

PCS - 5733

PASTANº:	51
QIDE.FIS.:	56

35

Design and Analysis of Fault-Tolerant Digital Systems

Barry W. Johnson

University of Virginia, Charlottesville

John B. Mitchell



Addison-Wesley Publishing Company

Reading, Massachusetts • Menlo Park, California • New York
Don Mills, Ontario • Wokingham, England • Amsterdam • Bonn
Sydney • Singapore • Tokyo • Madrid • San Juan




Fig. 4.15 Comparison of the reliability of a TMR system composed of three identical modules with the reliability of a single module.

The two solutions to the quadratic equation are 0.5 and 1.0, which implies that the reliability of the TMR system is equal to that of the corresponding nonredundant system when the reliability of the single module is 0.5 or the module is perfect ($R = 1$).

This further illustrates a point that we made when we defined fault tolerance and reliability. A system can be tolerant of faults and still have a low reliability. For example, a TMR system constructed from modules that have individual reliabilities of 0.5 can tolerate a fault in one of those modules, but the reliability of the TMR system is the same as the reliability of a single module. Conversely, a system can achieve a high reliability without being fault tolerant. Certainly, a system that consists of a perfect module will have the highest possible reliability but will not possess, or need, the at-

tribute of fault tolerance. This is, of course, an unrealistic example, but, in general, as the reliability of the components of a system increases, the reliability of the system also increases. It is possible for the reliability of a nonredundant system to approach that of a redundant system constructed from the same modules. The nonredundant system, however, will not be fault tolerant.

In many cases, we have systems that are of the M -of- N structure, but are not TMR; the general NMR system is a good example. In general, if there are N identical modules and M of those are required for the system to function properly, the system can tolerate $N - M$ module failures. The expression for the reliability of an M -of- N system can be written as ..

$$R_{M\text{-of-}N}(t) = \sum_{i=0}^{N-M} \binom{N}{i} R^{N-i}(t) (1 - R(t))^i$$

where

$$\binom{N}{i} = \frac{N!}{(N - i)! i!}$$

For example, the TMR system reliability is given by

$$R_{TMR}(t) = \sum_{i=0}^1 \binom{3}{i} R^{3-i}(t) (1 - R(t))^i$$

which reduces to

$$R_{TMR}(t) = 3R^2(t) - 2R^3(t)$$

which is identical to the expression derived earlier.

4.3.4 Markov Models

The primary difficulty with the combinatorial models is that many complex systems cannot be modeled easily in a combinatorial fashion. The reliability block diagrams can be extremely difficult to construct, and the resulting reliability expressions are often very complex. In addition, the fault coverage that we have seen to be extremely important in the reliability of a system is sometimes difficult to incorporate into the reliability expression in a combinatorial model. Finally, the process of repair that occurs in many systems is very difficult to model in a combinatorial fashion. For these reasons, we often use **Markov models**.

The purpose of the presentation in this text is not to delve into the mathematical details of Markov models but to understand how to use Markov models. For more explicit mathematical details, refer to the references ([Shooman 1968] and [Trivedi 1982]). The discussions here will provide sufficient mathematical background to apply the Markov model but will not pursue various techniques for solving the models.

The two main concepts in the Markov model are the **system state** and the **state transition**. The state of a system represents all that must be known to describe the system at any given instant of time. For reliability models, each state of the Markov model represents a distinct combination of faulty and fault-free modules. For example, suppose we have a TMR system with three identical computers in a majority voting arrangement with a perfect voter. We can define the state of this system as $S = (S_1, S_2, S_3)$ where $S_i = 1$ if module i is fault free and $S_i = 0$ if module i is faulty. The TMR system has eight distinct states in which it can operate: (000), (001), (010), (011), (100), (101), (110), and (111). Each state represents a unique combination of faulty and fault-free modules within the system. For TMR, we know that at least two of the modules must be fault free for the system to operate correctly. Therefore, the states (000), (001), (010), and (100) represent states in which the system has ceased to function correctly. The remaining states are those in which the system is functioning correctly.

The state transitions govern the changes of state that occur within a system. As time passes and failures and reconfigurations occur, the system goes from one state to another. For example, if the TMR system starts its operation in state (111) and at some time t module 1 fails, the system transitions to state (011). The state transitions are characterized by probabilities such as the probability of failure, fault coverage, and the probability of repair.

As an example of the state transitions that can occur, consider the TMR system. We have already defined the states that can exist in the system; now let us define the transitions that can occur. We construct our transitions using several assumptions. First, we assume that the system does not contain repair. In other words, once a module has failed, it remains failed permanently. Second, we assume that only one failure will occur at a time. In a TMR system, the single failure assumption implies that the system cannot go directly from the state corresponding to all modules operating correctly to a state that corresponds to the system having failed. In other words, no single failure can cause the complete TMR system to fail. Finally, we assume that the system starts in the perfect state (111) where all of the system's modules are operating correctly.

The state diagram that results for the TMR system is shown in Fig. 4.16. As can be seen, the system begins in state (111) and, upon the first module failure, transitions to state (110), (101), or (011), depending on whether module 1, 2, or 3 is the module that fails. Note that the transition exists for the module to remain in a state if a module failure does not occur. The state diagram shown in Fig. 4.16 is analogous to the state diagram of a synchronous digital circuit. When some event, a module failure in the case of a reliability model and the occurrence of a clock signal in the case of a synchronous machine, occurs, the system transitions from one state to another.

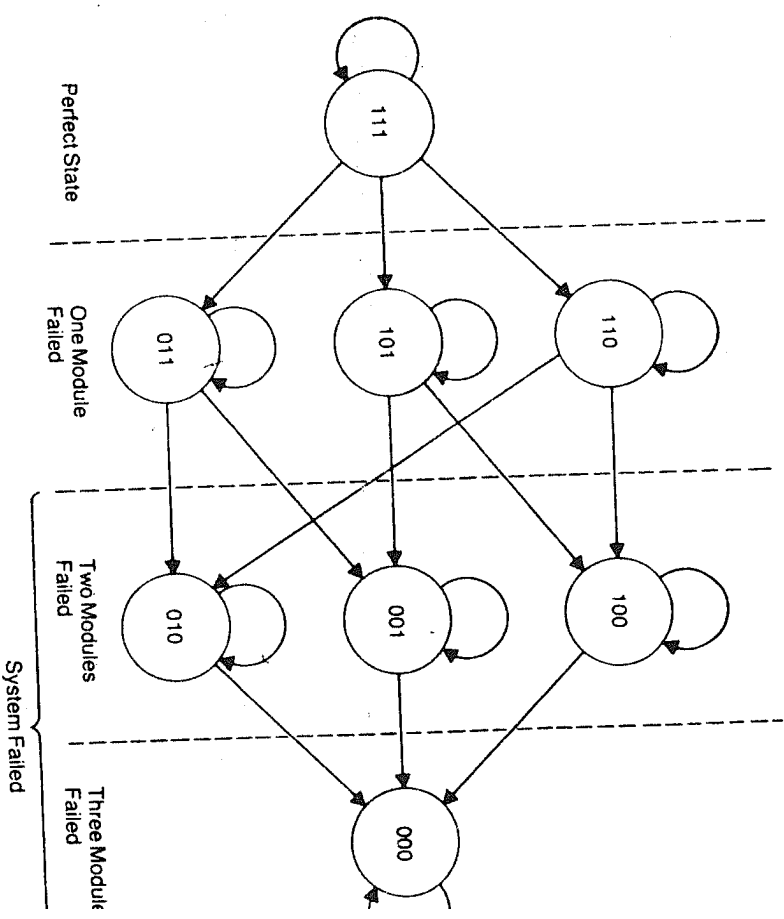


Fig. 4.16 State diagram showing possible states and state transitions for a system.

