

MODELING OF EXTERNALLY-INDUCED/COMMON-CAUSE FAULTS IN FAULT-TOLERANT SYSTEMS

Hagbae Kim and Kang G. Shin
Real-Time Computing Laboratory
Department of Electrical Engineering and Computer Science
The University of Michigan
Ann Arbor, MI 48109-2122.

Abstract: Modeling fault behaviors such as fault occurrences and active/benign durations is an essential step to the design and evaluation of fault-tolerant controller computers.

We use a beta-binomial distribution to model fault occurrences both in the presence and in the absence of environmentally-induced (thus common-cause) faults. A multinomial distribution is used to model fault active durations. The proposed model is validated by testing it against the data generated by a simulation program that mimics a common-cause fault environment. The model is then applied to the determination of an optimal time-redundancy recovery method for EMI-induced failures in an N -modular redundant controller computer, demonstrating its utility and power.

1 Introduction

It is commonplace to use digital computers for the control and monitoring of safety-critical real-time systems, such as avionics systems and nuclear reactors, where controller-computer failures may be life-threatening. Massive hardware redundancy has been widely used for such controller computers to meet the ultra reliability requirement. Three to five modules are often used to execute independently the copies of each critical task and the execution results are then voted on to mask the effects of the faulty modules. This method, called *N-Modular Redundancy* (NMR), is effective only when faults occur independently on different modules.

Most of the popular fault models used for the design and analysis of fault-tolerant systems and the evaluation of system reliability cover independently-occurring faults. However, all modules of an NMR system are subjected to the same environment (temperature, humidity, and electro-magnetic interferences) and often share the same clock and power supply, for which the assumption of independence of fault behaviors in different modules

no longer holds. The faults resulting from environmental disruptions such as electromagnetic interferences (EMI), radiation, temperature, or vibration are likely to be transient, because (1) the main effects of the abnormal operating environment may be functional error modes without actual component damages and (2) the adversary environmental conditions are in general temporary. The harsh environment resulting from EMI will affect the entire system and induce coincident, or common-source, faults in multiple modules of the NMR system. It is therefore important to develop a model (for fault occurrences and durations) that integrates *both* the correlated faults (due mainly to environmental disruptions) and independently-occurring faults. One can then use this fault model for more realistic design and evaluation of fault-tolerant controller computers.

Several researchers proposed models for correlated faults in hardware or software. In [4], a simple model of correlated faults treated the case of two modules using linear correlation but did not include coincident faults induced by a common source. The author of [5] considered correlated faults in multiversion programming for fault-tolerant software by using an appropriate intensity (or distribution) function. An example in [1] treated the problem of failure patterns in separately-located computers to investigate the effects of a common environment on reliability.

To develop a more effective model integrated for both correlated and independent faults in redundant hardware systems, we classify the operating conditions of a controller computer to be normal or abnormal, depending upon the occurrence of harsh environment like EMI. We propose a (compressed) beta-binomial distribution — whose range is less than one by excluding independent fault sources and using the parameters of this distribution one can measure the level of correlation — to model fault occurrences under both the operating conditions. The beta-binomial distribution was used to model the correlated incidence in the households of an infectious disease [2]. We also represent fault durations by a multinomial distribution, which discriminates the active faults

The work reported was supported in part by the Office of Naval Research under Grant N00014-91-J-1115 and by the NASA under Grant NAG-1-1120.

according to their durations.

The proposed model of fault occurrences is compared against a binomial distribution by fitting the data collected from a simulation program that imitates a harsh environment. (Unfortunately, there are few real data on correlated faults available in the open literature.) This model is then applied for the analysis and design of a fault-tolerant controller computer and the evaluation of system reliability, which are shown by an example NMR controller computer: (i) assessment of the reliability of a static-redundancy system and (ii) computation of the probability of successful recovery from faults when applying time redundancy.

