In most practical applications the MTTR is a small fraction of the MTTF, so the *approximation* that the MTBF and MTTF are equal is often quite good. Conceptually, however, it is crucial to understand the difference between the MTBF and the MTTF.

4.2.6. Fault Coverage

An extremely important parameter in the design and analysis of fault-tolerant systems is **fault coverage.** The fault coverage available in a system can have a tremendous impact on the reliability, safety, and other attributes of the system. There are several types of fault coverage, depending on whether the designer is concerned with fault detection, fault location, fault containment, or fault recovery. In addition, there are two primary definitions of fault coverage: one is intuitive, the other is more mathematical.

The intuitive definition is that *coverage* is a measure of a system's ability to perform fault detection, fault location, fault containment, and/or fault recovery. The four primary types of fault coverage are fault detection coverage, fault location coverage, fault containment coverage, and fault recovery coverage. **Fault detection coverage** is a measure of a system's ability to detect faults. For example, a system requirement may be that a certain fraction of all faults be detected; the fault detection coverage is a measure of the system's capability to meet such a requirement. **Fault location coverage** is a measure of a system's ability to locate faults. Once again, it is very common to require a system to locate faults to within easily replaceable modules, and the fault location coverage is a measure of the success with which fault location is performed. **Fault containment coverage** is a measure of a system's ability to contain faults; specifically, the fault containment coverage represents a system's ability to make the *extent* attribute of faults *local* instead of *global*. Finally, **fault recovery coverage** is a measure of a system's ability to recover from faults and maintain an operational status. Clearly, a high fault recovery coverage requires high fault detection, location, and containment coverages.

In the evaluation of fault-tolerant systems, the fault recovery coverage is the most commonly considered, and the general term "fault coverage" is often used to mean fault recovery coverage. In other words, fault coverage is interpreted as a measure of a system's ability to successfully recover after the occurrence of a fault, therefore tolerating the fault. Therefore, when using the term "fault coverage," make sure that the type of coverage—detection, location, containment, or recovery—is understood.

The remainder of this chapter uses the term "fault coverage" to imply fault recovery coverage since fault recovery is the most common form of coverage encountered. In all cases, however, it will be made clear whether detection, location, containment, or recovery coverage is being considered.

Fault coverage is mathematically defined as the conditional probability that, given the existence of a fault, the system recovers [Bouricius, Carter, and Schneider 1969]. In mathematical terms, fault coverage is written as

$$C = P(\text{fault recovery} \mid \text{fault existence})$$

where C is the fault coverage and $P(\text{fault recovery} \mid \text{fault existence})$ is read as the probability of fault recovery *given* the existence of a fault. Recall that fault recovery is the process of maintaining or regaining operational status after a fault occurs.

The fundamental problem with fault coverage is that it is extremely difficult to calculate. Probably the most common approach to estimating fault coverage is to develop a list of all the faults that can occur in a system and to form, from that list, a list of faults that can be detected, a list of faults that can be located, a list of faults that can be contained, and a list of faults from which the system can recover. The fault detection coverage factor, for example, is then computed as simply the fraction of faults that can be detected; that is, the number of faults detected divided by the total number of faults. The remaining fault coverage factors are calculated in a similar manner. As an example, consider the circuit shown in Fig. 4.3 which has fifteen potential sites of stuck-at-1 or stuck-at-0 faults; consequently, there are a total of 30 faults. Table 4.2 shows the input combinations that yield erroneous outputs when certain faults are present, therefore detecting the faults. Note that the circuit performs correctly even if a single stuck-at-0 fault on one of the lines F, G, or M occurs. In other words, a single stuck-at-0 fault on line F, G, or M cannot be detected. As a result, the fault detection coverage for the circuit of Fig. 4.3 is (30–3)/30, or 0.9. In other words, 90% of the stuck-at-1 and stuck-at-0 faults are detected by at least one of the input combinations.