The states in the diagram shown in Fig. 4.16 can be partitioned into three major categories: the *perfect* state (111) in which all modules function correctly; the *one-failed* states (110), (101), and (011) in which a single module has failed, and the *system-failed* states (100), (001), (010), and (000) in which enough modules have failed to cause the system to fail. State transitions will be useful later when we attempt to reduce the Markov model to a simpler form.

As stated earlier, each state transition has associated with it a transition probability that describes the probability of that state transition occurring within a specified period of time. In the case of the TMR example in Fig. 4.16, the transition probability from state (111) to state (110) is the probability of a single module failure. If we assume that each module in the TMR system obeys the exponential failure law and has a constant failure rate λ , the probability of a module being failed at some time $t + \Delta t$, given that the module was operational at time t , is given by

$$1 - e^{-\lambda \Delta t}$$

The exponential can be written in a series expansion as

$$e^{-\lambda \Delta t} = 1 + (-\lambda \Delta t) + \frac{(-\lambda \Delta t)^2}{2!} + \dots$$

such that we have

$$1 - e^{-\lambda \Delta t} = 1 - \left[1 + (-\lambda \Delta t) + \frac{(-\lambda \Delta t)^2}{2!} + \dots \right] = (\lambda \Delta t) - \frac{(-\lambda \Delta t)^2}{2!} - \dots$$

For small values of Δt , the expression reduces to simply

$$1 - e^{-\lambda \Delta t} \approx \lambda \Delta t$$

In other words, the probability that a module will fail within the time period Δt is approximately $\lambda \Delta t$.

Referring to our example on the TMR system, the transition probabilities can now be specified for each possible state transition. Figure 4.17

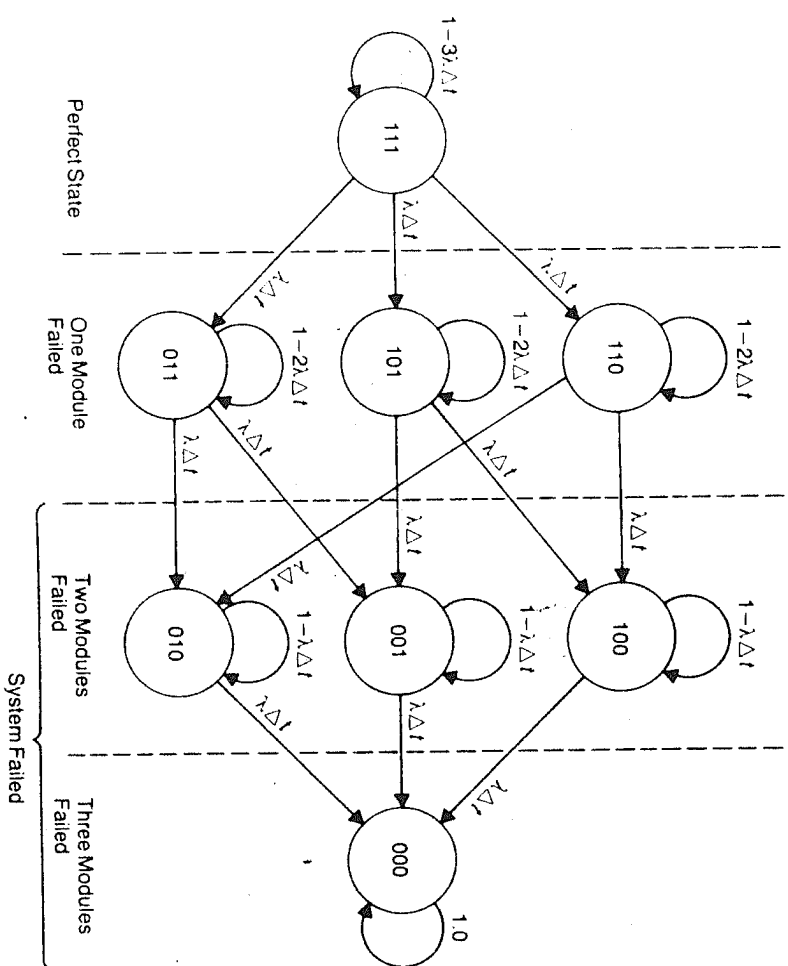


Fig. 4.17 Markov model of the TMR system showing possible states, state transitions, and failure categories.

shows the resulting state diagram for the Markov model of the TMR system including the specification of each state transition probability. It is possible, however, to reduce the Markov model of Fig. 4.17. As mentioned earlier, the states of the TMR model can be partitioned into three major classes: perfect state, the one-failed state, and the system-failed state. If we appropriately define the state transition probabilities, the several states within the TMR model can be combined.

Suppose that we let state 3 correspond to the state in which all modules in the TMR system are functioning correctly; state 2 is the state in which two modules are working correctly; state F is the failed state which two or more modules have failed. The resulting Markov model is illustrated as shown in Fig. 4.18. The state transition probabilities shown in Fig. 4.18 have been derived to account for one of several failures occurring. For example, the probability of transitioning from state 3 to state 2 depends on the probability of any one of three modules failing. Consequently, the transition probability assigned to the transition from state 3 to state 2 is $3\lambda(\Delta t)$. Likewise the transition probability assigned to the transition from state 2 to state F is $2\lambda(\Delta t)$.

The equations of the Markov model of the TMR system can be written easily from the state diagram shown in Fig. 4.18. The probability of the system being in any given state S at some time $t + \Delta t$ depends on the probability that the system was in a state from which it could transition to S and the probability of that transition occurring. For example, the probability that the TMR system will be in state 3 at time $t + \Delta t$ depends on the probability that the system was in state 3 at time t (since the system only transitions to state 3 from state 3) and the probability of the system transitioning from state 3 back into state 3.

In mathematical terms, we have

$$p_3(t + \Delta t) = (1 - 3\lambda \Delta t)p_3(t)$$

where $p_3(t)$ is the probability of being in state 3 at time t and $p_3(t + \Delta t)$ the probability of being in state 3 at time $t + \Delta t$. In a similar fashion, equations for the remaining two states can be written as




Fig. 4.18 Reduced Markov model of the TMR system with a minimal number of states.

$$p_2(t + \Delta t) = (3\lambda \Delta t)p_3(t) + (1 - 2\lambda \Delta t)p_2(t)$$

$$p_F(t + \Delta t) = (2\lambda \Delta t)p_2(t) + p_F(t)$$

assuming that the system will remain in the failed state if it ever enters the failed state. $p_2(t + \Delta t)$ is the probability of being in state 2 at time $t + \Delta t$; $p_2(t)$ is the probability of being in state 2 at time t ; $p_F(t + \Delta t)$ is the probability of being in state F at time $t + \Delta t$; and $p_F(t)$ is the probability of being in state F at time t .

