Improving Availability Bounds using the Failure Distance Concept

Juan A. Carrasco
Departament d'Enginyeria Electrònica
Universitat Politècnica de Catalunya
Diagonal 647, plta. 9
08028 Barcelona, Spain
juan.a.carrasco@upc.edu

Except for formatting details, this version matches exactly the version published with the same title and authors in *Dependable Computing for Critical Applications 4*, Springer-Verlag, 1995, pp. 479–497

Abstract

Continuous-time Markov chains are commonly used for dependability modeling of repairable fault-tolerant computer systems. Realistic models of non-trivial fault-tolerant systems easily have very large state spaces. An attractive approach which has been proposed to deal with the largeness problem is the use of pruning-based methods which provide error bounds. Using results from Courtois and Semal, a method for bounding the steady-state availability has been recently developed by Muntz, de Souza e Silva, and Goyal. This paper presents a new method based on a different approach which exploits the concept of failure distance to better bound the behavior out of the non-generated state space. The proposed method yields tighter bounds. Numerical analysis shows that the improvement is typically significant.

1 Introduction

Modeling plays an important role in the design, analysis and management of fault-tolerant computer systems. These systems are characterised by exhibiting an stochastic behavior and, accordingly, probabilistic measures are used for their quantitative assessment. Many systems are seen by their users as providing service or not. For these systems, dependability measures such as the availability and the reliability are appropriate. The steady-state availability is a useful measure for repairable systems when the long-term behavior is of interest. In some cases, this measure can be computed using combinatorial techniques [1] or closed-product solution queuing networks [6]. However, in general, the dependencies introduced by lack of coverage, failure propagation, operational configurations and maintenance are such that general-purpose, state level model solution techniques are

required. Continuous-time Markov chains (CTMC's) are often used to analyse systems with these dependencies and a number of dependability/performability tools based on these models have been developed in the past (see [11] for a recent review).

Numerical analysis of CTMC dependability models is hampered by the exponential growth of the number of states with the structural complexity of the system. Systems with moderate number of components easily yield CTMC's with millions of states and more. This problem has been attacked in three directions: a) hierarchical model solution [16], b) state lumping techniques [12], and c) pruning techniques. Only the last of them has general applicability. Recently, pruning-based solution methods providing error bounds have been developed for several dependability measures. Bounds for the reliability have been obtained in [2]. A method to bound the steady-state availability has been proposed in [15]. This method has been further developed in [14], [17]. The error bounds offered by these methods are qualitatively superior to the accuracy assessment offered by simulation methods recently developed to attack the largeness problem [10], [3], whose reliability depends on how well the variance is estimated. This makes of great interest the development of efficient bounding techniques.

This paper presents a new method to obtain steady-state unavailability bounds using CTMC models which, typically, gives significantly smaller bands than the method proposed in [15]. Our method uses an upper bound exploiting the fact that, very often, the system is operational a large portion of the time the model is out of the generated state space. The upper bound is developed using the failure distance concept. The rest of the paper is organised as follows. Section 2 describes the availability models under consideration, reviews the method proposed in [15] and, using a regenerative perspective, argues the potential looseness of the steady-state unavailability upper bound given by the method. Section 3 presents the theoretical developments yielding the upper bound used in our method. Section 4 illustrates with examples the reduction in the steady-state unavailability band which our method can achieve and discusses the computational overheads of our method in relation to the method proposed in [15]. Section 5 concludes the paper.

2 Preliminaries

The type of models addressed in this paper are those which result from conceptualising a fault-tolerant system as made up of components which fail and are repaired with constant rates. The system is operational or down as determined by a coherent structure function [1] on the unfailed/failed state of the components. This basically means that repairs cannot take down an operational system and that failures cannot bring operational a down system. A failure of a component can be propagated to other components. In addition, each component can be failed in a finite number of modes. Failure and repair rates and failure propagation can depend on the state of the system. Let $X = \{X(t); t \geq 0\}$ be the CTMC modeling the system, Ω its state space and o the (only) state in which all components of the system are unfailed. We assume that repair transitions involve only one component and that at least a repair transition exists from any state $\neq o$. It follows from the

hypotheses that X is finite and irreducible. It is assumed that a high-level description of the model is available from which it is possible to identify the bags of components (we allow component types with instances) which can be failed simultaneously in a single event. Those bags are called *failure events*. E will denote the set of failure events of the model and E_i the set of failure events including i components. It is also assumed that links to the high-level description of the model exist allowing to determine during generation of the CTMC the failure event associated to a failure transition and the component affected by a repair transition. Using this information, it is possible to compute the bag of failed components $F(x)$ in each generated state x.

Let D be the subset of down states and let p_i, $i \in \Omega$ the steady-state probability distribution of X, the steady-state unavailability is defined as $UA = \sum_{i \in D} p_i$. *UA* is a special case of the more general steady-state reward rate measure $R = \sum_{i \in \Omega} r(i) p_i$, where $r(i), i \in \Omega$ is an arbitrary reward rate structure imposed on X.

Since repair rates are usually several orders of magnitude higher than failure rates, X is highly skewed, i.e., it has a probability distribution concentrated in a small portion of the state space (the states with few failed components). Thus, in general, good approximations for R can be computed using only a small portion of Ω. However, assessing the accuracy of the solution is a difficult problem. The method proposed in [15] was the first to obtain tight bounds in the context of availability modeling. The method can be used to bound any steady-state reward rate measure R (see [17]). Let G be the generated portion of Ω, U the non-generated portion, and S the subset of G through which X can enter G (from U). As in [15], assume that G contains all states with up to a given number K of failed components. The method can be described in terms of the CTMC's X'_i, $i \in S$ which (conceptually) can be obtained from X as shown in Figure 1. First, X_i is obtained from X by redirecting to i all transitions from U to G (S). Second, U is replaced by the states $u_{K+1}, \ldots, u_N$, where each u_k accounts for the subset U_k of U including all the states with k failed components. Failure transitions from G to states in U with k failed components are directed to state u_k. Each state u_k, $k < N$ has transitions to states u_{k+j} with rates $f_j(k)$ chosen to be upper bounds for the sum of the failure transitions rates from any state with k failed components to U_{k+j}. Each state u_k, $k > K+1$ has also a repair transition to u_{k-1} with a rate $g(k)$ chosen to be a lower bound for the sum of the repair transition rates from any state with k failed components. A similar transition is also introduced from state u_{K+1} to state i. For $f_i(k)$ we can take $\sum_{e \in E_i} \lambda_{ub}(e)$, where $\lambda_{ub}(e)$ is an upper bound to the rate of the failure event e. For $g(k)$ we can take the slowest repair rate of the model. Let $|r|_{lb}$ and $|r|_{ub}$ be, respectively, lower and upper bounds for the reward rate in any state of X, the bounds for R are obtained using the following recipe:

1. for each state $i \in S$ find the steady-state distribution of X'_i and, assigning to the states in G the same reward rate as in X and to the states $u_{K+1}, \ldots, u_N$ a reward rate $|r|_{lb}$ ($|r|_{ub}$), compute the resulting steady-state reward rate $|R_i|_{lb}$ ($|R_i|_{ub}$),
2. $|R|_{lb} = \min_{i \in S} |R_i|_{lb}$, $|R|_{ub} = \max_{i \in S} |R_i|_{ub}$.

