NETWORK CONNECTIVITY FOR PERMANENT, TRANSIENT, INDEPENDENT, AND CORRELATED FAULTS

Allan L. White, NASA Langley, Hampton, Virginia
Courtney Sicher, Clemson University, Clemson, South Carolina
Courtney Henry, Hampton University, Hampton, Virginia

Abstract

This paper develops a method for the quantitative analysis of network connectivity in the presence of both permanent and transient faults. Even though transient noise is considered a common occurrence in networks, a survey of the literature reveals an emphasis on permanent faults. Transient faults introduce a time element into the analysis of network reliability. With permanent faults it is sufficient to consider the faults that have accumulated by the end of the operating period. With transient faults the arrival and recovery time must be included. The number and location of faults in the system is a dynamic variable. Transient faults also introduce system recovery into the analysis.

Introduction

The goal is the quantitative assessment of network connectivity in the presence of both permanent and transient faults. The approach is to construct a global model that includes all classes of faults: permanent, transient, independent, and correlated. A theorem is derived about this model that give distributions for (1) the number of fault occurrences, (2) the type of fault occurrence, (3) the time of the fault occurrences, and (4) the location of the fault occurrence. These results are applied to compare and contrast the connectivity of different network architectures in the presence of permanent, transient, independent, and correlated faults. The examples below use a Monte Carlo simulation, but the theorem mentioned above could be used to guide fault-injections in a laboratory.

Network performance in the presence of transients has been extensively studied. In [I], the authors stress the importance of considering both permanent and transient faults in a dependability analysis of a network. They conduct a failure-modes-and-effects analysis using 2080 transient fault injections in the host interface of a selected network. In [2], the authors propose a design to provide a transparent self-healing network that handles both permanent and transient faults. The goal is a system that continues to meet hard deadlines in the presence of fault occurrence. In [3], the authors describe an embedded system for real time control that uses error detection and cyclic operation to guard against permanent and transient faults. The goal is a system with a very high safety level. In [4], the authors present a comparative analysis of transient fault-tolerant techniques including end-to-end, node-by-node, and stochastic communication, but there is no quantitative assessment of reliability. In [5], the author studies masking the effects of transient faults. In [6], the authors study fault-tolerant cache coherence protocols that ensure the correct execution of programs when not all messages are correctly delivered.

The major difference between this paper and the papers referenced above is that this paper develops a method of quantitative assessment, but the results of this paper integrate well with the previous advances in the field. The results of a failure-modes-and-effects analysis (as in [l]) would be used as a model for system recovery that is a part of the quantitative assessment. This assessment method would be used to check the effectiveness of various proposed designs (as in [2] and [3]). The theorems below offer the possibility of putting this type of work (as in [4] and [5]) on a quantitative basis. More efficient methods (as in [6]) could be used to efficiently apply the methods below to larger networks.

The next section introduces a global fault model and proves the basic theorem about fault occurrence. The following section demonstrates fault occurrence can be considered in terms of three, independent, well-known distributions. The following section considers two complex networks and computes their connectivity under several fault conditions.