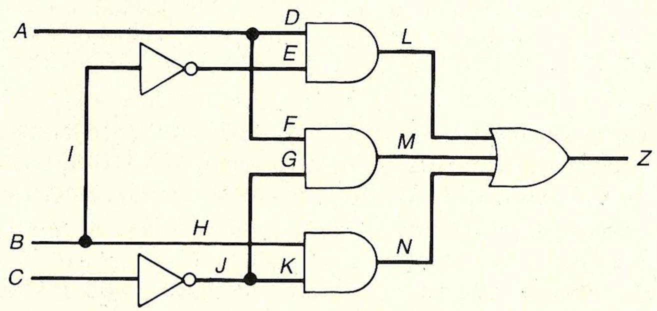

Fig. 4.3 Combinational circuit to illustrate fault detection coverage.

TABLE 4.2 Input patterns capable of detecting faults (test vectors) in the circuit of Fig. 4.3

Fault	Number of test vectors	Test vectors ABC
A_0	2	100, 101
A_1	2	000, 001
B_0	2	010, 111
B_1	2	000, 101
C_0	2	011, 111
C_1	2	010, 110
D_0	1	101
D_1	2	000, 001
E_0	1	101
E_1	1	111
F_0	0	—
F_1	2	000, 101
G_0	0	—
G_1	1	111
H_0	1	010
H_1	1	000
I_0	1	111
I_1	1	101
J_0	2	010, 110
J_1	2	011, 111
K_0	1	010
K_1	2	011, 111
L_0	1	101
L_1	4	000, 001, 011, 111
M_0	0	—
M_1	4	000, 001, 011, 111
N_0	1	010
N_1	4	000, 001, 011, 111
Z_0	4	010, 100, 101, 110
Z_1	4	000, 001, 011, 111

Several important points should be made about the estimation of coverage. First, the estimation of fault coverage requires the definition of the types of faults that can occur. Stating that the fault detection coverage is 0.9, for example, is meaningless unless the types of faults considered are identified. For example, the fault detection coverage for the circuit of Fig. 4.3 is 0.9 for all stuck-at-1 and stuck-at-0 faults, but the fault detection coverage may decrease substantially if stuck-open faults are included.

A second important point about the fault coverage is that it is typically assumed to be a constant. It is easy to envision applications in which the

probability of detecting a fault, for example, increases as a function of tim[e] [af]ter the occurrence of the fault. However, to simplify the analysis, the var[ious] fault coverages are normally assumed to be constants.

4.3 Reliability Modeling

Reliability is perhaps one of the most important attributes of systems. A[lmost] all specifications for systems mandate that certain values for reli[a]bility be achieved and in some way proved. We have seen in the previou[s] sections that reliability can be determined experimentally if a set of N sy[s]tems is operated over a period of time and the number of systems that fa[il] during that time period is recorded. One problem with the experimental ap[proach] is the number of systems that would be required to achieve a level [of] confidence in the experimental results. This is particularly a problem whe[n] costs limit the number of systems that can be built. For example, the spa[ce] shuttle program could not afford to build 1000 of its on-board processi[ng] systems such that reliability could be experimentally verified.

A second problem with the experimental approach is the time requir[ed] to run such experiments. Many systems today are being designed to achie[ve] reliabilities of 0.9_7, or higher, after ten hours of operation. Using the exp[o]nential failure law, a reliability of 0.9_7 corresponds to a failure rate of 10 failures per hour. Therefore, on the average, we would have to wait appro[xi]mately 100 million hours, or approximately 11,416 years for the first failu[re] to occur. Clearly, we need alternatives to the experimental approach.

The most popular reliability analysis techniques are the analytical a[p]proaches. Of the analytical techniques, combinatorial modeling and Mark[ov] modeling are the two most commonly used approaches.

4.3.1 Combinatorial Models

Combinatorial models use probabilistic techniques that enumerate the d[if]ferent ways in which a system can remain operational. The probabilities [of] the events that lead to a system being operational are calculated to form [an] estimate of the system's reliability.

The reliability of a system is generally derived in terms of the reliabi[li]ties of the individual components of the system. The two models of syste[ms] that are most common in practice are the series and the parallel. In a **seri[es] system**, each element of the system is required to operate correctly for t[he] system to operate correctly. In a **parallel system**, on the other hand, o[nly] one of several elements must be operational for the system to perform [its] functions correctly.

In practice, systems are typically combinations of series and paral[lel] subsystems. Once we have discussed both the series and parallel structur[es]