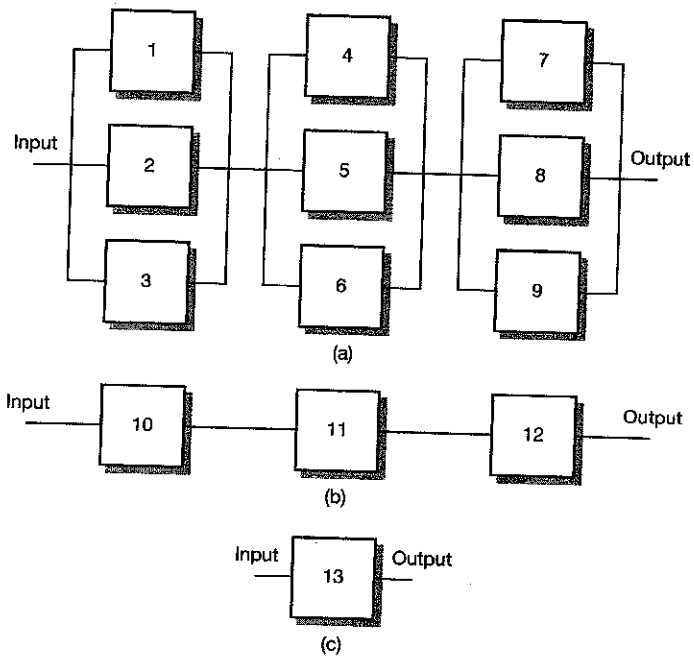

Figure 7.4 A series–parallel combination arrangement: (a) the original arrangement; (b) the result of combining the parallel modules; (c) the effect of combining the series elements.

Example 7.6

A system may be described by the reliability model shown at the top of the next page. Calculate the reliability of the system, given that the modules have the following values for their reliability:

Module 1	*0.99*
Modules 2, 3 and 4	*0.80*
Modules 5 and 6	*0.90*
Modules 7 and 8	*0.95*
Module 9	*0.94*

The reliability of the parallel combination of modules 2, 3 and 4 is given by

$$R(t) = 1 - [1 - R_m(t)]^N$$
$$= 1 - [1 - 0.8]^3$$
$$= 0.992$$

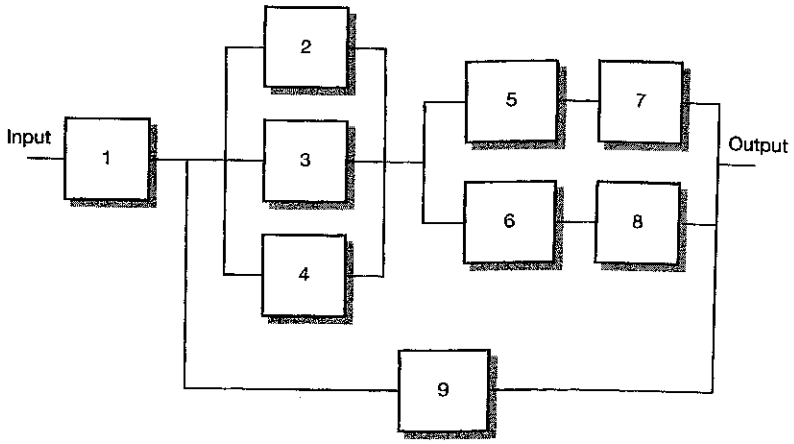

The series combinations of modules 5 and 7 and modules 6 and 8 each have a reliability given by the products of the reliability of the individual modules. Thus each pair has a reliability of 0.90 × 0.95, or 0.855. The parallel combination of these two paths is simply

$$\begin{aligned} R(t) &= 1 - [1 - R_m(t)]^N \\ &= 1 - [1 - 0.855]^2 \\ &= 0.979 \end{aligned}$$

The system may therefore be represented by the simpler form:

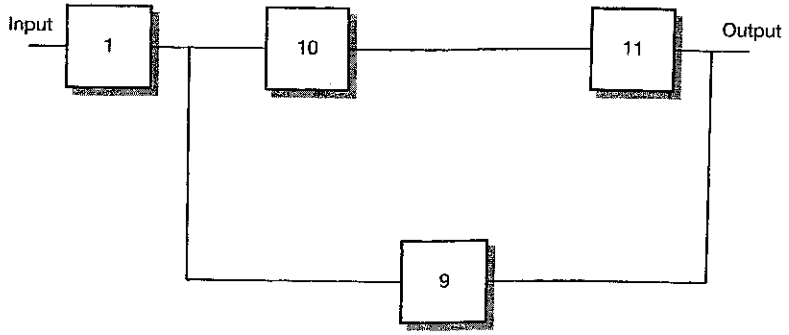

where module 10 represents the combination of modules 2, 3 and 4 and has a reliability of 0.992, and module 11 represents the combination of modules 5, 6, 7 and 8 and has a reliability of 0.979.

The series combination of modules 10 and 11 has a reliability of 0.992 × 0.979, or 0.971. This combination in parallel with module 9 produces an arrangement with a reliability given by

$$\begin{aligned} R(t) &= 1 - [1 - R_1(t)][1 - R_2(t)] \\ &= 1 - [1 - 0.971][1 - 0.94] \\ &= 0.998 \end{aligned}$$

The reliability of the complete system is now the product of this figure and the reliability of module 1. Thus the system reliability is 0.998 × 0.99, that is 0.988, or approximately 0.99.

Triple modular and N*-modular redundancy*

In Chapter 6 we looked at several forms of static redundancy that may be used to achieve fault tolerance. The voting mechanisms used within these arrangements may be represented within reliability block diagrams, as shown in Figure 7.5.

The simplest of the static configurations, triple modular redundancy (TMR), uses three parallel modules and will function correctly provided that at least two modules are operational. If we ignore the effects of the voting mechanism, the probability of the system working correctly may be expressed in words as follows:

Probability of correct operation = Probability of no failures
+ Probability of only module 1 failing
+ Probability of only module 2 failing
+ Probability of only module 3 failing

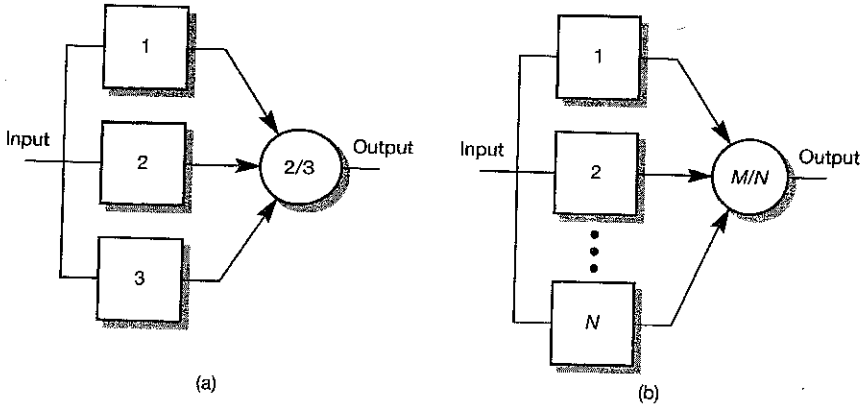

Figure 7.5 Representation of voting mechanisms within reliability block diagrams: (a) triple modular redundancy; (b) *N*-modular redundancy.

If the probability of a module working correctly is $R_m(t)$, then the probability of it failing is $[1 - R_m(t)]$, and the system reliability is given by

$$R_{TMR}(t) = R_1(t)R_2(t)R_3(t) + [1 - R_1(t)]R_2(t)R_3(t) + R_1(t)[1 - R_2(t)]R_3(t) + R_1(t)R_2(t)[1 - R_3(t)]$$

where $R_1(t)$, $R_2(t)$ and $R_3(t)$ represent the reliability of modules 1, 2 and 3 respectively.

If, as is often the case, the reliability of the three modules is identical, this simplifies to

$$R_{TMR}(t) = R_m{}^3(t) + 3R_m{}^2(t)[1 - R_m(t)] = 3R_m{}^2(t) - 2R_m{}^3(t) \quad (7.6)$$

Example 7.7

A TMR system consists of three identical modules each with a reliability of 0.95. Calculate the reliability of the resultant system, ignoring the effects of the voting arrangement.