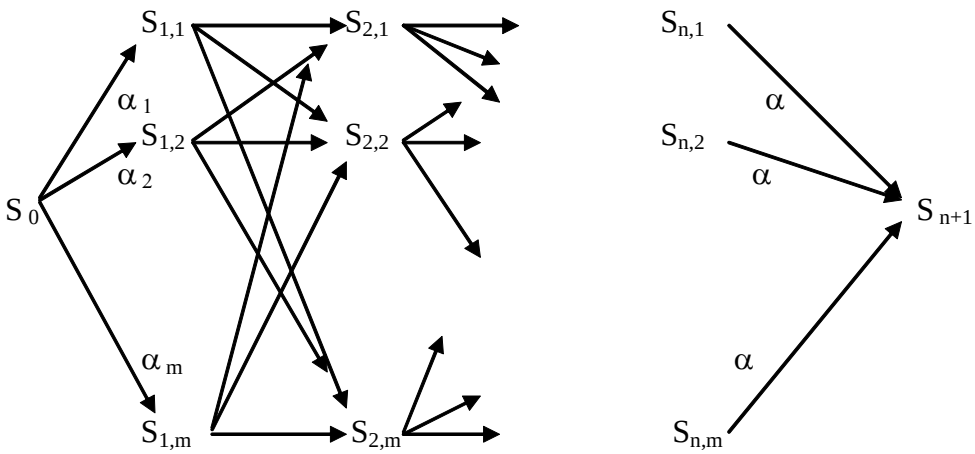

Figure 1. Global Fault Occurrence Model

The Global Fault Model

Faults are either permanent or transient. They are either independent or correlated in space or time or both. Suppose a certain system operating in a certain environment has m types of fault events, and event type j occurs at rate α_j. Let $\alpha = \alpha_{1+\ldots+}\alpha_m$. The global model for precisely n faults occurring during an operating period is given in figure 1, which displays all possible sequences of n fault events The global model for precisely n faults occurring during an operating period is given in figure 1, which displays all possible sequences of n fault events

The fan out from state S_0 includes all m classes of faults. The fan out from all the intermediate states is the same as the fan out from state S_0. There are n+2 columns beginning with column 0 and ending with column n+l. The first and last columns have a single row. The number of faults that have occurred corresponds to the column number. The process begins in state S_0 and ends in one of the states in the nth column. The final state S_{n+1} is included because the specification that precisely n faults have occurred (during the operating period) is expressed mathematically by requiring that the process reaches the nth column during the operating period, but the transition into S_{n+1} does not occur until after the operating period.

Obviously, if figure 1 is the fault occurrence model, then a device can collect several faults during an operating period. This phenomenon must be suitably interpreted during a simulation or laboratory fault-injection. If a component has already collected a permanent fault, then subsequent fault occurrences can be ignored. If a component has had previous transient fault occurrences, then a permanent fault renders it permanently faulty. Multiple transient fault occurrences are isolated if the component recovers from the previous transient before the next occurs. In a simulation or fault-injection experiment, an isolated transient is treated as if the others had not occurred. If transient faults occur while others are in the system, the simulation or experiment will have to observe system recovery from all the transient faults present.

The derivation of the formula for fault occurrence requires some bookkeeping. Suppose k_i faults of type i (i = 1,..., m) with $n = k_1 + \ldots + k_m$ have occurred in some specified order, and let β_j be the rate of the jth fault (j=1,..., n). We have,

$$\beta_1\,\beta_2\cdots\beta_n = (\alpha_1)^{k_1}(\alpha_2)^{k_2}\cdots(\alpha_m)^{k_m}$$

$$n = k_1 + k_2 + \ldots + k_m \tag{1}$$

and $\alpha = \alpha_1 + \ldots + \alpha_m$. The probability that the n designated faults occur in the designated order and that the jth fault occurs before time s_j is given by the convolution integrals below.

$$\int_0^{s_1} \beta_1\, e^{-\beta_1 t_1}\, e^{-(\alpha-\beta_1)t_1} \int_{t_1}^{s_2} \beta_2\, e^{-\beta_2 t_2}\, e^{-(\alpha-\beta_2)t_2} \vdots \int_{t_{n-1}}^{s_n} \beta_n\, e^{-\beta_n t_n}\, e^{-(\alpha-\beta_n)t_n} \int_{T-t_n}^{\infty} \alpha\, e^{-\alpha t_{n+1}}\, dt_{n+1}\, dt_n \ldots dt_2\, dt_1$$

$$= \beta_1\,\beta_2\cdots\beta_n\, e^{-\alpha T} \int_0^{s_1}\int_{t_1}^{s_2}\ldots\int_{t_{n-1}}^{s_n} dt_n \ldots dt_2\, dt_1$$

$$= (\alpha_1)^{k_1}(\alpha_2)^{k_2}\ldots(\alpha_m)^{k_m}\, e^{-\alpha T} \int_0^{s_1}\int_{t_1}^{s_2}\ldots\int_{t_{n-1}}^{s_n} dt_n \ldots dt_2\, dt_1 \tag{2}$$

This last expression is the occurrence probability for precisely n faults of a specified type and location in a specified order at specified times.