Section 2 addresses the generic properties of faults and their effects on a fault-tolerant controller computer. Section 3 presents the models for fault occurrences and durations by using the (compressed) beta-binomial distribution and the multinomial distribution, respectively. In Section 4, we justify the proposed distribution models with the simulated data of a harsh environment. An example of applying the proposed model is also presented there. The paper concludes with Section 5.

2 Generic Properties of Faults

Faults are classified into two types, internal and external, depending on the location (inside or outside the system) of their origin. Internal faults represent the malfunctioning/damaged parts of a system that induce errors. These faults occur due to physical defects during manufacture or due to component aging. Most internal faults are likely to be permanent or intermittent because the effects of physical defects (e.g., broken, short, or loose connections) tend to persist or cycle between active and inactive states. The sources of internal faults are usually independent, so their occurrences in different modules are assumed to be independent.

External faults, on the other hand, result from environmental interferences or disruptions, such as electromagnetic perturbation, radiation, temperature, or vibration. These external faults are transient because disruptive environmental conditions are temporary and may cause functional error modes without actual component damages. The harsh environment with lightning, HIRF (high intensity radio frequency fields), or NEMP (nuclear electromagnetic pulses) generally affects the entire system, regardless of the degree of redundancy used, thereby making the occurrences and durations of external faults dependent on one another.

In addition to the occurrence rate of external faults, their active durations are also important in characterizing them, because most external faults are known to be transient. Time redundancy can be applied effectively

against such faults using the knowledge of their durations. Any critical computation corrupted by transient faults can be recovered by re-executing the contaminated part of the computation after waiting a sufficient time for the transient faults to die away but soon enough to meet the application timing constraints.

3 Modeling of Fault Behaviors

A controller computer consists of multiple modules. Each module is a self-contained entity with input/output from/to others. Let N and N_f be the numbers of all modules and faulty modules, respectively, in the controller computer. Since we are mainly interested in the number of faulty modules, we do not consider how many faults are active in a single module. In other words, if a module contains at least one active fault regardless of the number of faults in that module, it contributes one faulty module to the calculation of N_f .

3.1 The Fault-Occurrence Model

As mentioned earlier, the operation of a controller computer is either normal or abnormal, depending on the absence or presence of EMI. EMI is generally characterized by a long latent period followed by a relatively short period of presence. We assume that EMI arrivals follow a time-invariant Poisson process with rate λ_e and each arrival remains active for an exponentially-distributed random period with rate μ_e . Consequently, environmental disruptions arrive at an exponential rate λ_e and disappear after an active duration with mean $\frac{1}{\mu_e}$.

Under the normal operating condition, we assume that no external faults occur since their occurring frequency is much smaller than that of internal faults. Internal faults were assumed to behave (occur/persist) independently in different modules. Thus, fault occurrences under the normal operating condition are modeled by a binomial distribution with constant mean p_0 . That is, the probability of k faulty modules is:

$$P_n(N_f = k) = \binom{N}{k} p_0^k (1 - p_0)^{N-k}, \quad (3.1)$$

where the expected number of faulty modules is Np_0 .

Under the abnormal operating condition, however, some external correlated faults may occur due to the presence of EMI. Since the sources of faults are still active under the abnormal operating condition, the number of faulty modules (N_f) is certainly larger (in a probabilistical sense) than that without EMI. Since N_f depends on the level of correlation — determined by the intensity of EMI and the structural/material properties of the system — it is modeled by a random intensity parameter p with a probability density function $f(p)$. That is, the

number of faulty modules becomes Np with probability $f(p)$. Thus, the expected fraction of faulty modules in the system is $\bar{p} = \int pf(p)dp$. Under the normal operating condition, the probability density function of p is simply represented by a unit impulse at p_0 , i.e., $f(p) = \delta(p-p_0)$. Thus, the random intensity parameter p of the abnormal operating condition must have the following properties: (i) $p_0 < p < 1$, (ii) \bar{p} increases with stronger EMI and weaker shielding, (iii) p spreads out over a wider range as inter-module correlation increases. Properties (ii) and (iii) can be satisfied by randomizing the parameter p of the binomial distribution and by using a beta distribution [3] for the probability density function of p :

$$f(p) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1} (1-p)^{\beta-1}, \quad 0 < p < 1, \quad \alpha, \beta > 0, \quad (3.2)$$

where Γ is the Gamma function.