The equations of the Markov model of the TMR system can be written in matrix form as

$$\begin{bmatrix} p_3(t + \Delta t) \\ p_2(t + \Delta t) \\ p_F(t + \Delta t) \end{bmatrix} = \begin{bmatrix} (1 - 3\lambda \Delta t) & 0 & 0 \\ 3\lambda \Delta t & (1 - 2\lambda \Delta t) & 0 \\ 0 & 2\lambda \Delta t & 1 \end{bmatrix} \begin{bmatrix} p_3(t) \\ p_2(t) \\ p_F(t) \end{bmatrix}$$

The resulting matrix equation can be written in a condensed form as

$$\mathbf{P}(t + \Delta t) = \mathbf{A}\mathbf{P}(t)$$

where

$$\mathbf{P}(t + \Delta t) = \begin{bmatrix} p_3(t + \Delta t) \\ p_2(t + \Delta t) \\ p_F(t + \Delta t) \end{bmatrix}$$

$$\mathbf{A} = \begin{bmatrix} (1 - 3\lambda \Delta t) & 0 & 0 \\ 3\lambda \Delta t & (1 - 2\lambda \Delta t) & 0 \\ 0 & 2\lambda \Delta t & 1 \end{bmatrix}$$

$$\mathbf{P}(t) = \begin{bmatrix} p_3(t) \\ p_2(t) \\ p_F(t) \end{bmatrix}$$

$\mathbf{P}(t)$ is the probability state vector at time t , $\mathbf{P}(t + \Delta t)$ is the probability state vector at time $t + \Delta t$, and \mathbf{A} is the transition matrix.

The matrix equations for the Markov model can be viewed as a difference equation for the purpose of obtaining a solution. By assuming some initial value of the probability state vector, $\mathbf{P}(0)$, the value of $\mathbf{P}(\Delta t)$ can be obtained as $\mathbf{P}(\Delta t) = \mathbf{A}\mathbf{P}(0)$. Similarly, the value of the probability state vector at time $2\Delta t$ can be written as $\mathbf{P}(2\Delta t) = \mathbf{A}\mathbf{P}(\Delta t) = \mathbf{A}^2\mathbf{P}(0)$. In general, the solution is given as

$$\mathbf{P}(n\Delta t) = \mathbf{A}^n\mathbf{P}(0)$$

The probability of the system failing is given by the probability of the system being in the failed state. For example, in the TMR illustration, the probability of the system failing is the element of $\mathbf{P}(t)$ given by $p_F(t)$. The reliability of the TMR system can be written as

$$R_{\text{TMR}}(t) = 1 - p_F(t) = p_3(t) + p_2(t)$$

The Markov models considered thus far have been *discrete-time* in which state transitions occur at fixed intervals Δt . It is possible to consider systems using *continuous-time* Markov models in which state transitions can occur at any point in time [Nelson 1986]. The continuous-time equations can be derived from the discrete-time equations by allowing the interval Δt to approach zero. For example, the equations of the discrete-time Markov model for the TMR system can be written as

$$\frac{p_3(t + \Delta t) - p_3(t)}{\Delta t} = -3\lambda p_3(t)$$

$$\frac{p_2(t + \Delta t) - p_2(t)}{\Delta t} = 3\lambda p_3(t) - 2\lambda p_2(t)$$

$$\frac{p_F(t + \Delta t) - p_F(t)}{\Delta t} = 2\lambda p_2(t)$$

through simple algebraic manipulations. Taking the limit as Δt approaches zero results in a set of differential equations given by

$$\frac{dp_3(t)}{dt} = -3\lambda p_3(t)$$

$$\frac{dp_2(t)}{dt} = 3\lambda p_3(t) - 2\lambda p_2(t)$$

$$\frac{dp_F(t)}{dt} = 2\lambda p_2(t)$$

The simultaneous differential equations can be solved using a variety of techniques. For example, if Laplace transforms are used, we have

$$sP_3(s) - p_3(0) = -3\lambda P_3(s)$$

$$sP_2(s) - p_2(0) = 3\lambda P_3(s) - 2\lambda P_2(s)$$

$$sP_F(s) - p_F(0) = 2\lambda P_2(s)$$

where $P_3(s)$ is the Laplace transform of $p_3(t)$, $P_2(s)$ is the Laplace transform of $p_2(t)$, $P_F(s)$ is the Laplace transform of $p_F(t)$, $p_3(0)$ is the initial value of $p_3(t)$ at time $t = 0$, $p_2(0)$ is the initial value of $p_2(t)$ at time $t = 0$, and $p_F(0)$ is the initial value of $p_F(t)$ at time $t = 0$. We assume in the analysis that the system starts in the perfect state at time $t = 0$, so $p_3(0) = 1$, $p_2(0) = 0$, and $p_F(0) = 0$. Consequently, the Laplace transform equations can be written as

$$P_3(s) = \frac{1}{s + 3\lambda}$$

$$P_2(s) = \frac{3\lambda}{(s + 2\lambda)(s + 3\lambda)}$$

$$P_f(s) = \frac{6\lambda^2}{s(s + 2\lambda)(s + 3\lambda)}$$

which can be rewritten as

$$P_3(s) = \frac{1}{s + 3\lambda}$$

$$P_2(s) = \frac{3}{(s + 2\lambda)} + \frac{-3}{(s + 3\lambda)}$$

$$P_f(s) = \frac{1}{s} + \frac{-3}{(s + 2\lambda)} + \frac{2}{(s + 3\lambda)}$$

Taking the inverse Laplace transform results in the solution given by

$$p_3(t) = e^{-3\lambda t}$$

$$p_2(t) = 3e^{-2\lambda t} - 3e^{-3\lambda t}$$

$$p_f(t) = 1 - 3e^{-2\lambda t} + 2e^{-3\lambda t}$$

Recall that the reliability of the TMR system is the probability of being in either state 3 or state 2, so

$$R_{TMR}(t) = p_3(t) + p_2(t) = e^{-3\lambda t} + 3e^{-2\lambda t} - 3e^{-3\lambda t} = 3e^{-2\lambda t} - 2e^{-3\lambda t}$$

which is exactly the same result obtained using the combinatorial techniques. Also note that the sum of $p_3(t)$, $p_2(t)$, and $p_f(t)$ is 1, as expected.

It is interesting to verify that the computer solution of the discrete-time Markov model yields the same results as the equations from the combinatorial model and the continuous-time model. For example, suppose we consider the TMR system. The primary reason for using the TMR system as an example is the relative ease with which both the Markov and the combinatorial models of the TMR system can be constructed. Recall that the combinatorial model of the TMR system that obeys the exponential failure law produces the reliability function

$$R(t) = 3e^{-2\lambda t} - 2e^{-3\lambda t}$$

Table 4.3 shows the values obtained from the combinatorial model of the TMR system at various points in time compared to the values obtained from the computer solution of the discrete-time Markov model of the TMR system. The Markov model was solved by assuming an initial state vector of $P(0) = (100)$ and using a time step, Δt of 0.1 seconds. The failure rate λ has been chosen as 0.1 failures per hour. The differences that exist between the

TABLE 4.3 Comparison of results from computer solution of the discrete-time Markov and the combinatorial model for the TMR system

Time (t) in minutes	Reliability	
	Combinatorial results	Markov results
1	0.99999177	0.99999171
2	0.99996674	0.99996686
3	0.99992549	0.99992561
4	0.99986792	0.99986809
5	0.99979424	0.99979442
6	0.99970472	0.99970472
7	0.99959898	0.99959916
8	0.99947786	0.99947786
9	0.99934101	0.99934095
10	0.99918842	0.99918854

Failure rate λ is 0.1 failures per hour, and time step Δt is 0.1 seconds.

two sets of numbers are within the computational accuracy used in the numbers.