Expression (2) does not depend on the specified order of the faults. Since any ordering yields the same probability, all orderings are equally likely. Hence, the occurrence probability for precisely n faults of specified type and location at specified times in any order whatsoever is

$$\frac{n!}{k_1!\,k_2!\cdots k_m!}(\alpha_1)^{k_1}(\alpha_2)^{k_2}\ldots(\alpha_m)^{k_m}\, e^{-\alpha T} \int_0^{s_1}\int_{t_1}^{s_2}\ldots\int_{t_{n-1}}^{s_n} dt_n \ldots dt_2\, dt_1 \tag{3}$$

Interpretation of the Fault Occurrence Formula

This section shows how to choose the number, occurrence time, type, and location of the faults for a trial (representing one operating period) in the experiment based on the result in the previous section. The first subsection presents a standard result, and the next three sections present the three distributions used to interpret the result in the previous section.

Competing Constant Rate Events

Suppose there are m events each with rate α_i as depicted in figure 2. These events represent all possible fault occurrences.

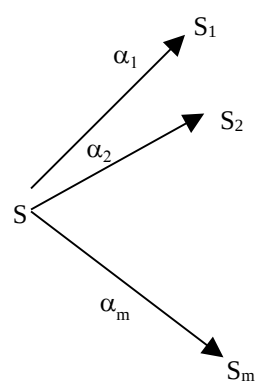

Figure 2. The Fan-Out for Fault Occurrence

For the model in figure 2, the probability that event i has occurred given some event has occurred is [7]

$$\frac{\alpha_i}{\alpha_1 + \ldots + \alpha_m} \tag{4}$$

The Poisson distribution

The Poisson distribution is a renewal process that occurs at a constant rate. The model with rate β is given in figure 3.

$$S_0 \xrightarrow{\beta} S_1 \xrightarrow{\beta} \cdots \xrightarrow{\beta} S_{k-1} \xrightarrow{\beta} S_k \xrightarrow{\beta} S_{k+1} \xrightarrow{\beta}$$

Figure 3. The Poisson Renewal Process

For the model in figure 3, the probability of being in state k at time T is [7]

$$\frac{(\beta T)^k \, e^{-\beta T}}{k!} \qquad (5)$$

As mentioned in the section above, the fault injection procedure will have to be adapted to the assumption of a constant rate used by the Poisson process. If the system removes failed components, the failure rate of the system does not remain constant. One method of handling this is to treat the removed components as virtual components. This means that the component is theoretically subject to later fault injections, but in practice these faults will not be injected if the system has already removed the component. If the system has not yet removed the faulty component, then the second fault can be injected into the same component. This double injection checks that the occurrence of a second fault does not interfere with the detection and removal of a faulty component

The Ordered Uniform Distribution

Choose a sample of size n ($x_1, x_2, \ldots, x_n$) from the uniform distribution on the interval [0 T]. Order it as $x_{(1)} \le x_{(2)} \le \ldots \le x_{(n)}$. The distribution

$$\text{Prob}\left\{ x_{(1)} \le s_1, x_{(2)} \le s_2, \ldots, x_{(n)} \le s_n \right\}$$

$$= \frac{n!}{T^n} \int_0^{s_1} \int_{t_1}^{s_2} \cdots \int_{t_{n-1}}^{s_n} dt_n \cdots dt_2 \, dt_1 \qquad (6)$$

is called the ordered uniform distribution [7].