Property (i) is also satisfied by modifying the range of the random variable (intensity parameter) in Eq. (3.2). The probability density function with the required compressed-range ($p_0 < p < 1$) is:

$$f(p) = \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \left(\frac{p-p_0}{1-p_0} \right)^{\alpha-1} \left(1 - \frac{p-p_0}{1-p_0} \right)^{\beta-1}. \quad (3.3)$$

Then, the resulting probability mass function of the number, N_f , of faulty modules is a (compressed) beta-binomial distribution [2]:

$$P_a(N_f = k) = \int_0^1 \binom{N}{k} \left(\frac{p-p_0}{1-p_0} \right)^k \left(1 - \frac{p-p_0}{1-p_0} \right)^{N-k} \frac{\Gamma(\alpha+\beta)}{\Gamma(\alpha)\Gamma(\beta)} \left(\frac{p-p_0}{1-p_0} \right)^{\alpha-1} \left(1 - \frac{p-p_0}{1-p_0} \right)^{\beta-1} dp, \quad (3.4)$$

where $0 \leq k \leq N$, $p_0 < p < 1$, and $\alpha, \beta > 0$.

This equation can be rewritten by using the beta function as:

$$P_a(N_f = k) = \frac{\binom{N}{k} B(\alpha+k, N+\beta-k)}{B(\alpha, \beta)}, \quad (3.5)$$

where

$$B(\alpha, \beta) = \int_0^1 p^{\alpha-1} (1-p)^{\beta-1} dp = \frac{\Gamma(\alpha)\Gamma(\beta)}{\Gamma(\alpha+\beta)}.$$

Accordingly, the mean and variance of this distribution are given by:

$$\begin{aligned} E(N_f) &= N\pi, \\ \text{Var}(N_f) &= N\pi\chi(1-p_0)^2 \frac{1+N\theta}{1+\theta}, \end{aligned} \quad (3.6)$$

where

$$\pi = (1-p_0) \frac{\alpha}{\alpha+\beta} + p_0, \quad \chi = 1-\pi, \quad \text{and} \quad \theta = \frac{1}{\alpha+\beta}. \quad (3.7)$$

In Eq. (3.7), $\pi > 0$ is defined as the mean value of the binomial parameter p for each individual module (i.e., \bar{p}). χ represents the mean probability of a module not being faulty. When $\theta \rightarrow 0$, the beta-binomial distribution becomes similar to the pure binomial distribution ($\theta = 0$), in which there is no correlation among modules. Thus, θ can be used to measure the level of inter-module correlation.

We proposed a probability distribution model for the number of faulty modules by using the beta-binomial distributions, whose parameter density functions are a unit impulse function (deterministic) and a beta distribution for the normal and abnormal operating conditions, respectively. Now, a proper model of fault occurrences over the whole operating period is obtained by combining the two conditions with the model of EMI behaviors:

$$\begin{aligned} P(N_f = k) &= P_n(N_f = k)P(\text{norm}) \\ &\quad + P_a(N_f = k)(1 - P(\text{norm})). \end{aligned} \quad (3.8)$$

$$P(\text{norm}) = e^{-\lambda_e t} + \int_0^t \lambda_e e^{-\lambda_e x} (1 - e^{-\mu_e(t-x)}) dx,$$

where t is the time interval from the beginning to the instant of interest.