Typically S will include all states in G with K failed components and the number of models

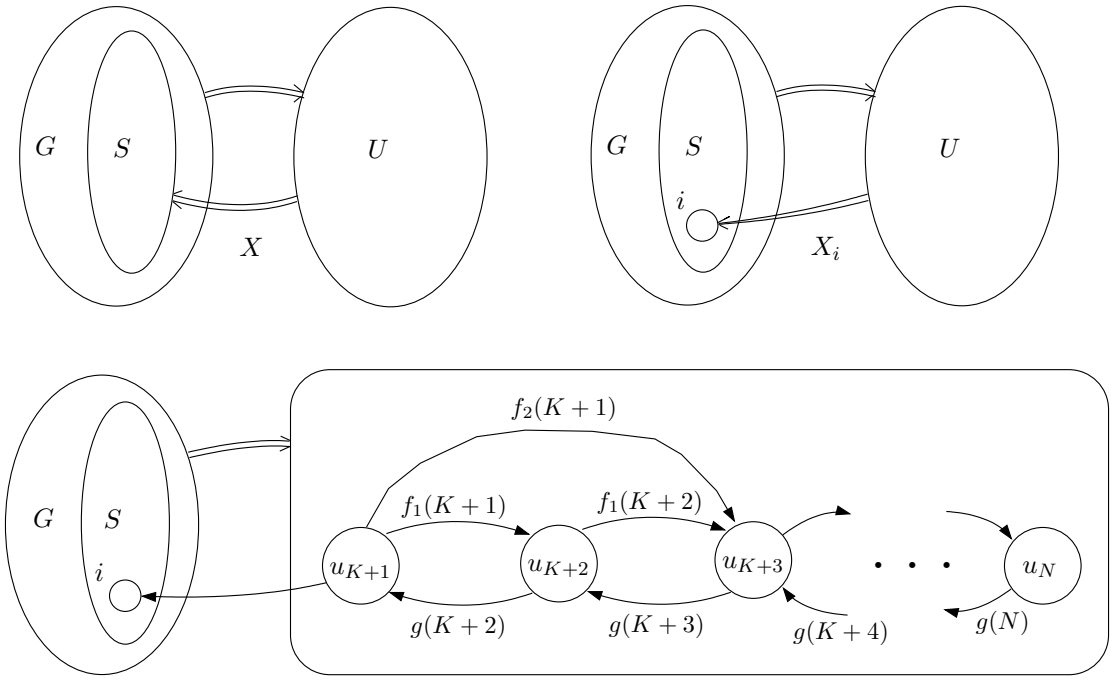

Figure 1: Construction of the CTMC's X_i' used by the bounding method proposed in [Mun89].

X_i' to be solved can be large. The computational cost of the method can be reduced at the expense of some looseness of the bounds by the state duplication technique proposed in [15]. In this technique, duplicates of all states in G with more than F failed components are (conceptually) added to U to account for the visits to these states after the number of failed components is made larger than K and before the number of failed components is made equal to F. This state duplication technique can be thought as a redefinition of the CTMC X to which the bounding method is applied. The resulting CTMC's X_i' have the structure depicted in Figure 1, except that the aggregate states u_i will run now from u_{F+1} to u_N and S will only include the states with F failed components which can be reached through repair transitions. Taking F small enough reduces arbitrarely the number of CTMC's X_i' to be solved. A final remark is that G does not have necessarely to include all states with up to a given number of failed components (see [17]). However, it does have to include all states with F or fewer failed components.

The reviewed bounding method was justified in [15] using the exact aggregation theorem for ergodic Markov chains and bounds on conditional steady-state distributions in subsets of Markov chains [7], [8]. Here, it will be discussed using a regenerative perspective. The motivation is to support theoretically our bounding method and ease the comparison between the method proposed in [15] and ours. Let C_i and T_i be, respectively, the expected reward and the expected time in X_i between consecutive jumps from U to i (regeneration points). Let R_i be the steady-state reward rate of X_i. Then, by regenerative theory, $R_i = C_i/T_i$. In addition, using semi-regenerative process theory [5] it is possible to obtain the following result:

Theorem 1. *Let $X = \{X(t); t \geq 0\}$ be a finite irreducible CTMC with state space Ω and reward*

rate structure $r(i)$, $i \in \Omega$. Let $\Omega = G \cup U$ be a non-trivial partition of Ω ($G, U \neq 0$) and let S be the subset of G through which X can enter G from U. Let R be the steady-state reward rate of X. Let X_i, $i \in S$, be the CTMC obtained from X by redirecting to i the transitions from U to S, assume $X_i(0) = i$, and let R_i be the steady-state reward rate of X_i. Then, $\min_{i \in S} R_i \leq R \leq \max_{i \in S} R_i$.

Theorem 1 has immediate application to the CTMC's X under consideration. The condition $X_i(0) = i$ is in general required because X_i could contain several closed sets. However, for the CTMC's X considered here, X_i is irreducible, and the steady-state reward rate R_i is independent on the initial distribution of X_i. An sketch of the proof of the theorem is given in the Appendix. The complete proof can be found in [4]. Using Theorem 1, the correctness of the bounds for R computed in the recipe follows from the correctness of the bounds $|R_i|_{lb}$ and $|R_i|_{ub}$ for R_i computed in the first step.