We have seen how the Markov model can be used to model systems that do not depend on fault coverage or a repair process. Now we want to describe the process of developing a Markov model that depends on the coverage factor. After examining coverage, we will investigate the Markov model system with repair.

The system that we wish to model is a triply redundant system that uses fault detection techniques to detect the occurrence of a fault within the three independent modules. The modules provide their output to a flux-summer such that only one of the three modules must perform correctly for the system to function correctly. Consequently, the system operates as many as two module failures provided that the failures are appropriately handled. The correct way for a module failure to be handled is to be removed from the flux-summing operation by opening the switch. As long as the switch is closed, the associated module provides input to the flux-summer. Once the switch is opened, however, the module is completely disconnected from the flux-summer and no longer affects the operation of the system. The probability that a failure will be correctly handled is the fault coverage and is denoted as C . The basic structure of the system to be analyzed is shown in Fig. 4.19.

The Markov model of this system is similar to that of the basic system with majority voting. In fact, if the coverage factor C becomes

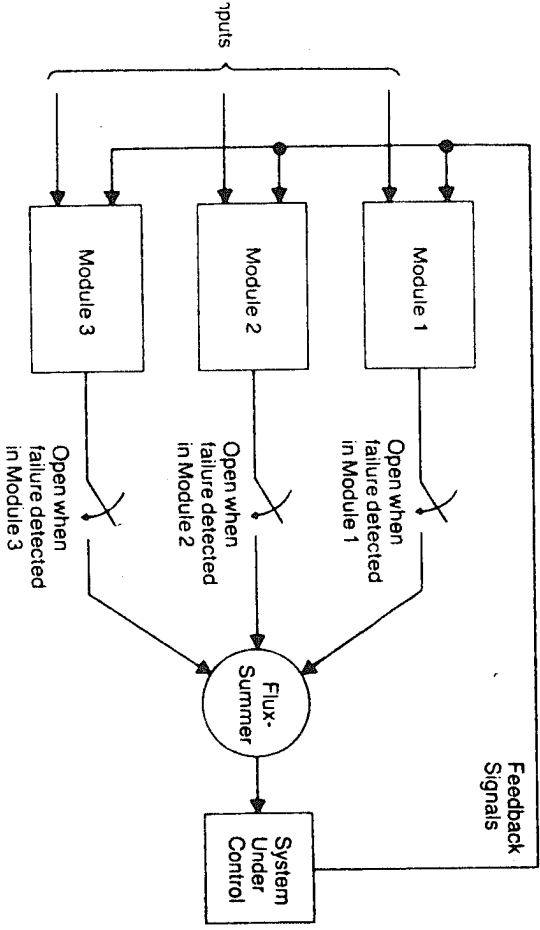


Fig. 4.19 An example hybrid redundancy technique to be used to illustrate the development of a Markov model that includes coverage.




Fig. 4.20 Discrete-time Markov model of the system shown in Fig. 4.19.

the system reduces to the basic TMR system. The Markov model of the system is shown in Fig. 4.20. The system is assumed to begin in the fault-free state, which is labeled as state 3. There are two paths by which the system can exit state 3. The first is shown as a transition to state 2 and corresponds to the failure of one of the three modules and the correct handling of that failure. The second transition that can occur from state 3 is the transition to state UD, which corresponds to one of the three modules failing and the failure going undetected or being handled inappropriately. Once the system enters state UD, it becomes equivalent to the basic TMR system with majority voting; the system cannot tolerate the second failure if the first failure is not handled correctly. The same types of transitions exist from state 2 to state 1 and the failed state. State 2 corresponds to the system having had one module to fail and having handled that failure correctly. While in state 2, the system can tolerate a failure and transition to state 1 provided that the failure is detected and handled correctly. Any undetected failures, however, take the system from state 2 to the failed state. Finally, any failure that occur while the system is in state 1 cause the system to transition immediately to the failed state.

While in any state, there is a nonzero probability that the system will remain in that particular state. For example, if the system is currently in state 3, the system remains in states 2, 1, and UD if the system is present in those states and new failures do not occur. The probability of being in the same state at the end of a Δt time period as at the beginning of that time period is calculated as $1 - P_{\text{exit}}(\Delta t)$, where $P_{\text{exit}}(\Delta t)$ is the probability of exiting a state during the Δt time period. For example, the probability of exiting state 3 during a Δt time period is the probability that any one of the three modules will fail. In other terms, the probability of exiting state 3 is $3\lambda\Delta t$. The probability of not exiting state 3 is, therefore, $1 - 3\lambda\Delta t$.

The equations of the Markov model for the system of Fig. 4.19 are developed as they were for the basic TMR system. The probability of being in state i at time $t + \Delta t$ depends on: (1) the probability of being in a state i from which the system can transition to state i , and (2) the associated transition probabilities. For example, the system can go to state 2 during Δt time period if and only if it is in either state 3 or state 2. Therefore, the probability of being in state 2 at time $t + \Delta t$ is

$$p_2(t + \Delta t) = 3\lambda\Delta t p_3(t) + (1 - 2\lambda\Delta t)p_2(t)$$

The complete set of equations for the Markov model of the system of Fig. 4.1 can be written as

TABLE 4.4 Reliability as a function of fault coverage for the system modeled using the Markov model of Fig. 4.20.

Fault coverage	Reliability (after 1 hour)
0.0	0.97460
0.1	0.97484
0.2	0.97558
0.3	0.97680
0.4	0.97852
0.5	0.98073
0.6	0.98343
0.7	0.98662
0.8	0.99030
0.9	0.99448
1.0	0.99914

Failure rate λ is 0.1 failures per hour, and time step Δt is 0.1 seconds.

$$\begin{bmatrix} p_3(t + \Delta t) \\ p_2(t + \Delta t) \\ p_1(t + \Delta t) \\ p_{UD}(t + \Delta t) \\ p_F(t + \Delta t) \end{bmatrix} = \begin{bmatrix} 1 - 3\lambda\Delta t & 0 & 0 & 0 & 0 \\ 3\lambda\Delta t C & 1 - 2\lambda\Delta t & 0 & 0 & 0 \\ 0 & 2\lambda\Delta t C & 1 - \lambda\Delta t & 0 & 0 \\ 3\lambda\Delta t(1 - C) & 0 & 0 & 1 - 2\lambda\Delta t & 0 \\ 0 & 2\lambda\Delta t(1 - C) & \lambda\Delta t & 2\lambda\Delta t & 1 \end{bmatrix} \begin{bmatrix} p_3(t) \\ p_2(t) \\ p_1(t) \\ p_{UD}(t) \\ p_F(t) \end{bmatrix}$$

The reliability of the system described by the Markov model of Fig. 4.20 is the probability of being in states 3, 2, 1, or *UD*. In other words, the reliability can be written as

$$R(t) = p_3(t) + p_2(t) + p_1(t) + p_{UD}(t)$$

It is interesting to note the effect that fault coverage has on the reliability of this system. Table 4.4, for example, shows the reliability of the system after one hour as a function of the coverage factor. The time step Δt has been selected as 0.1 seconds and the failure rate λ has been chosen as 0.1 failures per hour. For perfect coverage, the system achieves a reliability of approximately 0.99914, which is the same reliability that can be achieved by the perfect parallel system with three identical modules. When the cov-

erage is zero, the system is identical to the basic TMR system with no voting. In other words, the system will be a 2-of-3 system where two three modules must work for the system to work. The reliability achieved by this system when the coverage is zero is approximately 0.9746. Table 4.4 that the impact of changes in the fault coverage is more significant at the higher values of fault coverage.