3.2 The Fault Duration Model

The duration of a fault begins with the moment of physical or logical deviations from the specified values. The duration of a transient fault measured from that moment is generally a random variable with a distribution governed by the source of the fault. To distinguish permanent faults from transient faults, we define a *threshold duration*, i.e., all long-lasting transient faults beyond the threshold are treated as permanent faults. All faults of duration longer than this threshold are treated as permanent by the recovery methods in a given fault-tolerant controller computer.

To model fault durations we first consider a transient fault in one module and then extend this resulting model to include the effects of correlated faults in multiple (redundant) modules. A control mission can be divided into a number of mission segments/phases. The active duration of a fault can also be represented by an integer multiple of the unit time interval like a sampling period. This integer representing the fault duration is a random variable. Let t_d , T_s , and T_t be the fault duration, the unit time, and the threshold duration, respectively, and let M be the largest integer such that $MT_s \leq T_t$. Then, the probability that the integer multiple is i ($0 \leq i \leq M$) is derived by using the probability distribution,¹ $F_a(t)$,

¹The exponential distribution is the most commonly-encountered model for the active duration of a transient fault, i.e., $F_a(t) = 1 - e^{-\mu_a t}$ with mean active duration $\frac{1}{\mu_a}$.

of the active duration ($1 \leq j \leq (M-1)$):

$$\begin{aligned} q_0 &= Pr[0 \leq t_d < T_s] = F_a(T_s) \\ q_j &= Pr[jT_s \leq t_d < (j+1)T_s] \\ &= F_a((j+1)T_s) - F_a(jT_s) \\ q_M &= Pr[t_d \geq MT_s] = 1 - F_a(MT_s), \end{aligned} \quad (3.9)$$

where the probability of a permanent fault is q_M .

Now, we use a multinomial distribution to model fault durations in multiple modules. Let N_i be the number of faulty modules with an active duration, t_d , such that $(i-1)T_s \leq t_d < iT_s$, among the N_f faulty modules. If there is no correlated fault, the probabilities of all possible durations of N_f (independent) faulty modules are obtained by:

$$P_d(N_0 = k_0, N_1 = k_1, \dots, N_M = k_M) = \frac{N_f!}{k_1!k_2! \dots k_M!} q_0^{k_0} q_1^{k_1} \dots q_M^{k_M}, \quad (3.10)$$

where $\sum_{i=0}^M k_i = N_f$ and $\sum_{i=0}^M q_i = 1$.

To include the effects of external (correlated) faults, we must modify Eq. (3.10). If the present faults were caused by the same source, they are unlikely to become inactive until the source disappears. We assume that all faults disappear according to a certain probability distribution model, i.e., the duration of faults is distributed with a *pdf* such as an exponential distribution. We can also assume independence among (active) faults after the common source of faults (EMI) disappears.

If T_e is defined as the active duration of EMI, which has a distribution function F_e , the probabilities of all possible durations of N_f faulty modules caused by the EMI are modified from Eq. (3.10) to:

$$\begin{aligned} &P_d(N_0 = k_0, \dots, N_M = k_M) = \\ &\sum_{j=1}^{j=M} Pr(N_0 = k_0, \dots, N_M = k_M | (j-1)T_s \leq T_e < jT_s) \\ &Pr[(j-1)T_s \leq T_e < jT_s] + Pr(N_0 = k_0, N_1 = k_1, \dots, \\ &\dots, N_M = k_M | T_e \geq MT_s) Pr[T_e \geq MT_s] \\ &= \sum_{j=1}^{j=M} \left(F_e(jT_s) - F_e((j-1)T_s) \right) \frac{N_f!}{k_j!k'_{(j+1)}! \dots k'_M!} \\ &p_0^{k_0} p_1^{k'_1} \dots p_{(M-j)}^{k'_M} + \left(1 - F_e(MT_s) \right) \\ &Pr(N_0 = k_0, \dots, N_M = k_M | T_e \geq MT_s), \end{aligned} \quad (3.11)$$

where

$$\begin{aligned} &Pr(N_0 = k_0, \dots, N_M = k_M | T_e \geq MT_s) = \\ &P_d(0, 0, \dots, N_M = N_f) = 1, \text{ and } \sum_{i=j}^M k'_i = N_f. \end{aligned}$$