Let $C_{G,i}$ and $C_{U,i}$ denote, respectively, the contributions of the states in G and U to C_i and assume a similar notation for the contributions of the states in G and U to T_i. Then, we have $C_i = C_{G,i} + C_{U,i}$, $T_i = T_{G,i} + T_{U,i}$, and

$$R_i = \frac{C_{G,i} + C_{U,i}}{T_{G,i} + T_{U,i}} .$$

Consider now the regenerative behavior of X_i' defined by the times at which X_i' hits i from u_{F+1} (analogous to the regenerative behavior considered for X_i). As it will be shown later, the mean time in the states $u_{F+1}, \ldots, u_N$ between regenerations upper bounds $T_{U,i}$, so we can properly call it $|T_{U,i}|_{ub}$. Notice that, since X_i and X_i' enter G through the same state and are identical in G, the mean reward and time in G between regenerations are identical for X_i and X_i'. Then, the lower and upper bounds for R_i computed in the first step of the recipe can be written as

$$|R_i|_{lb} = \frac{C_{G,i} + |r|_{lb}|T_{U,i}|_{ub}}{T_{G,i} + |T_{U,i}|_{ub}} , \tag{1}$$

$$|R_i|_{ub} = \frac{C_{G,i} + |r|_{ub}|T_{U,i}|_{ub}}{T_{G,i} + |T_{U,i}|_{ub}} . \tag{2}$$

The correctness of these bounds can be justified as follows. Let $g_{lb}(x) = (C_{G,i} + |r|_{lb}x)/(T_{G,i} + x)$, $g_{ub}(x) = (C_{G,i} + |r|_{ub}x)/(T_{G,i} + x)$. Their first derivatives are $dg_{lb}/dx = (|r|_{lb}T_{G,i} - C_{G,i})/(T_{G,i} + x)^2$, $dg_{ub}/dx = (|r|_{ub}T_{G,i} - C_{G,i})/(T_{G,i} + x)^2$. Using $|r|_{lb}T_{G,i} \leq C_{G,i} \leq |r|_{ub}T_{G,i}$, we have $dg_{lb}/dx \leq 0$, $dg_{ub}/dx \geq 0$. Then, since $|r|_{lb}T_{U,i} \leq C_{U,i} \leq |r|_{ub}T_{U,i}$,

$$R_i = \frac{C_{G,i} + C_{U,i}}{T_{G,i} + T_{U,i}} \geq \frac{C_{G,i} + |r|_{lb}T_{U,i}}{T_{G,i} + T_{U,i}} = g_{lb}(T_{U,i}) \geq g_{lb}(|T_{U,i}|_{ub}) = \frac{C_{G,i} + |r|_{lb}|T_{U,i}|_{ub}}{T_{G,i} + |T_{U,i}|_{ub}} = |R_i|_{lb} ,$$

$$R_i = \frac{C_{G,i} + C_{U,i}}{T_{G,i} + T_{U,i}} \leq \frac{C_{G,i} + |r|_{ub}T_{U,i}}{T_{G,i} + T_{U,i}} = g_{ub}(T_{U,i}) \leq g_{ub}(|T_{U,i}|_{ub}) = \frac{C_{G,i} + |r|_{ub}|T_{U,i}|_{ub}}{T_{G,i} + |T_{U,i}|_{ub}} = |R_i|_{ub} .$$

For the particular case of the steady-state unavailability $|r|_{lb} = 0$, $|r|_{ub} = 1$ and the bounds (1), (2) can be written as

$$|UA_i|_{lb} = \frac{C_{G,i}}{T_{G,i} + |T_{U,i}|_{ub}} , \tag{3}$$

$$|UA_i|_{ub} = \frac{C_{G,i} + |T_{U,i}|_{ub}}{T_{G,i} + |T_{U,i}|_{ub}}. \tag{4}$$

The examples given in [15] indicate that $|UA|_{ub}$ tends to be much looser than $|UA|_{lb}$. An intuitive explanation for this is the following. Since X_i tends to be highly skewed, typically $|T_{U,i}|_{ub} \ll T_{G,i}$. Since $UA_i = (C_{G,i} + C_{U,i})/(T_{G,i} + T_{U,i})$, the tightness of $|UA_i|_{lb}$ (3) and $|UA_i|_{ub}$ (4) depend mainly on the closeness of $C_{G,i}$ and $C_{G,i} + |T_{U,i}|_{ub}$ to $C_{G,i} + C_{U,i}$. Down states tend to be sparse and, typically, $C_{U,i} \ll T_{U,i}$. Then, $C_{G,i} + |T_{U,i}|_{ub}$ tends to be less closer to $C_{G,i} + C_{U,i}$ than $C_{G,i}$, making $|UA_i|_{ub}$ significantly looser than $|UA_i|_{lb}$.

3 Proposed bounding approach

3.1 Setup

Our method differs from the method given in [15] in the use of tighter upper bounds for UA_i, $i \in S$. We start considering the more general steady-state reward rate measure R and showing how a different upper bound for R_i, $|R_i|'_{ub}$, can be established using an upper bound for $C_{U,i}$. First, $C_{U,i} \leq |r|_{ub} T_{U,i}$ implies

$$R_i = \frac{C_{G,i} + C_{U,i}}{T_{G,i} + T_{U,i}} \leq \frac{C_{G,i} + C_{U,i}}{T_{G,i} + C_{U,i}/|r|_{ub}} = h(C_{U,i}),$$

with $h(x) = (C_{G,i} + x)/(T_{G,i} + x/|r|_{ub})$. In addition, $dh/dx = (T_{G,i} - C_{G,i}/|r|_{ub})/\ (T_{G,i} + x/|r|_{ub})^2 \geq 0$, since $C_{G,i} \leq |r|_{ub} T_{G,i}$. Thus, $h(x)$ is monotonically increasing and

$$R_i \leq h(|C_{U,i}|_{ub}) = \frac{C_{G,i} + |C_{U,i}|_{ub}}{T_{G,i} + |C_{U,i}|_{ub}/|r|_{ub}} = |R_i|'_{ub}. \tag{5}$$

Regarding the tightness of $|R_i|_{ub}$ and $|R_i|'_{ub}$, we have the following result:

Theorem 2. *Assume $C_{G,i} < |r|_{ub} T_{G,i}$. Then, $|R_i|'_{ub} < |R_i|_{ub}$ if and only if $|C_{U,i}|_{ub} < |r|_{ub} |T_{U,i}|_{ub}$.*

Proof: Consider again the function $h(x) = (C_{G,i} + x)/(T_{G,i} + x/|r|_{ub})$. For $C_{G,i} < |r|_{ub} T_{G,i}$, $dh/dx = (T_{G,i} - C_{G,i}/|r|_{ub})/(T_{G,i} + x/|r_{ub}|)^2 > 0$. This implies that $h(x)$ is strictly monotonically increasing and, since (5) $|R_i|'_{ub} = h(|C_{U,i}|_{ub})$ and (2) $|R_i|_{ub} = h(|r|_{ub} |T_{U,i}|_{ub})$, the result follows. □

For the steady-state unavailability ($|r|_{ub} = 1$) $|R_i|'_{ub}$ (5) is reduced to

$$|UA_i|'_{ub} = \frac{C_{G,i} + |C_{U,i}|_{ub}}{T_{G,i} + |C_{U,i}|_{ub}}, \tag{6}$$

where C has the meaning of "mean down time". Also, the fact that the state o is operational and, therefore, has reward rate 0 ensures $C_{G,i} < T_{G,i}$. Then, Theorem 2 establishes that $|UA_i|'_{ub} < |UA_i|_{ub}$ if and only if $|C_{U,i}|_{ub} < |T_{U,i}|_{ub}$.