We now consider systems that incorporate repair as a form of recovery. For example, many applications require that the repair rate's effect on system be modeled. A system that possesses a poor repair rate can be modeled to have fault tolerance to the extent that the system can function while elements are being repaired. The Markov model is an extremely useful tool for analyzing the effect that repair has on a system.

Consider the Markov model of a simple system consisting of one computer and no redundancy. The single computer might be a banking system for example, and we wish to model the failure and recovery process of a single computer. Further assume that the computer has a constant failure rate λ and a constant repair rate μ . During the time interval Δt , the computer will have a probability of failure given by $\lambda\Delta t$. Since the repair rate is analogous to the failure rate and represents the number of repairs expected to occur in a specific time period, the probability of a repair occurring within the time period Δt is $\mu\Delta t$. Using this information allow us to formulate the simple Markov model shown in Fig. 4.21 for the computer. State *O* represents the condition of the computer being completely operational, and state *F* represents the failed condition. If the computer is in state *O*, the probability of the system transitioning to state *F* during the time period Δt is $\lambda\Delta t$. Likewise, if the system is in state *F*, the probability of transitioning to state *O* is $\mu\Delta t$. As we discovered during previous examples if the system is in state *O* and a failure does not occur, the system remains in state *O*. Similarly, if the system is in state *F* and a repair does not occur, the system remains in state *F*.

The equations for the Markov model of Fig. 4.21 can be written as




Fig. 4.21 Markov model of a simple nonredundant system with repair

$$\begin{bmatrix} p_o(t + \Delta t) \\ p_f(t + \Delta t) \end{bmatrix} = \begin{bmatrix} 1 - \lambda\Delta t & \mu\Delta t \\ \lambda\Delta t & 1 - \mu\Delta t \end{bmatrix} \begin{bmatrix} p_o(t) \\ p_f(t) \end{bmatrix}$$

It is interesting to solve this Markov model to determine the effect that the repair rate has on the probability of the system being operational. Figure 4.22 shows the plot of the probability of the system being in state O versus the repair rate. The failure rate λ was selected as 0.1 failures per hour, and the time step Δt was chosen as 0.1 seconds for this example. The system was assumed to start in the operational state.

It is also instructive to determine the continuous-time solution for the model shown in Fig. 4.21. Using a procedure identical to that used to determine the continuous-time equations for the TMR system, the equations from the discrete-time Markov model are manipulated algebraically to obtain

$$\begin{aligned} \frac{p_o(t + \Delta t) - p_o(t)}{\Delta t} &= -\lambda p_o(t) + \mu p_f(t) \\ \frac{p_f(t + \Delta t) - p_f(t)}{\Delta t} &= \lambda p_o(t) - \mu p_f(t) \end{aligned}$$




Fig. 4.22 Probability of remaining operational for the system described by the Markov model of Fig. 4.21.

Taking the limit as Δt approaches zero results in the differential equations given by

$$\begin{aligned} \frac{dp_o(t)}{dt} &= -\lambda p_o(t) + \mu p_f(t) \\ \frac{dp_f(t)}{dt} &= \lambda p_o(t) - \mu p_f(t) \end{aligned}$$

Assuming that the initial conditions are $p_o(0) = 1$ and $p_f(0) = 0$ and using Laplace transforms results in the simultaneous equations

$$\begin{aligned} sP_o(s) &= 1 - \lambda P_o(s) + \mu P_f(s) \\ sP_f(s) &= \lambda P_o(s) - \mu P_f(s) \end{aligned}$$

Solving the simultaneous equations for $P_o(s)$ and $P_f(s)$ yields

$$\begin{aligned} P_o(s) &= \frac{1}{s + (\lambda + \mu)} + \frac{\mu}{s(s + (\lambda + \mu))} \\ P_f(s) &= \frac{\lambda}{s(s + (\lambda + \mu))} \end{aligned}$$

which can be rewritten as

$$\begin{aligned} P_o(s) &= \frac{\mu}{s} + \frac{\lambda + \mu}{s + (\lambda + \mu)} \\ P_f(s) &= \frac{\lambda}{\lambda + \mu} - \frac{\lambda + \mu}{s + (\lambda + \mu)} \end{aligned}$$

Taking the inverse Laplace transform results in the time-domain solution given by

$$\begin{aligned} p_o(t) &= \frac{\mu}{\lambda + \mu} + \frac{\lambda}{\lambda + \mu} e^{-(\lambda + \mu)t} \\ p_f(t) &= \frac{\lambda}{\lambda + \mu} - \frac{\lambda}{\lambda + \mu} e^{-(\lambda + \mu)t} \end{aligned}$$

Several interesting features are apparent in the time-domain expressions for $p_o(t)$ and $p_f(t)$. For example, as time approaches infinity, $p_o(t)$ approaches the constant value of

$$p_o(\infty) = \frac{\mu}{\lambda + \mu} = \frac{1}{\frac{\lambda}{\mu} + 1}$$

and $p_F(t)$ approaches

$$p_F(\infty) = \frac{\lambda}{\lambda + \mu} = \frac{1}{\frac{\mu}{\lambda} + 1}$$

As we will discover when discussing availability modeling, the value of $p_O(t)$ as time approaches infinity is the steady-state availability.

4.4 Safety Modeling

The safety of a system, as defined in Chapter 1, is the probability that the system will either perform correctly or will fail in a safe manner. The concepts of *safe* and *unsafe* are highly dependent upon the application. In many cases, for example, a safe course of action is to simply turn the system off after a failure occurs. In some applications, however, turning the system off can be a disastrous course of action. In any case, however, the fundamental concept of safety analysis is that a system will possess two different ways in which it can fail: one system failure is defined as *safe*, and the other is categorized as *unsafe*. The definition of safe and unsafe failures must be created uniquely for each application.

Safety can be modeled using Markov models by splitting the system failed state into two separate states. One failed state is normally labeled FS for *failed unsafe*. A Markov model for a simple system containing one hardware module with a failure rate of λ and self-diagnostics with a fault detection coverage of C is shown in Fig. 4.23. Safe failures are defined in this example as those that are detected by the self-diagnostics. Consequently, unsafe failures are defined as those that are not detected by the self-diagnostics. If a failure occurs, the system transitions to either state FS or FU depending on whether or not the condition is detected.




Fig. 4.23 Three-state Markov model allowing safety calculations.