From Equation 7.6 we have

$$\begin{aligned} R_{TMR}(t) &= 3R_m{}^2(t) - 2R_m{}^3(t) \\ &= 3(0.95)^2 - 2(0.95)^3 \\ &= 0.993 \end{aligned}$$

Example 7.8

Repeat the calculations of the previous example using modules with a reliability of 0.4.

As before,

$$\begin{aligned} R_{TMR}(t) &= 3R_m{}^2(t) - 2R_m{}^3(t) \\ &= 3(0.4)^2 - 2(0.4)^3 \\ &= 0.352 \end{aligned}$$

Note that in this case the reliability of the TMR arrangement is lower than that of an individual module. This illustrates that the use of redundancy does not necessarily increase the reliability of a system. It also shows that a system may be fault tolerant without being reliable. In a TMR arrangement the resultant reliability is only greater than that of the modules themselves if the module reliability is greater than 0.5. Reliability cannot be produced by combining unreliable modules.

A similar analysis to that given above can be applied to systems with N identical modules in which M modules must function correctly in order to prevent a system failure. This results in an expression for the reliability of the form

$$R_{M\text{-of-}N}(t) = \sum_{i=0}^{N-M} \left(\frac{N!}{(N-i)!i!} \right) R_m{}^{N-i}(t)[1 - R_m(t)]^i \quad (7.7)$$

Substituting values of 3 and 2 for N and M respectively yields a result equivalent to that derived above for the TMR arrangement.

Example 7.9

An N-modular redundant system consists of five identical modules each with a reliability of 0.95. Calculate the reliability of the resultant system, ignoring the effects of the voting arrangement.

From Equation 7.7 we have

$$\begin{aligned} R_{3\text{-of-}5}(t) &= \sum_{i=0}^{2} \left(\frac{5!}{(5-i)!i!} \right) R_m{}^{5-i}(t)[1 - R_m(t)]^i \\ &= R_m{}^5(t) + 5R_m{}^4(t)[1 - R_m(t)] + 10R_m{}^3(t)[1 - R_m(t)]^2 \\ &= 10R_m{}^3(t) - 15R_m{}^4(t) + 6R_m{}^5(t) \\ &= 10(0.95)^3 - 15(0.95)^4 + 6(0.95)^5 \\ &= 0.9988 \end{aligned}$$

So far in our discussions of static redundant systems we have ignored the effects of the voting mechanism. As the voter is usually a very simple circuit (as described in Chapter 6) it is possible that its reliability may be sufficiently high to allow a single, non-redundant voter to be used. Such a system may be modelled as a series combination of the M-of-N arrangement described earlier, with a single element representing the reliability of the voter. In more critical applications a fault-tolerant voting arrangement will be needed. Such a voting arrangement will then represent a parallel combination of elements in series with the M-of-N network. This series–parallel combination can then be reduced and analysed as discussed earlier.

Dynamic redundancy and fault coverage

In Chapter 6 we looked at various forms of dynamic redundancy. These use some form of fault detection to switch between a number of redundant modules. This arrangement may be modelled as a parallel combination of elements as described above, in which correct operation will be maintained provided that at least one of the modules functions correctly. This may be represented within a reliability block diagram, as shown in Figure 7.6.

Dynamic systems may function in many ways. In a standby spare configuration a primary module is used exclusively unless fault detection circuitry determines that there is a problem with this unit, when control will switch to a standby module. Thus in the arrangement of Figure 7.6, module 1 would be used exclusively unless a fault were detected, when control would switch to module 2. The success of such an arrangement is critically dependent on the effectiveness of the fault detection circuitry.

