AS

ital to
mobiles, to
engineers are

any
knowledge of

ion, through

the

of computer-

the nuclear,

engineers
nsider

ublished a
man ears

7

5 >

SAFETY-CRITICAL COMPUTER SYSTEMS

Storey

ADDISON
WESLEY

42787

SAFETY-CRITICA COMPUTER SYSTEMS

Neil Storey

ADDISON-WESLEY

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

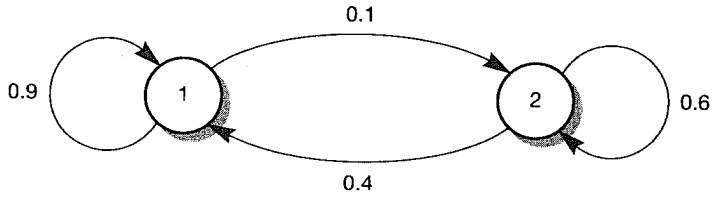

Figure 7.8 A two-state system.

two states could represent any aspects of the system and could, for example, represent working and non-working states.

If we assume that the system is initially in state 1, the diagram shows that at the end of the first time interval it has a probability of 0.9 of remaining in state 1 and a probability of 0.1 of leaving state 1 and entering state 2. Note that the sum of these probabilities is unity, as the system must follow one of these courses of action. Therefore, at the beginning of the second time interval it has a probability of 0.9 of being in state 1 and of 0.1 of being in state 2. At the end of the second time interval the probabilities of leaving or remaining in its current state are again defined by the diagram, and this process continues for successive time steps. The possible sequences of transitions taken by the system, together with the probabilities of following each route, can be represented in a tree diagram, as shown in Figure 7.9.

It can be seen that from the end of the second time interval onwards there are several ways of ending up in either state. The probability of being in each state is therefore the sum of the probabilities of each route leading to that state. If we sum these probabilities at the end of each time interval we quickly see a pattern emerging, as shown in Table 7.1.

Clearly, at the end of each time interval the sum of the probabilities of being in each state must equal unity, as the system must be in one or other state. However, as time progresses the probabilities of each state tend to particular values, depending on the transition probabilities. In this example the probability of being in state 1 tends towards a value of 0.8, and that of being in state 2 tends to a value of 0.2. Within just a few transitions the state probabilities are very close to their limiting values.

The state probabilities shown in Table 7.1 are determined by the transition probabilities between the states, and also by the initial conditions

Table 7.1 Successive state probabilities for the two-state system.

	Time interval				
	1	*2*	*3*	*4*	*5*
State 1	0.9	0.85	0.825	0.8125	0.806 25
State 2	0.1	0.15	0.175	0.1875	0.193 75