The safety of the system described by the Markov model of Fig. 4.23 can be written as

$$S(t) = p_O(t) + p_{FS}(t)$$

where $S(t)$ is the safety, $p_O(t)$ is the probability of being in the O state at time t , and $p_{FS}(t)$ is the probability of being in the failed state at time t . The complete equations of the discrete-time Markov model can be written as

$$p_O(t + \Delta t) = (1 - \lambda \Delta t)p_O(t)$$

$$p_{FS}(t + \Delta t) = \lambda \Delta t C p_O(t) + p_{FS}(t)$$

$$p_{FU}(t + \Delta t) = \lambda \Delta t (1 - C)p_O(t) + p_{FU}(t)$$

As we have done previously, the differential equations of the corresponding continuous-time Markov model can be written as

$$\frac{dp_O(t)}{dt} = -\lambda p_O(t)$$

$$\frac{dp_{FS}(t)}{dt} = \lambda C p_O(t)$$

$$\frac{dp_{FU}(t)}{dt} = \lambda (1 - C)p_O(t)$$

Taking the Laplace transform results in

$$P_O(s) = \frac{p_O(0)}{s + \lambda}$$

$$P_{FS}(s) = \frac{\lambda C p_O(0)}{s(s + \lambda)} + \frac{p_{FS}(0)}{s}$$

$$P_{FU}(s) = \frac{\lambda (1 - C)p_O(0)}{s(s + \lambda)} + \frac{p_{FU}(0)}{s}$$

where $p_O(0)$, $p_{FS}(0)$, and $p_{FU}(0)$ are the initial values of the respective probabilities. If we assume that the system begins in state O such that $p_O(0) = 1$, $p_{FS}(0) = 0$, and $p_{FU}(0) = 0$, we obtain

$$P_O(s) = \frac{1}{s + \lambda}$$

$$P_{FS}(s) = \frac{\lambda C}{s(s + \lambda)} = \frac{C}{s} - \frac{C}{s + \lambda}$$

$$P_{FU}(s) = \frac{\lambda (1 - C)}{s(s + \lambda)} = \frac{(1 - C)}{s} - \frac{(1 - C)}{s + \lambda}$$

The time domain solutions can now be written as

$$\begin{aligned}
 P_0(t) &= e^{-Nt} \\
 P_{FS}(t) &= C - Ce^{-Nt} \\
 P_{FU}(t) &= (1 - C) - (1 - C)e^{-Nt}
 \end{aligned}$$

Intuitively, the equations are as expected. For example, the reliability of the system is

$$R(t) = P_0(t) = e^{-Nt}$$

and the probability of being in one of the two failed states is

$$P_{FS}(t) + P_{FU}(t) = 1 - e^{-Nt} = 1 - R(t)$$

The safety of the system is written as

$$S(t) = P_0(t) + P_{FS}(t) = C + (1 - C)e^{-Nt}$$

At time $t = 0$, the safety of the system is 1, as expected. As time approaches infinity, however, the safety approaches

$$S(\infty) = C$$

In other words, if the fault detection coverage is perfect ($C = 1$), the system has perfect safety. However, if the fault detection coverage is nonexistent ($C = 0$), the system will eventually fail in an unsafe manner. The safety of the system, in this example, is directly dependent upon the fault detection coverage. In subsequent sections, we investigate the safety of more complex systems.

4.5 System Comparisons

Now that we have several tools at our disposal, we can begin to examine the process of comparing two or more systems. When we make comparisons, we must be careful about the parameters that we choose to compare. For example, if we choose to compare the MTTF of two systems or the reliability of the systems, the results can be surprising. Suppose that we wish to compare a simplex system consisting of a single computer to a TMR system with three computers in a majority voting arrangement. Assume for simplicity that the majority voter is perfect. The two systems are shown in Fig. 4.24. The computers in each system are identical and are assumed to obey the exponential failure law.

Recall that the MTTF of a system is defined as

$$MTTF = \int_0^{\infty} R(t) dt$$

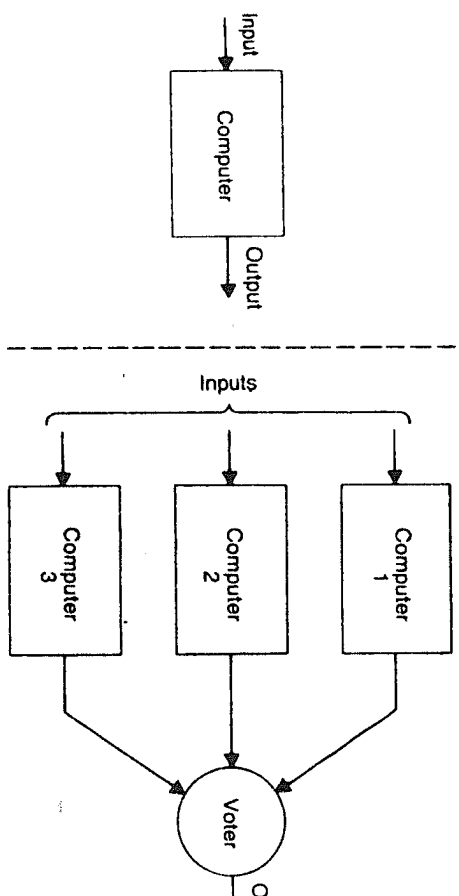


Fig. 4.24 Simplex and TMR systems to be compared to assess the benefits of redundancy.

where $R(t)$ is the reliability function of the system. The reliability function of the simplex system is

$$R_{\text{simplex}}(t) = e^{-Nt}$$

whereas that of the TMR system is

$$R_{\text{TMR}}(t) = 3e^{-2Nt} - 2e^{-3Nt}$$

If we integrate $R_{\text{simplex}}(t)$ and $R_{\text{TMR}}(t)$, we find that the MTTF of each system is given by

$$\begin{aligned}
 MTTF_{\text{simplex}} &= \int_0^{\infty} e^{-Nt} dt = \frac{1}{N} \\
 MTTF_{\text{TMR}} &= \int_0^{\infty} (3e^{-2Nt} - 2e^{-3Nt}) dt = \frac{5}{6N}
 \end{aligned}$$

Thus, the MTTF of the TMR system is lower than the MTTF of the simplex system.

Based on these calculations, we might conclude that the TMR system is not as good as the simplex system. This may or may not be a correct conclusion depending on the application and the length of time the system is expected to operate correctly. Figure 4.25 provides a good insight into the reason why the MTTF of the TMR system is less than that of the simplex system. Figure 4.25 shows the reliability of the simplex and the TMR




Fig. 4.25 Simplex and TMR reliabilities versus λt —a crossover point occurs where the two reliabilities are equal.

terms as functions of λt . As can be seen, a point in time is reached when the TMR system becomes less reliable than the simplex system. The MTTF is the area under the reliability curve, and that area is larger for the simplex than for the TMR configuration. The point of this discussion is that the MTTF can sometimes misrepresent the quality of a system. For certain values of the product λt , the TMR system's reliability will be superior to that of the simplex system. Regardless of the reliability, if fault tolerance is necessary, the TMR system will be superior to the simplex system.

Sometimes, a single parameter comparison that is better than the MTTF is the **mission time**, denoted as $MT[r]$. The mission time is the time at which the reliability of a system falls below the level r . For example, a simplex system that obeys the exponential failure law has a reliability of r when

$$r = e^{-\lambda t}$$

The time at which the reliability value of r occurs can be found by taking the natural logarithm of both sides of the preceding equation and solving the resulting equation. The solution yields

$$MT[r] = \frac{-\ln(r)}{\lambda}$$

A **mission time improvement** can be calculated as the ratio of mission times of the two systems being compared. For example, suppose we wish to compare two computer systems: (1) a simplex computer with a single computer that has a failure rate of 0.01 failures per (2) a TMR system constructed using three of that same computer. Puters are assumed to obey the exponential failure law. We wish to determine the mission time improvement of the TMR system over the system for a reliability of 0.86.