The Mutinomial Distribution

Suppose we sample with replacement from a population with m classes of objects. Suppose the probability of choosing an object from class i is p_i . For a sample of size n the probability of choosing k_i objects from class i (i=1,...,n) is [7].

$$\frac{n!}{k_1! \, k_2! \cdots k_m!} \, p_1^{k_1} p_2^{k_2} \cdots p_m^{k_m} \qquad (7)$$

In particular, if the class of objects is the set of faults given in figure 1 then the probability of k_j faults of type j occurring given n faults have occurred is given by the expression

$$\frac{n!}{k_1! \, k_2! \cdots k_m!} \left(\frac{\alpha_1}{\alpha}\right)^{k_1} \left(\frac{\alpha_2}{\alpha}\right)^{k_2} \cdots \left(\frac{\alpha_m}{\alpha}\right)^{k_m} \qquad (8)$$

where $\alpha = \alpha_1 + \cdots + \alpha_m$. In the formulas (8) and (9) some of the k_j's can be zero.

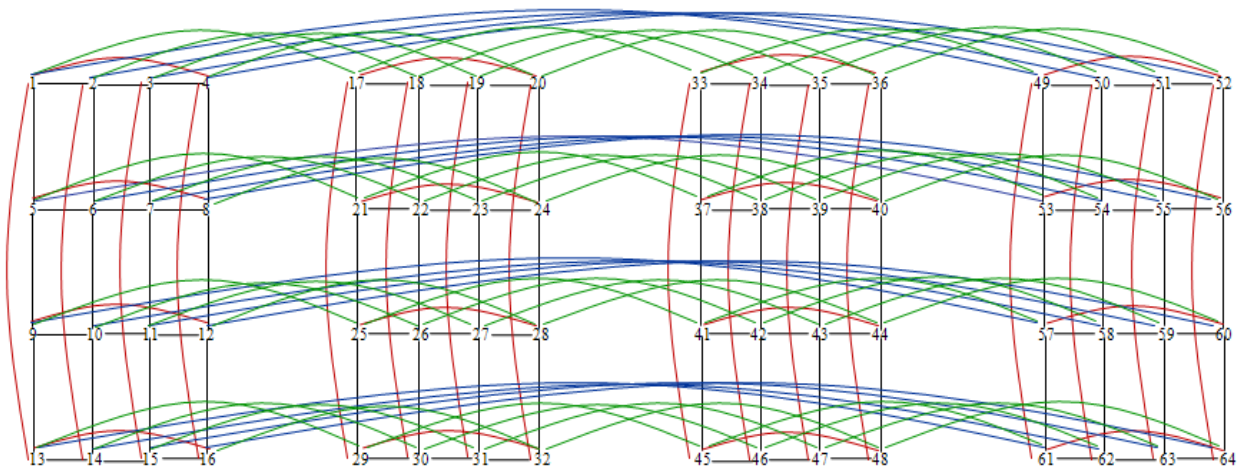

Figure 4. A Four-by-Four-by-Four Planar Array

The Fault Injection Procedure

The expression (3) is the probability that precisely n faults of some specified types have occurred at some specified times in an operating period of the system. The expression (3) (and hence the probability) is algebraically equivalent to a product of three probability distributions as given in expression (9). The multiplicative property implies these three distributions act independently.

$$\left[\frac{(\alpha T)^n \, e^{-\alpha T}}{n!}\right] \left[\frac{n!}{T^n}\int_0^{s_1}\int_{t_1}^{s_2}\cdots\int_{t_{n-1}}^{s_n} dt_n \ldots dt_2 \, dt_1\right] \left[\frac{n!}{k_1!\,k_2!\cdots k_m!}\left(\frac{\alpha_1}{\alpha}\right)^{k_1}\left(\frac{\alpha_2}{\alpha}\right)^{k_2}\cdots\left(\frac{\alpha_m}{\alpha}\right)^{k_m}\right] \quad (9)$$

The three distributions (in order) are the Poisson renewal process, the ordered uniform, and the multinomial distribution. Since they act independently, the faults for any trial (representing one operating period) can be chosen in the following three steps.