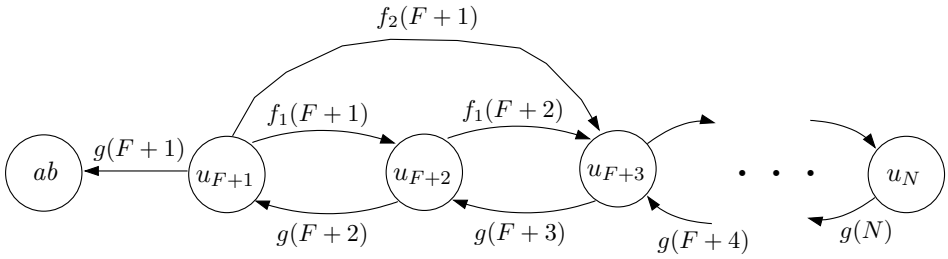

Figure 2: Transient CTMC Y used to derive the bounds $T(k)$ and $C(k)$.

The bounds $|UA_i|_{lb}$ and $|UA_i|_{ub}$ are computed in [15] from the steady-state solution of X'_i. In our bounding method, $C_{G,i}$, $T_{G,i}$, $|T_{U,i}|_{ub}$ and $|C_{U,i}|_{ub}$ are computed independently and then combined using (3), (6) to obtain $|UA_i|_{lb}$ and $|UA_i|'_{ub}$.

3.2 Computation of $T_{G,i}$, $C_{G,i}$ and $|T_{U,i}|_{ub}$

In the following $\tau(v, Z)$ will denote the mean time to absorption in the state or subset of states v of the transient CTMC Z with given initial distribution. Let $\mathbf{A}$ be the restriction of the transition rate matrix of Z to its transient states, $\mathbf{q}$ the column vector giving the initial probability distribution of Z, and $\boldsymbol{\tau}$ the solution of $\mathbf{A}\boldsymbol{\tau} = -\mathbf{q}$. As it is well-known, $\tau(i, Z) = \tau_i$.

$T_{G,i}$ and $C_{G,i}$ can be computed solving the transient CTMC Y_G^i with initial state i tracking X from i to exit of G:

$$T_{G,i} = \sum_{j \in G} \tau(j, Y_G^i)\,,$$

$$C_{G,i} = \sum_{j \in G \cap D} \tau(j, Y_G^i)\,.$$

The upper bound $|T_{U,i}|_{ub}$ can be computed using the transient CTMC Y depicted in Figure 2. The boundness of $|T_{U,i}|_{ub}$ will be justified using exact aggregation results for transient CTMC's [9] and the following lemma (see [4] for the proof), closely related to the mean holding time lemma of [15].

Lemma 1. *Let a transient CTMC Y with the structure depicted in Figure 2 and consider another transient CTMC, Y', with the same structure and such that $f'_i(k) \leq f_i(k)$ and $g'(k) \geq g(k)$. Also assume that Y and Y' have the same initial distribution. Then, $\tau(u_i, Y') \leq \tau(u_i, Y)$, $i = F+1, \ldots, N$.*

Let T_U^s be the mean time spent by X during a visit to U conditioned to entry through state s. T_U^s is the mean time to absorption of the transient CTMC Y_U^s with initial state s tracking X from s to exit of U. Let $Y_U^{s\prime}$ be the result of the exact aggregation in Y_U^s of the subsets U_k, $k = F+1, \ldots, N$. $Y_U^{s\prime}$ has the structure of Y and initial state $u_{|F(s)|}$. From exact aggregation results for transient CTMC's [9], $\tau(u_k, Y_U^{s\prime}) = \tau(U_k, Y_U^s)$ and the transition rates of $Y_U^{s\prime}$ are convex

linear combinations of the transition rates of Y_U^s. More specifically, $\lambda'_{u_k,u_l} = \sum_{j\in U_k} w_j^{k,l}\lambda_{j,U_l}$ and $\lambda'_{u_{F+1},ab} = \sum_{j\in U_{F+1}} w_j^{F+1,ab}\lambda_{j,ab}$, with $w_j \geq 0$, $\sum_j w_j^{k,l} = 1$, $\sum_j w_j^{F+1,ab} = 1$. Consider the "failure" transition rates of $Y_U^{s'}$, $\lambda'_{u_k,u_{k+i}}$, $k \geq F+1$, $i > 0$. Since $f_i(k)$ upper bounds $\lambda_{j,U_{k+i}}$, $j \in U_k$,

$$\lambda'_{u_k,u_{k+i}} = \sum_{j\in U_k} w_j^{k,k+i}\lambda_{j,U_{k+i}} \leq \max_{j\in U_k} \lambda_{j,U_{k+i}} \leq f_i(k)\,.$$

Using $g(k) \leq \lambda_{j,U_{k-1}}$, $j \in U_k$, $k > F+1$ and $g(F+1) \leq \lambda_{j,ab}$, $j \in U_{F+1}$, it can be similarly shown that $\lambda'_{u_k,u_{k-1}} \geq g(k)$, $k > F+1$ and $\lambda'_{u_{F+1},ab} \geq g(F+1)$. In summary, the transition rates of Y and $Y_U^{s\prime}$ satisfy the conditions of Lemma 1. Denoting by Y^k the transient CTMC Y with initial state u_k and by $T(k)$ the mean time to absorption of Y^k and using Lemma 1,

$$T_U^s = \sum_{j=F+1}^{N} \tau(U_j, Y_U^s) = \sum_{j=F+1}^{N} \tau(u_j, Y_U^{s\prime}) \leq \sum_{j=F+1}^{N} \tau(u_j, Y^{|F(s)|}) = T(|F(s)|)\,.$$

Let ϕ_s^i be the conditional entry probability distribution of X_i in U through state s. ϕ_s^i can be computed from the mean times to absorption of Y_G^i as

$$\phi_s^i = \sum_{j\in G} \tau(j, Y_G^i)\lambda_{j,s}\,. \tag{7}$$

Let π_k^i be the probability that X_i enters U through U_k. We have

$$\pi_k^i = \sum_{s\in U_k} \phi_s^i\,. \tag{8}$$

Then, $|T_{U,i}|_{ub}$ can be computed as

$$|T_{U,i}|_{ub} = \sum_{k=F+1}^{N} \pi_k^i T(k)\,. \tag{9}$$

The upper boundness of $|T_{U,i}|_{ub}$ can be easily justified using $T_U^s \leq T(|F(s)|)$:

$$T_{U,i} = \sum_{s\in U} \phi_s^i T_U^s = \sum_{k=F+1}^{N} \sum_{s\in U_k} \phi_s^i T_U^s \leq \sum_{k=F+1}^{N} \sum_{s\in U_k} \phi_s^i T(k) = \sum_{k=F+1}^{N} \pi_k^i T(k)\,.$$

Giving the relationships between Y and X_i', it is clear that $|T_{U,i}|_{ub}$ is the upper bound for $T_{U,i}$ implicitely used in [15].