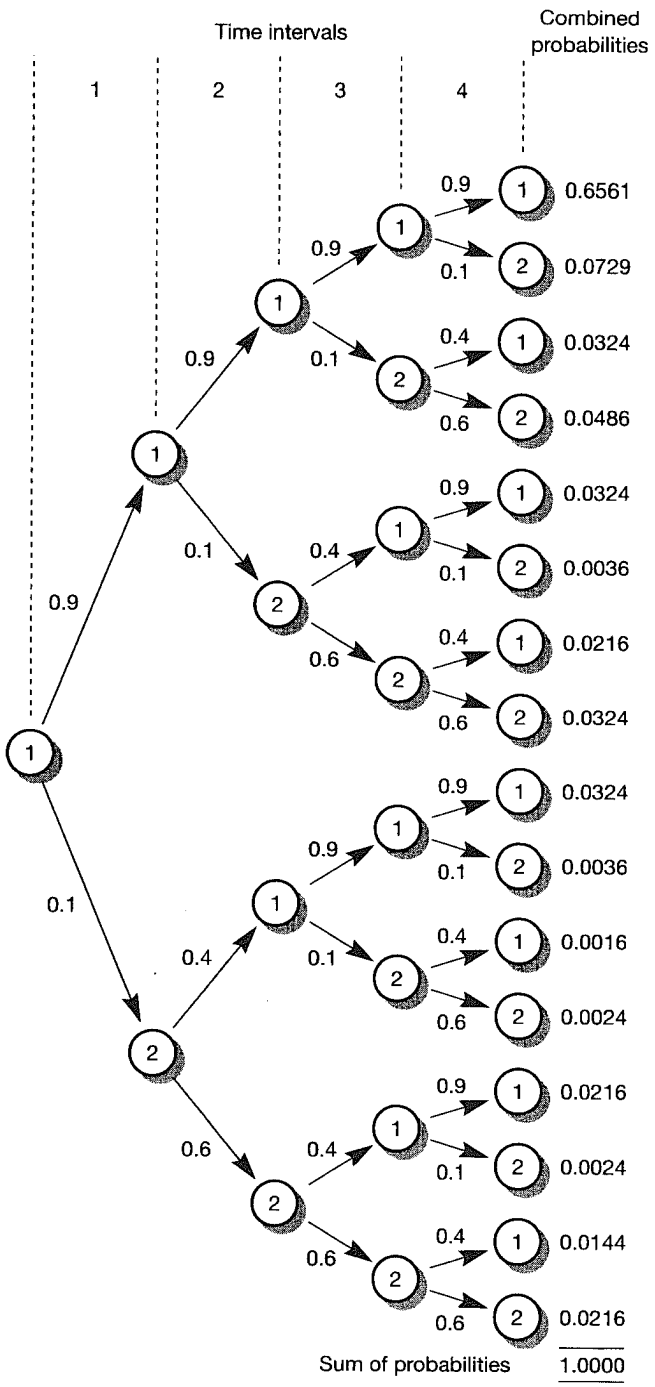

Figure 7.9 A tree diagram of the two-state system.

of the system. In this example the system started from state 1; the initial state probabilities would be markedly different had it started from the other state. However, a very important characteristic of this model is that the limiting values of the state probabilities are independent of the initial conditions, as these have only a transient effect. In other words, the long-term characteristics of the system are not affected by its initial conditions.

It was suggested earlier that the two states of our simple system could represent the working and failed states. In this case the probability of being in either state is clearly related to the reliability of the system and to its availability.

Continuous Markov modelling

In many cases it is more sensible to consider a system in a continuous time domain rather than as a series of discrete time intervals. This can be done using continuous Markov modelling, where the probabilities of state transitions are replaced by transition rates. Let us again consider our simple two-state system, where one state corresponds to the system working correctly and the other to its having failed. Here the rates of transition between the two states represent the failure rate λ and the repair rate μ, as defined in Section 7.1. The resulting model is shown in Figure 7.10.

Using methods similar to those given above for discrete Markov models, it can be shown that the limiting probabilities of being in each state are given by

$$P_1 = \frac{\mu}{\lambda + \mu}$$

and

$$P_2 = \frac{\lambda}{\lambda + \mu}$$

You might like to compare these results with the limiting values obtained for the discrete Markov model obtained earlier.

As the MTTF of the system is $1/\lambda$ and the MTTR is $1/\mu$, P_1 represents the availability of the system; P_2 its unavailability.

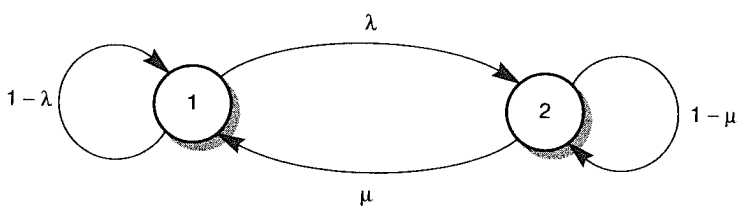

Figure 7.10 A continuous Markov model of the two-state system.