(1) The number of faults is given by the Poisson with rate the sum of all fault occurrence rates.
(2) The occurrence times are given by the ordered uniform.
(3) At each occurrence time, the location and type of fault is chosen by a random sample (without replacement, according to occurrence probability) from the set of faults.

Connectivity Example

We consider a two-dimensional and a three-dimensional planar array. The examples have the same number of nodes, an eight-by-eight and a four-by-four-by-four, although they have a different number of links: 128 and 192, respectively. A four-by-four-by-four planar array is depicted in figure 4. The connectivity for a two-dimensional four-by-four is given by the first 16 nodes and the links connecting them.

Independent node and link failures can be either permanent or transient. For this example, faults correlated in space are all transient faults, and they are node-centric with a correlated fault bringing down a node, a group of adjacent nodes, and the connecting links. There are two levels of severity for faults correlated in space. For a two-dimensional array, a fault of severity level one brings down a node and the four adjacent nodes along with the associated 14 links. A fault of severity level two brings down a node and the eight adjacent nodes along with the associated 24 links. These faults are shown in figures 5 and 6. Level one and two faults in a three-dimensional array bring down 7 and 27 nodes respectively along with the associated links.

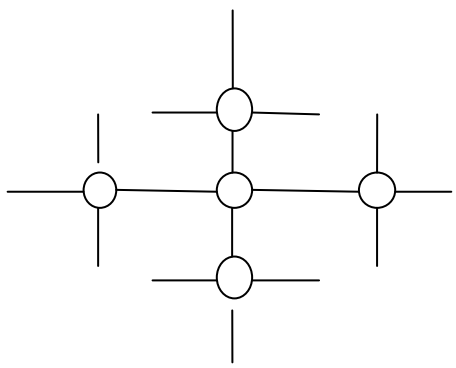

Figure 5. Correlated Fault of Severity One

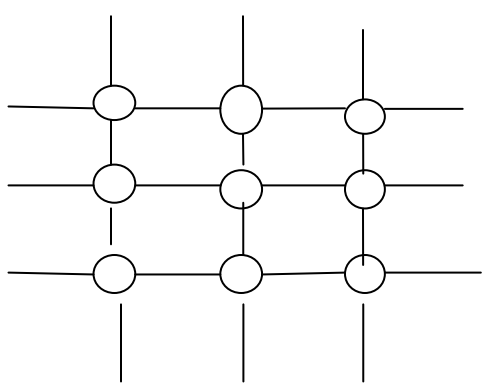

Figure 6. Correlated Fault of Severity Two

We also consider faults correlated in time, and for this study, they are treated as arising from external phenomena that also cause correlation in space. For these computations, we assume they are transient faults of severity level one that repeatedly hit the same collection of nodes. The number of repeated hits is given by a Poisson distribution, and their occurrence times are given by an exponential distribution.

We give the systems a long operating time of more than 10 years. The nodes are assumed simple (perhaps computers on a chip) and highly reliable. The node transient failure rate is a magnitude larger than the node permanent failure rate. The links are simpler pieces of hardware, but they are more exposed. Their permanent failure rate is equal to the node permanent failure rate, and their transient failure rate is an order of magnitude greater than their permanent rate. The correlated-in-space rates are proportional to the number of links in the system. The correlated-in-time faults average two additional hits, and the additional hits occur quickly. In this first study, the transient recovery times for all types of transient faults are equal.