Although the bounds $|T_{U,i}|_{ub}$ can be computed directly as the mean times to absorption of Y with initial distributions $P[Y(0) = u_k] = \pi_k^i$, this procedure requires $|S|$ solutions of Y (one for each state $i \in S$) and a more efficient approach when $|S| > 1$ is to compute $|T_{U,i}|_{ub}$ from $T(k)$, $k = F+1, \ldots, N$ using (9). $T(N)$ can be computed solving Y^N as $T(N) = \sum_{j=F+1}^{N} \tau(u_k, Y^N)$. Denoting by $\lambda(k)$ the output rate of u_k in Y, the remaining $T(k)$'s can be computed exploiting the following relations, which result from a conditional path analysis of Y.

$$T(k) = \frac{1}{\lambda(k)} + \frac{g(k)}{\lambda(k)}T(k-1) + \sum_i \frac{f_i(k)}{\lambda(k)}T(k+i)\,, \quad F+1 < k < N\,,$$

$$T(N) = \frac{1}{g(N)} + T(N-1)\,,$$

yielding

$$T(N-1) = T(N) - \frac{1}{g(N)}\,,$$

$$T(k) = \frac{1}{g(k+1)}[\lambda(k+1)T(k+1) - 1 - \sum_i f_i(k+1)T(k+1+i)]\,, \qquad k = N-2,\ldots,F+1\,.$$

3.3 Computation of $|C_{U,i}|_{ub}$

The strategy to find a bound $|C_{U,i}|_{ub}$ potentially smaller than $|T_{U,i}|_{ub}$ is to exploit the fact that many of the states in U are operational and, thus, do not contribute to $C_{U,i}$. As we shall show, the strategy can be implemented using the concept of *failure distance*, which has been useful to speed up the simulation of the type of models considered in this paper [3]. The failure distance from an state x, $d(x)$, is defined as the minimum number of components which have to fail (in addition to $F(x)$) for the system to go down ($d(x) = 0$ for $x \in D$).

Let $U_{k,d}$ be the subset of U including the states with k failed components and failure distance d and let $\pi^i_{k,d}$ be the probability that X_i enters U through $U_{k,d}$. We have

$$\pi^i_{k,d} = \sum_{s \in U_{k,d}} \phi^i_s\,. \tag{10}$$

Assume that upper bounds $C(k,d)$ to the mean down time in U conditioned to entry in U through any state $\in U_{k,d}$ are available. Then, an upper bound for $C_{U,i}$ can be computed as

$$|C_{U,i}|_{ub} = \sum_{k,d} \pi^i_{k,d} C(k,d)\,. \tag{11}$$

Since $\pi^i_k = \sum_d \pi^i_{k,d}$, it is clear (9) that $C(k,d) \le T(k)$ implies $|C_{U,i}|_{ub} \le |T_{U,i}|_{ub}$. If, in addition, $C(k,d) < T(k)$ for some pair (k,d) with $\pi^i_{k,d} \neq 0$, $|C_{U,i}|_{ub} < |T_{U,i}|_{ub}$.

Our approach to obtain bounds $C(k,d) \le T(k)$ includes two steps. In the first step, we obtain upper bounds to the mean down time in U conditioned to entry in U through U_k. Then, we let $C(k,d) = C(k)$ and improve iteratively $C(k,d)$. The bounds $C(k)$ are $\le T(k)$ and, as a result, $C(k,d) \le T(k)$. Thus, our bounds $|C_{U,i}|_{ub}$ are always $\le |T_{U,i}|_{ub}$ and our upper bound $|UA|'_{ub}$ is never worse than $|UA|_{ub}$.

Let L be the minimum number of components which have to fail to take the system down ($L = d(o)$). With the reward rate structure

$$r(u_j) = \begin{cases} 0 & \text{if} \quad j < L \\ 1 & \text{if} \quad j \ge L \end{cases}\,,$$

the mean reward to absorption of Y^k provides a suitable bound $C(k)$. To justify this, let C^s_U be the mean down time in a stay in U since entry through state s. C^s_U is the mean down time of the

transient CTMC Y_U^s. Using exact aggregation results for transient CTMC's, Lemma 1, and the fact that all states in U_k, $k < L$ are operational,

$$C_U^s = \sum_{j \in U \cap D} \tau(j, Y_U^s) \leq \sum_{k \geq L} \tau(U_k, Y_U^s) = \sum_{k \geq L} \tau(u_k, Y_U^{s\prime}) \leq \sum_{k \geq L} \tau(u_k, Y^k) = C(k)\,.$$

For $F + 1 \geq L$, $C(k) = T(k)$. Otherwise, $C(k) < T(k)$. $C(N)$ can be easily computed from the mean times to absorption vector of Y^N as $C(N) = \sum_{i=L}^{N} \tau(u_k, Y^N)$. The remaining $C(k)$'s can be computed using the following recursive equations (analogous to the equations giving $T(k)$, $k < N$), where $I(c)$ is the indicator function which returns 1 if c is true and 0 otherwise.

$$C(N-1) = C(N) - \frac{1}{g(N)}\,,$$

$$C(k) = \frac{1}{g(k+1)}[\lambda(k+1)C(k+1) - I(k+1 \geq L) - \sum_i f_i(k+1)C(k+1+i)]\,,$$
$$k = N-2, \ldots, F+1\,.$$

Let *FC* be the set of different cardinalities of the failure events of the model. Let $F(k, d, i, r)$, $i \in$ *FC*, be upper bounds for the sum of failure rates involving i components from any state in U with k failed components and failure distance d to states with failure distance $\leq r$, let $w = \min\{i, d\}$, and let

$$f_{i,j}(k, d) = F(k, d, i, d-j) - F(k, d, i, d-j-1)\,, \quad 0 \leq j < w\,,$$
$$f_{i,w}(k, d) = F(k, d, i, d-w)\,.$$

The iterative improvement procedure of $C(k, d)$ is based on the following result (proved in [4]), where in the expression for $C'(k, d)$ the terms $C(k, d)$ corresponding to unfeasible pairs (k, d) have to be set to 0. The feasible pairs (k, d) are given by $F + 1 \leq k \leq N$, $\max\{0, L - k\} \leq d \leq \min\{L, N - k\}$.