$N_f \setminus$	SD	BD	BBD
0	0.99088	0.990862	0.99088
2	$5.5e-5$	$3.1474e-5$	$5.42667 \times e-5$
4	0	$1.35397e-10$	$9.22787e-10$
6	0	$9.40008e-17$	$6.18479e-15$
8	0	$1.04884e-23$	$1.13426e-20$

Table 1: $\hat{\pi} = 1.019e-3$ and $\hat{\theta} = 4.81e-4 \approx 0$ when $\mu = 0$: SD=Simulation Data, BD=data of the fitted Binomial Distribution (with parameter $\hat{\pi}$), BBD=data of the fitted Beta-Binomial Distribution (with parameters $\hat{\pi}$ and $\hat{\theta}$).

$N_f \setminus$	SD	BD	BBD
0	0.990875	0.990689	0.990879
2	$2.05e-4$	$3.85728e-5$	$2.17055e-4$
4	0	$1.46012e-10$	$1.4066e-7$
6	0	$1.05278e-16$	$5.50459e-11$
8	0	$1.21996e-23$	$7.22844e-15$

Table 2: $\hat{\pi} = 1.039e-3$ and $\hat{\theta} = 5.278e-3$ when $\mu = 5e-3$.

In the above equation, p_i ($0 \leq i \leq M-j$) is obtained similarly to Eq. (3.9).

$$\begin{aligned} p_i &= Pr[iT_s \leq t_d < (i+1)T_s] \\ &= F_{ae}((i+1)T_s) - F_{ae}(iT_s) \text{ for } 0 \leq i \leq M-j-1, \end{aligned}$$

where $\sum_{i=0}^{M-j} p_i = 1$ and F_{ae} is the probability distribution of active durations of faults caused by EMI after EMI disappeared. From Eqs. (3.10) and (3.11), we can obtain the probabilities of all possible durations of N_f faults involving both internal and external faults.

4 Validation and Application

In this section, the proposed distribution of fault occurrences is justified by comparing fitted distributions with the data collected from a simulation program for both the beta-binomial distribution and the binomial distribution. An example of applying the proposed model is also presented.

4.1 Fitting the Model to Simulation

Due to the unavailability of actual data on common-cause faults, we use the data on the number of faulty modules² collected from a simulation program which imitates a common-cause fault environment. When the sys-

²A fault may disappear without inducing any error/failure and a failure can be detected only after a certain interval (the fault and error latencies) following the occurrence of a fault. Since it is not feasible to directly detect faults, we approximate the number of modules in which an error/failure is detected, as the number of faulty modules.

$N_f \setminus$	SD	BD	BBD
0	0.991145	0.990723	0.991157
2	$3.95e-4$	$3.82856e-5$	$4.19986e-4$
4	$5e-6$	$1.43841e-10$	$1.30365e-6$
6	0	$1.02937e-16$	$2.61283e-9$
8	0	$1.1839e-23$	$1.83465e-12$

Table 3: $\hat{\pi} = 1.035e-3$ and $\hat{\theta} = 1.2553e-2$ when $\mu = 1e-2$.