The $MT_{\text{simplex}}[0.86]$ is fairly easy to calculate from the exponential ability function as

$$MT_{\text{simplex}}[0.86] = \frac{-\ln(r)}{\lambda} = \frac{-\ln(0.86)}{0.01} = 15.08 \text{ hours}$$

The $MT_{\text{TMR}}[0.86]$ is found from the solution of the equation

$$R_{\text{TMR}}(t) = 3e^{-2 \times 0.01t} - 2e^{-3 \times 0.01t} = 0.86$$

for the time t , which results in $MT_{\text{TMR}}[0.86] = 27$ hours. The mission time improvement of the TMR system over the simplex system is approximately 1.8. In other words, the TMR system, in this example, can operate as long as the simplex system while still maintaining a reliability greater than 0.86. The graph shown in Fig. 4.26 illustrates the concepts of mission time and the mission time improvement for the comparison of a simplex and the TMR systems with the failure rate of 0.01 failures per

4.6 Availability Models

Thus far, we have considered only the modeling of the reliability of a system. However, we have seen in the discussions of Chapters 1 and 2 that parameters such as availability and maintainability are also important in the analysis of fault-tolerant systems. Many computer companies are concerned more with the probability of their systems being available to their customers than with the probability of their systems failing. Availability is the time the system can operate without failure (related to reliability) rather than with failure. As a result, the rate at which a system can be repaired becomes a critical parameter of the design. The repair rate can dramatically affect the availability of a system.

Recall that the availability $A(t)$ of a system is defined as the probability that a system will be available to perform its tasks at the instant t . Intuitively, we can see that the availability can be approximated by the total time that a system has been operational divided by the total elapsed time since the system was initially placed into operation. In other words, the availability is the percentage of time that the system is available.




Fig. 4.26 Mission time improvement comparison between simplex and TMR.

form its expected tasks. Suppose that we place a system into operation at time $t = 0$. As time moves along, the system performs its functions, perhaps fails, and is repaired. At some time $t = t_{\text{current}}$, suppose that the system has operated correctly for a total of t_{op} hours and has been in the process of repair or waiting for repair for a total of t_{repair} hours. The time t_{current} is then the sum of t_{op} and t_{repair} . The availability can be determined as

$$A(t_{\text{current}}) = \frac{t_{\text{op}}}{t_{\text{op}} + t_{\text{repair}}}$$

where $A(t_{\text{current}})$ is the availability at time t_{current} .

The preceding expression lends itself well to the experimental evaluation of the availability of a system; we can simply place the system into operation and measure the appropriate times required to calculate the availability of the system at a number of points in time. Unfortunately, the experimental evaluation of the availability is often not possible because of the time and expense involved. Also, we would like to have some means of estimating the availability before we actually build the system so that availability considerations can be factored into the design process. We will consider two approaches. The first is based on the single parameter measures such as MTTF and MTTR and yields what is typically called the

steady-state availability A_{ss} . The second approach uses the failure rate the repair rates in a Markov model to calculate the availability as a function of time.

We have seen that availability is basically the percentage of time system is operational. Using knowledge of the statistical interpretation of the MTTF and the MTTR, we expect that, on the average, a system will operate for MTTF hours and then encounter its first failure. Once the failure has occurred, the system will then, again on the average, require λ hours to be repaired and placed into operation once again. The system then operate for another MTTF hours before encountering its second failure. This concept has been illustrated in Fig. 4.2.

If the average system experiences N failures during its lifetime, then the time that the system is operational is $N(\text{MTTF})$ hours. Likewise, the time that the system is "down" for repairs is $N(\text{MTTR})$ hours. In other words, the operational time t_{op} is $N(\text{MTTF})$ hours and the down-time t_{down} is $N(\text{MTTR})$ hours. The average, or steady-state, availability is

$$A_{ss} = \frac{N(\text{MTTF})}{N(\text{MTTF}) + N(\text{MTTR})} = \frac{\text{MTTF}}{\text{MTTF} + \text{MTTR}}$$

We know, however, that the MTTF and the MTTR of a simplex system are related to the failure rate and the repair rate, respectively, as

$$\begin{aligned} \text{MTTF} &= \frac{1}{\lambda} \\ \text{MTTR} &= \frac{1}{\mu} \end{aligned}$$

Therefore, the steady-state availability of a simplex system is given

$$A_{ss} = \frac{\frac{1}{\lambda}}{\frac{1}{\lambda} + \frac{1}{\mu}} = \frac{1}{1 + \frac{\lambda}{\mu}}$$

Recall that the repair rate μ is expressed in repairs per hour, while the failure rate λ is in failures per hour. We would expect that if the failure rate goes to 0, implying that the system never fails, or the repair rate goes to infinity, implying that no time is required to repair the system, then the availability will go to 1. Looking at the expression for the steady-state availability, we can see that this is true.

As an example calculation, consider a computer system that has a failure rate of one failure every 100 hours and a repair rate of one repair every 10 hours. The failure rate of this system is $\lambda = 0.01$ failures per hour

repair rate is $\mu = 0.1$ repairs per hour. The steady-state availability is calculated as

$$A_{ss} = \frac{1}{1 + \frac{0.01}{0.1}} = 0.90909$$

This implies that the system is available for use an average of slightly more than 90% of the time.


Now suppose we investigate the use of the Markov model as a means of determining the availability of a system. We already have the necessary tools to accomplish this task. The Markov model of a system with repair is in fact the model required to calculate the availability of a system. Recall the two-state model of a simple system with repair and having a failure rate of λ and a repair rate of μ . The state diagram of this model is repeated in Fig. 4.27 for convenience. State O represents the state in which the system is completely operational, whereas state F is the state in which the system has failed and is in the process of being repaired. The probability of the system failing during the time interval Δt is given by $\lambda \Delta t$, whereas the probability of the system being repaired during the time interval Δt is $\mu \Delta t$.

The equations of the Markov model are given by

$$\begin{bmatrix} p_o(t + \Delta t) \\ p_f(t + \Delta t) \end{bmatrix} = \begin{bmatrix} 1 - \lambda \Delta t & \mu \Delta t \\ \lambda \Delta t & 1 - \mu \Delta t \end{bmatrix} \begin{bmatrix} p_o(t) \\ p_f(t) \end{bmatrix}$$

where $p_o(t)$ is the probability, at time t , that the system is in the operational state and is, therefore, available to perform its tasks. Consequently, $p_o(t)$ is the availability of the system.

As an example, the Markov model shown in Fig. 4.27 has been solved for the failure rate of $\lambda = 0.01$ failures per hour and the repair rate of $\mu = 0.1$ repairs per hour. The plot of the resulting availability is shown in Fig. 4.28. Note that the availability approaches the value of 0.90909 that was previ-



O = Operational State
F = Failed State

Fig. 4.27 Markov model of a simple nonredundant system with repair.




Fig. 4.28 Availability versus time for a simple nonredundant system with 0.01 failures per hour and $\mu = 0.1$ repairs per hour.

ously determined as the steady-state availability of a system with this particular failure rate and repair rate.

4.7 Maintainability Models

As defined in Chapter 1, the maintainability is the probability that a system will be restored to working order within a specified time. We use the notation that $M(t)$ is the maintainability for time t . In other words, $M(t)$ is the probability that a system will be repaired in a time less than or equal to t .