Proposition 1. *Let $C(k, d)$ be upper bounds for C_U^s, $s \in U_{k,d}$ and assume that $C(k, d)$ is decreasing on d. Then, for any $s \in U_{k,d}$,*

$$C_U^s \leq C'(k, d) = \frac{I(d=0)}{g(k)} + \max\{C(k-1, d), C(k-1, d+1)\} + \sum_{i \in FC} \sum_{j=0}^{w} \frac{f_{i,j}(k, d)}{g(k)} C(k+i, d-j)\,. \quad (12)$$

The iterative improvement procedure can be implemented using (12). At each step, $C'(k, d)$ is computed for each feasible (k, d) pair and accepted as new $C(k, d)$ if $C'(k, d) < C(k, d)$. The procedure can be finished when no bound $C(k, d)$ has been reduced significantly during a step. It is important to note that the correctness of the bounds $C'(k, d)$ requires that the available set of $C(k, d)$ bounds be decreasing on d. It is proved in [4] that this is satisfied if 1) the bounds $F(k, d, i, r)$ are decreasing on d, and 2) the bounds $C(k, d)$ are reviewed grouped by k. In our implementation the

bounds are reviewed by increasing values of k and, for a given k, by increasing values of d. This ordering has been proved effective, in the sense that very few improvement steps (typically < 10) are required to reach stable values for the bounds.

It is possible to argue that the bounds $C(k,d)$ obtained at the end of the iterative improvement procedure for $d > 0$ are potentially much smaller than the original $C(k)$ if $\sum_{j=0}^{w} f_{i,j}(k,d) = F(k,d,i,d) \ll g(k)$. Consider $C'(F+1,d)$ with $d > 0$ and $C(k,d) = C(k)$. For such a case, the first two terms of $C'(F+1,d)$ are 0 ($C(k,d) = 0$ for non-feasible (k,d) pairs) and only the last term remains, but even considering that $C(k) > C(F+1)$ for $k > F+1$, the last term can be much smaller than $C(F+1)$ if $\sum_{j=0}^{w} f_{i,j}(k,d) \ll g(k)$. Consider now $C'(F+2,d)$ with $d > 0$. A similar discussion can be made except that the second term will not be null, but since this term corresponds to revised values $C'(F+1,d)$ with $d > 0$, it is potentially much smaller than $C(F+1)$, and thus than $C(F+2)$. The argument can be iterated for increasing values of k.

Combining (7), (8), (9) and (10) $|T_{U,i}|_{ub}$ and $|C_{U,i}|_{ub}$ can be formulated as

$$|T_{U,i}|_{ub} = \sum_{j \in G} \tau(j, Y_G^i)\alpha(j)\,,$$

$$|C_{U,i}|_{ub} = \sum_{j \in G} \tau(j, Y_G^i)\beta(j)\,,$$

with

$$\alpha(j) = \sum_{s \in U} \lambda_{j,s} T(|F(s)|)\,,$$

$$\beta(j) = \sum_{s \in U} \lambda_{j,s} C(|F(s)|, d(s))\,.$$

Note that $\alpha(j)$, $\beta(j)$ are independent on i and the above formulations are used with advantage when $|S| > 1$.

3.4 Computation of failure distances and bounds $F(k,d,i,r)$

The computation of $|C_{U,i}|_{ub}$ requires the knowledge of the failure distances from the states in the frontier of U. The failure distance $d(x)$ from a state x can be computed from $F(x)$ if the minimal cuts of the structure function of the system [1] are known. Let *MC* be the set of all minimal cuts of the structure function of the system, using standard bag notation, we have

$$d(x) = \min_{m \in MC} |m - F(x)|\,. \tag{13}$$

Although (13) can be used to compute all the required failure distances, most of the transitions from G to U will be of the failure type (all if G contains all states up to a given number of failed components K) and a more efficient procedure can be established introducing the notion of "after" minimal cuts associated with a given failure event e. Let $MC_e = \{m' \mid m' = m - e, m \in$

$MC, m \bigcap e \neq \phi\}$ be the set of "after" minimal cuts associated to e, the failure distance from any state reached from x through a failure transition with failure event e can be computed as

$$ad(x, e) = \min\{d(x), \min_{m \in MC_e} |m - F(x)|\}. \quad (14)$$

The cardinality of MC_e is in general much smaller than the cardinality of MC. Then, for each state x in the frontier of G we can compute its failure distance using (13), and use (14) to compute the failure distances for the states in U reached from x through failure transitions. If some state y in U is reached from x through a repair transition, then we can construct $F(y)$ and compute $d(y)$ using (13).

The tightness of the bounds $C(k, d)$ depends on the tightness of the bounds $F(k, d, i, r)$. In general, better bounds $F(k, d, i, r)$ require a more detailed analysis of the model and thus their computation requires more effort. The bounds $F(k, d, i, r)$ used here are relatively easy to compute and, as the examples in the next section will show, provide good results. The bounds are based on two structural properties of failure events. The *importance* $I(e)$ of a failure event e is defined as the minimum number of components which are left unfailed in any minimal cut affected by the failure event. The *activity* $A(e)$ of a failure event e is defined as the maximum number of components of the failure event in any minimal cut. From their definitions, $I(e)$ and $A(e)$ can be computed by

$$I(e) = \min_{m \in MC, m \cap e \neq \phi} |m - e|,$$

$$A(e) = \max_{m \in MC} |m \cap e|.$$

Consider a state with k failed components and failure distance d and another state reached from it through a failure event e. The number of components left unfailed in any minimal cut m is $\geq |m - e| - k$, since at most k components not included in $m \cap e$ were failed before e. Then, $d' \geq I(e) - k$. Also, $d' \geq d - A(e)$, since at most $A(e)$ components in the same minimal cut will be failed by e. Imposing $d' \leq r$ results in:

$$I(e) - k \leq r,$$

$$d - A(e) \leq r.$$

Then, the failure rate from any state with k failed components and failure distance d due to failure events with i components leading to states with failure distance $\leq r$ is bounded above by

$$F(k, d, i, r) = \sum_{e \in E_i, A(e) \geq d-r, I(e) \leq k+r} \lambda_{ub}(e).$$

It is easy to check that these bounds are decreasing on d, as required for the correctness of the iterative improvement procedure for $C(k, d)$.