$N_f \setminus$	SD	BD	BBD
0	0.991375	0.989484	0.991412
2	$1.165e-3$	$4.92023e-5$	$1.23227e-3$
4	$6e-5$	$2.37864e-10$	$5.56469e-5$
6	$5e-6$	$2.19034e-16$	$1.92125e-6$
8	0	$4.23295e-27$	$3.07243e-9$

Table 4: $\hat{\pi} = 1.174e-3$ and $\hat{\theta} = 6.1497e-2$ when $\mu = 5e-2$.

tem is composed of n modules, the probability of fault occurrences is represented by:

$$P(m_1, m_2, \dots, m_n) = P(m_1)P(m_2|m_1)P(m_3|m_2, m_1), \\ \dots, P(m_n|m_{n-1}, \dots, m_1) = \prod_{k=1}^n P(m_k|m_{k-1}, \dots, m_1),$$

where $m_i \in \{0, 1\}$ is a random variable indicating the occurrence of a fault in the i -th module. We consider a simple example for common-cause fault environments, in which fault occurrences are governed by $P(m_1) = p_0$ and

$$P(m_k|m_{k-1}, \dots, m_1) = p_0 + \sum_{i=1}^{k-1} m_i \mu \quad \text{for } 2 \leq k \leq n. \quad (4.1)$$

Under this condition of fault occurrences and $p_0 = 1e-3$, we collected the data in Tables 1–4 from simulation runs for a period of 2×10^5 , while incorporating the various levels of correlation. When $\mu = 0$ (i.e., independent fault occurrences under the normal operating condition), both the binomial distribution and the beta-binomial distribution represent the number of faulty modules with little difference. However, it is evident that the binomial distribution becomes inappropriate to model fault occurrences as μ increases, as shown in Tables 2–4. By contrast, the beta-binomial distribution fits the data quite well regardless of the level of correlation. Since one can easily see the significant improvement of the BBD fitting over the BD fitting, we do not perform any hypothesis testing like χ^2 -test.

4.2 Application of the Model

We now present a simple example to demonstrate the usefulness of this model by using an N -modular redundant (NMR) system, which will fail to form a majority at the time of voting if more than $\frac{N-1}{2}$ module outputs are




Figure 1: Reliability of a 7-MR system while varying the mission lifetime from $2e+04(\lfloor \frac{T}{T_s} \rfloor)$ to $2e+05(\lfloor \frac{T}{T_s} \rfloor)$ with the fault-occurrence model equal to (a) the BD and (b)–(e) the BBD, where (b) $\theta = 0.005$, (c) $\theta = 0.01$, (d) $\theta = 0.05$, (e) $\theta = 0.1$ with $p_0 = 1e-04$, $\pi = 1e-03$, $P(norm) = 1 - 1e-07$.

erroneous. In the example, (i) the reliability of a static redundant system is assessed, and (ii) the necessary information on fault durations is derived to determine an “optimal” time-redundancy strategy to enhance system reliability using the proposed fault model.

The probability of more than $\frac{N-1}{2}$ module faults is derived by using the fault occurrence model as:

$$Pr[N_f \geq \frac{N-1}{2}] = 1 - \sum_{k=0}^{\frac{N-1}{2}} P(N_f = k), \quad (4.2)$$

where $P(N_f = k)$ is calculated using Eq. (3.8). Thus, the NMR system reliability for a control mission lifetime T , $R(T)$, is:

$$R(T) = 1 - \left(Pr[N_f \geq \frac{N-1}{2}] \right)^{\lfloor \frac{T}{T_s} \rfloor}, \quad (4.3)$$

where $\lfloor A \rfloor$ is the smallest integer such that $\lfloor A \rfloor \geq A$, and T_s is substituted for the variable t in Eq. (3.8) to compute $P(normal)$. As shown in Fig. 1, the beta-binomial distribution is used for the fault-occurrence model, and the merit of static redundancy is significantly diminished, especially for the case of high-degree correlation (i.e., large θ).