The operating time and failure rates for the computations are

Operating time = 100,000 hours

Node permanent = 1e-6/hour

Node transient = 1e-5/hour

Link permanent = 1e-6/hour

Link transient = 1e-5/hour

Correlated level one for two dimensional

= 1.28e-4/hour

Correlated level two for two dimensional

= 6.4e-5/hour

Correlated level one for three dimensional

= 1.92e-4/hour

Correlated level two for three dimensional

= 9.6e-5/hour

Occurrence rate for two-dimensional

time dependent faults = 6.4e-5/hour

Occurrence rate for three-dimensional

time dependent faults = 9.6e-5/hour

Possion parameter for time dependent faults = 2

Exponential rate for time dependent faults = 3.6e+4/hour

Recovery time for transients = 1 second.

Table 1 gives the failure frequencies and 95% confidence intervals for two and three dimensional planar arrays when the number of trials is 10,000 for each fault scenario.

Summary and Further work

This paper derives a result about fault occurrence in a global fault model where the faults appear at a constant rate. A fault can be permanent or transient. It can be independent or correlated in space or time or both. Additional analysis shows this basic result on fault occurrence can be written as the product of three, well-known distributions, which renders it convenient for fault injection experiments. The examples demonstrate applying the result to all possible classes of faults: permanent or transient and independent or correlated in space or time or both.

In addition to connectivity, there are a number of other elements that can now be examined quantitatively. There is the study of the detection and identification of faults with rerouting of messages. There is the comparison of architectures and protocols in different operating environments, and there is sensitivity analysis to guide the design of systems and the gathering of field data

Table 1. Statistics of Fault Injection

Fault Types Present	**Eight-by-Eight**	**Four -by- Four -by- Four**
Node perm	52 [38, 66]	0 [0, 3]
Node tran	0 [0, 3]	0 [0, 3]
Node perm Node tran	431 [391, 471]	4 [0, 8]
Link perm	42 [29, 55]	1 [0, 3]
Link tran	0 [0, 3]	0 [0,3]
Link perm Link tran	470 [429, 511]	2 [0, 5]
Node per Node tran Link perm Link tran	3674 [3579, 3768]	177 [151, 203]
Node perm Node tran Link perm Link tran Correlated type 1	5881 [5784, 5978]	382 [344, 420]
Node perm Node tran Link perm Link tran Correlated type 1 Correlated type 2	5905 [5808, 6001]	470 [428, 511]
Node perm Node tran Link perm Link tran Correlated type 1 Correlated type 2 Time correlated	6534 [6441, 6627]	641 [593, 689]

References

1. D. T. Stott, M.-C. Hsueh, G. L. Ries, and R. K. Iyer, "Dependability Analysis of a Commercial High-Speed Network," Twenty-Seventh Annual International Symposium on Fault-Tolerant Computing, June 1997.

2. T. Sauter and N. Kero, "Majority Manchester Decoding for Active Redundant Data Transmission," IEEE Symposium on Computers and Communications, June 1999.

3. B. Rostamzadeh, H. Lonn, R. Snedsbol, and J. Torin, "DACAPO: A Distributed Computer Architecture for Safety-Critical Control Applications," Intelligent Vehicle '95 Symposium, September 1995.

4. Muhammad Ali, Awais Adnan, "Comparative Analysis of Transient-fault Tolerant Schemes for Network on Chips," World Academy of Science, Engineering and Technology, Volume 43, 2008, pp. 386-391.

5. N. J. Wang, "Characterizing the effects of transient faults on a high-performance processor pipeline," 2004 International Conference on Dependable Systems and Networks, June 2004, pp. 61-70.

6. Ricardo Fernandez-Pascual, Jose M. Garcia, Manuel E. Acacio, and Jose Duato, "Dealing with Transient Faults in the Interconnection Network of CMPs at the Cache Coherence level," IEE Transactions on Parallel and distributed Systems, Vil. 21, No. 8, August 2010, pp. 1117-1131.

7. Samuel Karlin and Howard Taylor, *A First Course in Stochastic Processes*, Academic Press, New York, 1975