The system will function correctly if module 1 is fault-free. Alternatively, it will function correctly if module 1 fails provided that the fault circuitry detects this fault and module 2 is fault-free. This may be expressed in probability terms as

$$R(t) = R_1(t) + [1 - R_1(t)]C_1R_2(t)$$

where $R(t)$ is the system reliability, $R_1(t)$ and $R_2(t)$ represent the reliability of module 1 and module 2 respectively, and C_1 is the fault coverage of module 1. The fault coverage represents the probability of a fault within the module being detected. For modules with identical reliability $R_m(t)$ and fault coverage C_m, this becomes

$$R(t) = R_m(t) + [1 - R_m(t)]C_mR_m(t) \tag{7.8}$$

If we assume perfect fault coverage, that is $C_m = 1$, this reduces to

$$R(t) = 1 - [1 - R_m(t)]^2$$

which is equivalent to the expression obtained earlier in Equation 7.5.

It is interesting to note that if we assume a fault coverage of 0, the expression of Equation 7.8 reduces to $R_m(t)$ and the reliability of the system

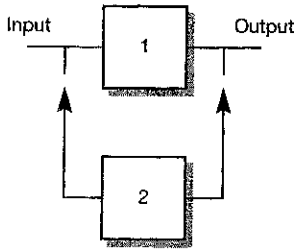

Figure 7.6 A simple dynamic redundant system.

becomes simply the reliability of the primary module. Clearly nothing is gained by having redundant modules if faults are never detected to bring them into service.

Alternative dynamic configurations, such as the use of self-checking pairs, use different methods of fault detection, resulting in slightly different reliability models. Hybrid systems combine the features of static and dynamic systems and again result in variations on the models given above. However, the basic techniques outlined in this section may be used to model all the various fault-tolerant architectures.

Cut and tie sets

Networks that consist entirely of active elements arranged without feedback may be analysed using cut and tie sets (Bansal *et al.*, 1982). These methods do not give an exact value for the reliability of the system, but give upper and lower bounds for its value. This is often useful when dealing with complex arrangements where exact calculation would be difficult.

Cut sets are formed by drawing lines through the reliability block diagram to represent combinations of elements in which simultaneous failure would lead to system failure. Of particular interest are **minimal cut sets**, which represent cut sets in which no subset will result in system failure. Examples of minimal cut sets are shown in Figure 7.7(a). Here it can be seen that failure of component 1 will result in a system failure, whereas failure of component 2 must be accompanied by failure of either 3 or 4 to affect the overall system.

The overall reliability of the complete system may be approximated by considering the influence of each of the minimal cut sets separately. Each cut set

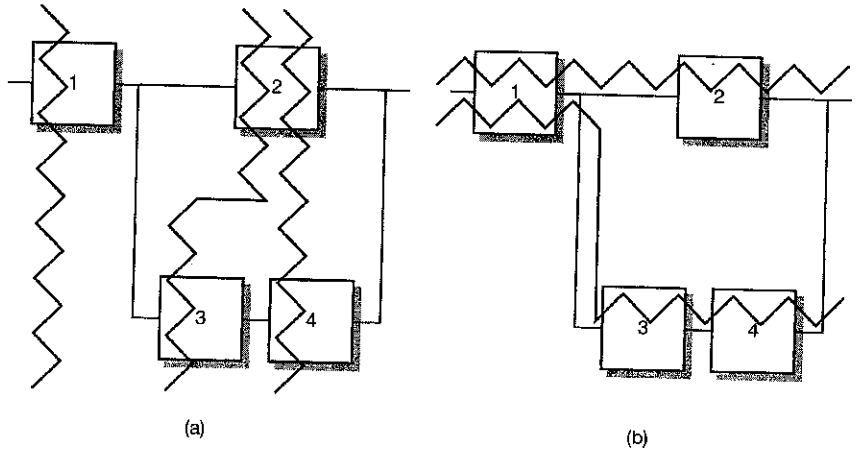

Figure 7.7 Minimal cut and tie sets: (a) minimal cut sets; (b) minimal tie sets.

represents a parallel combination of components, and from Equation 7.4 we know that the reliability of this combination is given by

$$R(t) = 1 - Q(t) = 1 - [1 - R_1(t)][1 - R_2(t)] \ldots [1 - R_N(t)]$$

or simply

$$R(t) = 1 - \prod_{i=1}^{N} [1 - R_i(t)]$$

By combining the influences of all the minimal cut sets it can be shown that the overall reliability must be greater than a value given by the expression

$$R(t) > 1 - \sum_{j=1}^{N_C} \prod_{i=1}^{n_j} [1 - R_i(t)] \qquad (7.9)$$

where N_C is the number of minimal cut sets and n_j is the number of elements in the jth cut set.