An important parameter in maintainability modeling is the repair rate μ . The repair rate is the average number of repairs that can be performed per time unit. The inverse of the repair rate is the MTTR, which is the average time required to perform a single repair. Mathematically, the relationship between the repair rate and the MTTR is given by

$$MTTR = \frac{1}{\mu}$$

In industry, the MTTR, and as a result μ , is usually derived in an experimental fashion. A system can be constructed and faults injected; the average time required to repair the system is measured and recorded as the MTTR. A good estimate of the MTTR can be obtained only if a sufficient number of different faults are injected and repair personnel with a variety of skill levels are used.

An expression for the maintainability of a system can be derived in a manner similar to that used to develop the exponential failure law for the reliability function. Suppose that we have N systems. We inject one unique fault into each of the systems, and we allow one maintenance person to repair each system. We begin this experiment by injecting the faults into the systems at time $t = 0$. Later, at some time t , we determine that $N_r(t)$ of the systems have been repaired and $N_{nr}(t)$ have not been repaired. Since the maintainability of a system at time t is the probability that the system can be repaired by time t , an estimate of the maintainability can be computed as

$$M(t) = \frac{N_r(t)}{N} = \frac{N_r(t)}{N_r(t) + N_{nr}(t)}$$

If we differentiate $M(t)$ with respect to time, we obtain

$$\frac{dM(t)}{dt} = \frac{1}{N} \frac{dN_r(t)}{dt}$$

which can also be written as

$$\frac{dN_r(t)}{dt} = N \frac{dM(t)}{dt}$$

The derivative of $N_r(t)$ is simply the rate at which components are repaired at the instant of time t .

At time t , we have $N_{nr}(t)$ systems that have not been repaired. If we divide $dN_r(t)/dt$ by $N_{nr}(t)$, we obtain

$$\frac{1}{N_{nr}(t)} \frac{dN_r(t)}{dt}$$

which is called the *repair rate function* and is assumed to have a constant value of μ , the repair rate; μ has the units of repairs per unit of time.

Using the expression for the repair rate and the expression for the derivative of $N_r(t)$, we can write

$$\mu = \frac{1}{N_{nr}(t)} \frac{dN_r(t)}{dt} = \frac{N}{N_{nr}(t)} \frac{dM(t)}{dt}$$

which yields a differential equation of the form

$$\frac{dM(t)}{dt} = \mu \frac{N_{nr}(t)}{N}$$

We know, however, that $N_{nr}(t)/N$ is $1 - M(t)$, so we can write

$$\frac{dM(t)}{dt} = \mu(1 - M(t))$$

The solution to the differential equation is well known and is given by

$$M(t) = 1 - e^{-\mu t}$$

The relationship developed for $M(t)$ has the desired character. First, if the repair rate is zero, the maintainability is also zero since the system cannot be repaired in any length of time. Second, if the repair rate is infinite, the maintainability is one since repair can be performed in any time. A final interesting feature of the maintainability function is its relationship to the MTTR. At $t = \text{MTTR}$, the maintainability function will be

$$M(t = \text{MTTR}) = 1 - e^{-\mu/\mu} = 1 - e^{-1} = .632$$

which implies that there is a probability of 0.632 that a system will be repaired in a time less than or equal to its MTTR.

As we have seen, the repair rate plays a crucial role in the maintainability of a system. The repair rate can differ depending on the type of repair that must be performed. For example, a banking system that is repaired on location by a local maintenance person will have a better repair rate than one that must be returned to the factory or some third party for repair. In addition, certain types of faults can be easily repaired on location whereas others require facilities that are not practical to bring to the site of the electronic system. For example, the replacement of a memory card can be performed easily at the site of the system, but the replacement of the power supply and cooling system can be much more difficult.

Because of the preceding issues, the repair rate is typically specified in terms of several levels of repair. The most common partitioning is to provide several levels of repair. The first is called the *organizational level* and consists of repairs that can be performed at the site where the system is located. Organizational repairs typically include all faults that can be located to a specific circuit cards such that the cards can simply be replaced and the system made operational once again. For example, if an aircraft can be repaired without bringing it off the runway, it is considered an organizational repair. The key to organizational repairs is the ability to locate the fault; it is seldom feasible to bring sophisticated fault detection and location equipment to the site of the system. Repairs at the organizational level often depend on the built-in test provided by the system to locate specific problems.

The second level of repair is called the *intermediate level*. Intermediate level repairs cannot be performed at the organizational level, but they can be performed in the immediate vicinity of the system. For example,

puter firm can have a local repair facility to which the faulty pieces of equipment are taken for repair. Intermediate level repair is not as good as being able to perform the repair on site, but it is better than having to return a piece of equipment to the factory. In the case of an airplane, for example, an intermediate level repair might be made in the hangar as opposed to on the runway.

The final level of repair is called the *depot level* or the *factory level*. In depot level repairs, the equipment must be returned to a major facility for the repair process. For example, if a calculator cannot be repaired at home (organizational level), it is taken to the store from which it was purchased (intermediate level). If the store is unable to perform the repair, they send the calculator to a site designated by the manufacturer as a major repair facility (depot level). The length of time required to perform the repair depends on the level at which it is performed. It may take less than an hour to repair a device at the organizational level, several hours or perhaps days at the intermediate level, and as much as several weeks or months at the depot level.

As an example, assume that the MTTTR for a computer system is 2.0 hours at the organizational level, 8.0 hours at the intermediate level, and one week (168 hours) at the depot level. The resulting maintainability functions are plotted versus time in Fig. 4.29. Note the tremendous difference that exists between the maintainability of the system for the different levels of repair.

4.8 Redundancy Ratios

A system that is more reliable or more available than another is better with regards to that one attribute. However, to achieve the improved reliability, availability, or maintainability, the system may contain an excessive amount of redundancy. The cost of the extra redundancy will appear in the weight, size, power consumption, volume, and financial costs of the improved system. In many applications, the improvements in the reliability, for example, may not be worth the extra weight that the system contains.

One good measure of the impact that improvements in reliability, availability, and maintainability have on a system is the **redundancy ratio**. The redundancy ratio is defined simply as the amount of hardware, information, time, or software that the redundant system requires divided by the amount required in a nonredundant system that performs the same function. The redundancy ratio can be specified for each type of hardware component; for example, the processors, memory, buses, interface units, power supplies, and displays. The redundancy ratio gives a measure of the extra resources required for a given application.




Fig. 4.29 Maintainability $M(t)$ versus time for example values of the M

A good example of the redundancy ratio can be found in a TMR processor system. If the nonredundant system requires only a single processor the TMR system will clearly require three and will have a redundancy ratio of three. Redundancy ratios are equally applicable for the processors of three. Redundancy ratios are equally applicable for software as well as hardware. If a nonredundant system requires lines of code and the redundant system requires 40,000 lines, the so redundancy ratio is 1.33. In other words, one-third more software required to implement the redundant system than is required in the redundant system.

4.9 Qualitative Methods

Thus far, we have discussed methods of evaluation that generate numbers to use to compare two or more systems. For example, one system may have a higher reliability, another may be less expensive, a third weigh less, and still another may consume less power. Often, we find certain attributes of a system that enter into the design process are extremely difficult to quantify. We may anticipate a drop in the prices of hardware at a future date, or we may feel that the production plan