When the timing constraints are stringent³ or the cost of spatial redundancy is not high, the faulty modules will

³One can determine the timing constraints of the controller computer by analyzing the controlled processes in the context of certain fault models [6].

be replaced with healthy ones from a large pool of spares. This increases the cost of spatial redundancy significantly and the time overhead of switching in new spares. In case the external fault durations are correlated, most of them may become inactive upon disappearance of the source of faults, and thus, re-executing part or whole of the program (retry, rollback, or restart) may recover the system from the errors caused by external faults without avoiding the premature retirement of (transient) faulty modules. The information about the fault durations is required to apply these time-redundancy recovery methods effectively. That is, re-execution must be initiated after a sufficient backoff time, during which the faults are likely to disappear, and thus, at least $\frac{N+1}{2}$ modules become nonfaulty. When the number of faulty modules is $N_f > \frac{N-1}{2}$, we can derive the probability that the number of faulty modules after a backoff time t (denoted as $N_p(t)$) by using the proposed fault duration model. Let q be the smallest integer such that $qT_s \geq t$, or $q = \lfloor \frac{t}{T_s} \rfloor$, and let K_s denote the set of all combinations of $\{k_1, k_2, \dots, k_M\}$ such that $\sum_{i=q}^M k_i > \frac{N-1}{2}$, then:

$$Pr[N_p(t) \leq \frac{N-1}{2}] = 1 - \sum_{K_s} P_d(N_1 = k_1, N_2 = k_2, \dots, N_M = k_M), \quad (4.4)$$

where $P_d(\dots)$ is calculated using Eq. (3.11). The probability of successful recovery in applying time redundancy depends on Eq. (4.4). Some numerical examples of $Pr[N_p(t) \leq \frac{N-1}{2}]$ are plotted in Fig. 2, which are derived from Eqs. (3.11) and (4.4) by varying backoff times or mean durations of fault/EMI. In this figure, one can see that the probability of successful recovery is more dependent upon the mean duration of EMI than that of fault when comparing the curves of $\mu_f = 1/5, \mu_e = 1/3$ and $\mu_f = 1/3, \mu_e = 1/5$. This is because no fault is likely to become inactive before EMI disappears.

5 Conclusion

Our model using a beta-binomial distribution (and a multinomial distribution) for fault occurrences (and fault durations) to represent fault behaviors in NMR systems differs from the previous models which usually assumed independence of modules or dealt only with two or three modules.

Our future work includes the problems of (i) developing a fault model which can be validated by analytical and experimental tools and cover sources of correlated faults other than EMI and (ii) detecting the correlated faults caused by EMI and assessing their effects. We are planning to conduct experiments in the HIRF Lab currently being set up at the NASA Langley Research Center which will shed some light on these problems.




Figure 2: Probability of successful time-redundancy recovery in a 7-MR system with $f_e(t) = \mu_e e^{-\mu_e t}$, $f_{ae}(t) = \mu_f e^{-\mu_f t}$, and $M = 20$.

Acknowledgement

The authors would like to thank Allan White and Chuck Meissner of the NASA Langley Research Center for their technical and financial assistance.

References

- [1] D. R. Cox and P. A. W. Lewis, "Multivariable point processes," in *Proc. 6th Berkeley Symp. Math. Statist. Prob.*, 3, 1972.
- [2] D. A. Griffiths, "Maximum likelihood estimator for the beta-binomial distribution and an application to the household distribution of the total number of cases of a disease," *Biometrics*, vol. 29, pp. 637-648, December 1973.
- [3] F. A. Haight, *Applied Probability*, Plenum Press, New York and London, 1981.
- [4] Y. K. Malaiya, "Linearly correlated intermittent failures," *IEEE Trans. on Reliability*, vol. R-31, no. 2, pp. 211-215, June 1982.
- [5] V. F. Nicola and A. Goyal, "Modeling of correlated failures and community error recovery in multiversion software," *IEEE Trans. on Software Eng.*, vol. SE-16, no. 3, pp. 350-359, March 1990.
- [6] K. G. Shin and H. Kim, "Derivation and application of hard deadlines for real-time control systems," *IEEE Trans. on Systems, Man, and Cybernetics*, vol. 22, no. 6, pp. 1403-1413, Nov./Dec. 1992.