Tie sets (or path sets) are formed by drawing lines through the reliability block diagram to represent groups which, if all the elements were working, would guarantee the functioning of the system. **Minimal tie sets** represent tie sets in which no subset will perform this function. Figure 7.7(b) illustrates the concept of minimal tie sets.

By combining the effects of all the minimal tie sets it is possible to place a maximum value on system reliability. This is given by the expression

$$R(t) < \sum_{j=1}^{N_T} \prod_{i=1}^{n_j} R_i(t) \qquad (7.10)$$

where N_T is the number of minimal tie sets and n_j is the number of elements in the jth tie set.

By combining the use of cut sets and tie sets it is possible to place bounds on the possible value of the reliability of a system.

Example 7.10

Determine upper and lower bounds to the reliability of the arrangement of Figure 7.7, if the reliability of the individual modules is as follows:

Module 1	*0.8*
Module 2	*0.7*
Modules 3 and 4	*0.9*

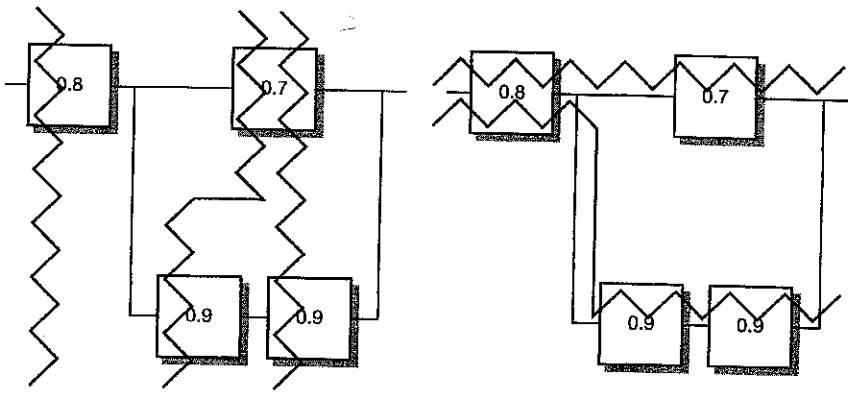

Minimal cut sets

Minimal tie sets

From Equation 7.9 we have

$$\begin{aligned} R(t) &> 1 - [1 - R_1(t)] - [1 - R_2(t)][1 - R_3(t)] - [1 - R_2(t)][1 - R_4(t)] \\ &> 1 - [1 - 0.8] - [1 - 0.7][1 - 0.9] - [1 - 0.7][1 - 0.9] \\ &> 0.74 \end{aligned}$$

and from Equation 7.10 we have

$$\begin{aligned} R(t) &< [R_1(t)][R_2(t)] + [R_1(t)][R_3(t)][R_4(t)] \\ &< [0.8][0.7] + [0.8][0.9][0.9] \\ &< 1.208 \text{ (that is, } < 1.0) \end{aligned}$$

In this case the actual value of the reliability is given by

$$\begin{aligned} R(t) &= R_1(t)\big(1 - [1 - R_2(t)][1 - R_3(t)R_4(t)]\big) \\ &= 0.8\big(1 - [1 - 0.7][1 - 0.9 \times 0.9]\big) \\ &= 0.7544 \end{aligned}$$

Cut and tie set analyses would not normally be used for arrangements as simple as that in the above example. However, in complex systems the calculation of an exact value for the reliability may be impractical. In such cases, an estimate based on cut and tie sets may represent a sufficient and efficient method of assessing reliability. As the complexity of the arrangement increases, the values obtained from the two approaches tend to converge, this convergence being more rapid when the reliability of the individual blocks is high. In practice, analysis often uses only a cut set approach to obtain a lower limit on system reliability.