4 Numerical Analysis

In this section our bounding method is compared with the method proposed in [15] using the large model described there and a variation of it to explore the impact of the redundancy level L on the

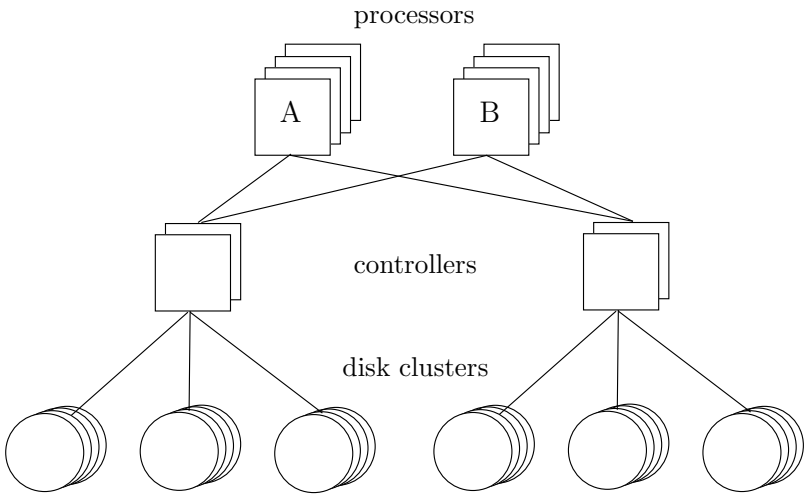

Figure 3: Fault-tolerant database system from [Mun89] (model 1).

relative tightness of the bounds given by the methods. We use the same state generation strategy as in [15], i.e., G includes all the states with up to K failed components. The large model considered in [15] is the fault-tolerant database system shown in Figure 3. The system includes two processor types (A and B), two sets of dual-ported controllers with two controllers per set and six disk clusters with four disks. Each set of controllers controls three clusters. Each processor type has three spares. The system is operational if at least one processor of any type is unfailed, at least one controller in each set is unfailed, and at least three disks in each cluster are unfailed. Thus, $L = 2$. A failure in the active processor A is propagated to the active processor B with probability 0.10. Processors and controllers of one set fail with rate $1/2000$, controllers of the other set fail with rate $1/4000$. Disks fail with different rates from one cluster to another. These rates are $1/6000$, $1/8000$, $1/10000$, $1/12000$, $1/14000$, and $1/16000$. Any component fails in one of two modes with equal probabilities. The repair rate is 1 for one mode and 0.5 for the other. Components are repaired by a single repairman who chooses components at random from the set of failed components. Unfailed components continue to fail when the system is down. This model has about 9×10^{10} states, clearly illustrating the "largeness" problem. A slight variation of this example is also considered. We call the original model from [15] model 1, and call model 2 its variation. Model 2 is obtained from model 1 by increasing the number of controllers in each set to 3 and the number of disks in each cluster to 5, without modifying any other aspect. For model 2, $L = 3$.

Tables 1 and 2 give the number of generated states, the steady-state unavailability bounds and bands under both methods, and the improvement measured as the band ratio. Our method always gives significantly smaller bands. Thus, for model 1, our method for $K = 2$ (231 states) gives bounds wich can be considered tight enough for most purposes, whereas the bounds given by the method proposed in [15] are quite loose. Using that method, 1,763 states ($K = 3$) should be generated to achieve bounds of acceptable quality. For model 2, our method with any K gives tigther bounds than the other method with $K + 1$. The improvement of our method decreases for larger values of K and is considerably larger for model 2. Both behaviours can be explained by the

relative sparsness of down states in U.

Table 1: Comparison of the bounding methods for model 1.

$\|G\|$	K	F	Muntz et al.	proposed	improvement
231	2	0	3.2313×10^{-6}	3.2313×10^{-6}	23.7
			8.9886×10^{-6}	3.4746×10^{-6}	
			5.7573×10^{-6}	2.4322×10^{-7}	
1,763	3	0	3.3167×10^{-6}	3.3167×10^{-6}	14.2
			3.4182×10^{-6}	3.3239×10^{-6}	
			1.0155×10^{-7}	7.1676×10^{-9}	
10,464	4	0	3.3192×10^{-6}	3.3192×10^{-6}	9.44
			3.3208×10^{-6}	3.3194×10^{-6}	
			1.5547×10^{-9}	1.6469×10^{-10}	

Table 2: Comparison of the bounding methods for model 2.

$\|G\|$	K	F	Muntz et al.	proposed	improvement
231	2	0	0	0	522
			8.5262×10^{-6}	1.6324×10^{-8}	
			8.5262×10^{-6}	1.6324×10^{-8}	
1,771	3	0	4.5418×10^{-9}	4.5418×10^{-9}	202
			1.6621×10^{-7}	5.3420×10^{-9}	
			1.6167×10^{-7}	8.0016×10^{-10}	
10,616	4	0	4.7214×10^{-9}	4.7214×10^{-9}	95.6
			7.4986×10^{-9}	4.7504×10^{-9}	
			2.7773×10^{-9}	2.9058×10^{-11}	
52,916	5	0	4.7277×10^{-9}	4.7277×10^{-9}	55.2
			4.7736×10^{-9}	4.7286×10^{-9}	
			4.5912×10^{-11}	8.3150×10^{-13}	

It has been observed that the tightness of the bounds derived in [15] increases with F. This is not typically the case with ours. Table 3 gives $|S|$ and the lower and upper bounds obtained with both methods for model 1, $K = 3$ and all possible values for F. The lower bound (identical for both methods) does not experiment variations at the level of the 6th significant digit. The upper bound given by the method proposed in [15] experiments some improvement when F increases. Our upper bound experiments a slight improvement from $F = 0$ to $F = 1$, but deteriorates considerably with further increase of F. This behavior can be explained as follows. Given the orders of magnitude difference between failure and repair transitions, the model reaches state o with high probability and in short time for any $i \in S$ and $T_{G,i}$ tends to depend vary little on the "return" state i. Then, the

dependency of $|UA_i|'_{ub}$ (6) on i comes mainly through $C_{G,i}$ and $|C_{U,i}|_{ub}$. The latter is determined by the exit distribution $\pi^i_{k,d}$ (11). For larger F, S includes states with more failed components and smaller failure distances, the distributions $\pi^i_{k,d}$ are more shifted to high values of k and smaller values of d, and since $C(k,d)$ increases with higher k and smaller d, the corresponding $|C_{U,i}|_{ub}$ are larger. When $F \geq L$, S includes down states and the shift of $\pi^i_{k,d}$ for such states is specially significant, since all failure transitions from i give contributions to $\pi^i_{k,d}$ with $d = 0$. Also, the corresponding transient CTMC's Y^i_G include visits to down states with probability 1, yielding larger $C_{G,i}$ values. For $0 < F < L$ the small dependency of $T_{G,i}$ on i may outweight the other factors and yield a slightly tighter upper bound than for $F = 0$. Both behaviors are clearly supported by the results in Table 3 ($L = 2$ for model 1). The cost in time of the bounding method is very sensitive to F, since $|S|$ CTMC's Y^i_G have to be solved, and $F = 0$ should be the reasonably choice for our method.