Although cut and tie set methods are usually simpler than obtaining an exact value for system reliability, they may themselves involve a considerable

amount of calculation. A system of moderate complexity may result in a reliability block diagram containing several millions of minimal cut and tie sets. Such complexity inevitably requires the use of computer-based tools to perform the analyses. The estimations of reliability produced may be improved by systematic removal of the multiple counting of elements. This is a process ideally suited to automation, making this approach particularly suitable for computer-based tools.

Minimal cut sets that contain only a single element represent potential sites for **single-point failure** of the system. Such sets are referred to as 'first-order' cut sets. If the failure of two components can result in system failure this will be represented by a cut set containing two elements – a second-order cut set. Analysis of the system can be simplified by considering only certain low-order cut sets. This simplification is based on the assumption that simultaneous failure of a large number of components is unlikely and can therefore be neglected. The validity of this assumption will depend on the reliability of the components concerned.

In addition to their use with reliability block diagrams, cut and tie sets may also be used in the analysis of **fault trees**. Here minimal cut sets represent combinations of events that can result in the top event of the tree.

Reliability and MTTF

In this section we have often referred to a numerical value for the reliability of a module and used this figure to calculate system reliability. It is important to remember that reliability is a function of time, and that these figures depend on the time for which the system must operate. Different applications place varying constraints on the system designer in terms of reliability and length of service. In some military situations mission duration might be measured in hours, whereas in other situations equipment might have to operate dependably for several decades. Thus one designer might be looking for a reliability of 0.999 for a period of only 24 hours, whereas another might require a reliability of 0.99 for 10 years.

By contrast, the mean time to failure (MTTF) of a system is a fixed characteristic that does not change with time. It might seem logical that reliability and a high MTTF would go hand in hand, but this is not necessarily true. Adding redundancy to a system may increase its reliability over a given time period, but by increasing its complexity it may also reduce its MTTF. As the period of service of a module approaches its MTTF its reliability falls. We noted earlier in the case of a TMR system that adding redundancy only increases the system's reliability if the module reliability is greater than 0.5. In Section 7.1 we saw, in Equation 7.2, that the reliability of a system at a time equal to its MTTF is only about 0.37. Therefore, when the period of service of a unit is comparable with its MTTF, adding redundancy will not increase the overall reliability.

In highly critical systems we normally wish to achieve high levels of reliability. This is usually achieved by using redundant, fault-tolerant designs

with modules that are operated for periods which are short compared with their MTTF.

Independence of failures

In the analysis of series and parallel systems given above it has been assumed that all failures are independent. This assumption is normally valid in the case of random component failures, but is not so for systematic faults. Consider, for example, the case of a parallel system consisting of three identical modules each containing a software fault. In this situation, because each module receives the same input data, it is likely that all the modules would fail simultaneously, thereby removing any benefit from the redundancy. Design faults of other kinds are also likely to produce correlated faults in different modules, resulting in common failures. Similarly, intermittent faults may be caused by interference or other transient events that affect more than one module, leading to simultaneous failures.

Because faults of these kinds produce correlated errors in a number of modules, the assumptions made in the analysis within this section are invalid. For these reasons the combinational modelling techniques described above are frequently restricted to the analysis of random component failures.

Markov models

The combinational modelling techniques described above determine the overall reliability of a system by using measured or predicted values for the reliability of its constituent parts. An alternative approach is to assign various states to a system and to determine the probability of being in any of these states. This is termed Markov modelling (Lewis, 1996). As an example, one might assign two possible states to a system, representing the working and not working conditions. The probability of being in either state would then indicate the availability of the system. One of the advantages of this approach is that it provides a more powerful way of modelling systems that are repairable, allowing variables such as the time taken to repair a system to be incorporated. A detailed treatment of Markov modelling is beyond the scope of this text. However, it is instructive to consider a simple example.

Discrete Markov modelling

Consider a simple two-state system as shown in Figure 7.8. In this system the two states are assigned the designations 1 and 2, and the model assumes that the probabilities of leaving or remaining in a particular state are constant for all time, at the values indicated in the diagram. Transitions between states occur in discrete steps, and thus this is termed a discrete Markov model of the system. The