Table 3: Impact of F on the bounds for model 1 and $K = 3$.

F	$\|S\|$	Muntz et al.	proposed
0	1	3.31670×10^{-6}	3.31670×10^{-6}
		3.41825×10^{-6}	3.32386×10^{-6}
1	20	3.31670×10^{-6}	3.31670×10^{-6}
		3.39292×10^{-6}	3.32373×10^{-6}
2	210	3.31670×10^{-6}	3.31670×10^{-6}
		3.39275×10^{-6}	3.34645×10^{-6}
3	1532	3.31670×10^{-6}	3.31670×10^{-6}
		3.39272×10^{-6}	3.36944×10^{-6}

Our bounding method is more complex both theoreticaly and computationaly. The last aspect requires some discussion. The only storage and time overheads of our method which can be significant are related to the computation of the failure distances: generation and storage of minimal cuts and computation of failure distances using (13), (14). Efficient algorithms (see [13] for a review) exist which will find all minimal cuts very fast even when their number is of the order of several thousands. Thus generation "per se" does not seem to be an important problem. Since a minimal cut requires less storage than a state and the information associated to it, the requirement of storing the minimal cuts can only be significant when the number of minimal cuts is substantially larger than the number of states of the model. Regarding the cost in time associated to the computation of failure distances, it represented a 5% overhead for the examples used here which have 9 minimal cuts. When the number of minimal cuts is large the method described here for failure distances computation can be time consumming. However, the techniques proposed in [3] can be used to reduce drastically the number of minimal cuts which have to be "touched" to compute the failure and "after" failure distances from a particular state. Using these techniques, storage and time overheads will only be significant when a number of minimal cuts in the order of several thousands has to be managed. We also note that knowing *all* minimal cuts is not a requirement of the method. We can

simply consider the minimal cuts with up to a given number M of components and assume that the system is down for all combinations of more than M failed components to obtain a looser upper bound (but never worse than the bound obtained using [15]). Thus, a tradeoff can be made between tightness of the bounds and overhead caused by the management of the minimal cuts.

5 Conclusions

In this paper we have proposed a method to bound the steady-state unavailability of repairable fault-tolerant systems using CTMC's which, with the same number of generated states, can give significantly smaller (and *never* worse) bands than a method previously proposed [15]. Using the failure distance concept we have obtained an upper bound exploiting the fact that, typically, the system is operational a large portion of the time the model is out of the generated state space. The quality of our upper bound depends on the tightness of the failure rate bounds $F(k,d,i,r)$. The bounds $F(k,d,i,r)$ we have used here are relatively simple and we plan to consider in the future the use of more precise $F(k,d,i,r)$ bounds. We are also interested in studying the behavior of our bounding method and how it compares with the method proposed in [15] in combination with state exploration techniques recently proposed [17].

APPENDIX

Sketch of the proof of Theorem 1

Let C_i (T_i), $i \in S$ be the expected reward (time) in X between entry in i and the next entry in S from U. Using results from semi-regenerative process theory (Theorem 6.12 of [Cin75, Chapter 10]) and using the fact that X is irreducible and finite, it is easy to show that

$$R = \frac{\sum_{i \in S} \psi_i C_i}{\sum_{i \in S} \psi_i T_i} ,$$

where ψ_i, $i \in S$ is any invariant measure of the embedded discrete-time Markov chain Π of X. Being Π finite and irreducible, there exists an invariant measure for Π satisfying $\psi_i > 0$, $\sum_{i \in S} \psi_i = 1$. Using this, it can be shown by induction on $|S|$ that

$$\min_{i \in S}\{C_i/T_i\} \leq R \leq \max_{i \in S}\{C_i/T_i\} .$$

Being X irreducible, S is reached in X from i with probability 1. Then, i is recurrent in X_i. Assuming $X_i(0) = i$, it is easy to check that X_i is recurrent aperiodic. C_i and T_i are, respectively, the expected reward and time between recurrences. Then, by regenerative theory, $R_i = C_i/T_i$ and the result follows. □

References

[1] R. E. Barlow and F. Proshan, *Statistical Theory of Reliability and Life Testing. Probability Models*, McArdle Press, Silver Spring, 1981.

[2] M. A. Boyd, M. Veeraraghavan, J. B. Dugan, and K. S. Trivedi, “An approach to solving large reliability models,” in *Proc. AIAA Components in Aerospace Conference*, pp. 245–258, 1988.

[3] J. A. Carrasco, “Failure distance-based simulation of repairable fault-tolerant systems,” in *Proc. 5th Int. Conf. on Modelling Techniques and Tools for Computer Performance Evaluation*, pp. 351–365, Elsevier, 1992.

[4] J. A. Carrasco, “Tight steady-state availability bounds using the failure distance concept,” Research report, September 1993, submitted for publication.

[5] E. Çinlar, *Introduction to Stochastic Processes*, Prentice-Hall, Inc., New Jersey, 1975.

[6] M. Dal Cin, “Availability analysis of a fault-tolerant computer system,” *IEEE Trans. on Reliability*, vol. R-29, no. 3, pp. 265–268, August 1980.

[7] P. J. Courtois and P. Semal, “Bounds for the positive eigenvectors of nonnegative matrices and for their approximations,” *Journal of the ACM*, vol. 31, no. 4, pp. 804–825, October 1984.

[8] P. J. Courtois and P. Semal, “Computable bounds for conditional steady-state probabilities in large Markov chains and queueing models,” *IEEE J. Selected Areas in Communications*, vol. SAC-4, no. 6, pp. 926–937, September 1986.

[9] P. J. Courtois and P. Semal, “Bounds for transient characteristics of large or infinite Markov chains,” in W. J. Stewart, editor, *Proc. 1st Int. Conf. on Numerical Solution of Markov Chains*, pp. 413–434, Marcel Dekker, 1991.

[10] A. Goyal, P. Shahabuddin, P. Heidelberger, V. F. Nicola, and P. W. Glynn, “A unified framework for simulating Markovian models of highly dependable systems,” *IEEE Trans. on Computers*, vol. 41, no. 1, pp. 36–51, January 1992.

[11] A.M. Johnson Jr. and M. Malek, “Survey of software tools for evaluating reliability, availability and serviceability,” *ACM Computing Surveys*, vol. 20, pp. 227–271, 1988.

[12] J. G. Kemeny and J. L. Snell, *Finite Markov Chains*, Springer-Verlag, New York, 2nd edition, 1978.

[13] W. S. Lee, D. L. Grosh, F. A. Tillman and C. H. Lie, “Fault Tree Analysis, Methods and Applications—A Review,” *IEEE Trans. on Reliability*, vol. R-34, no. 3, pp. 194–203, August 1985.

[14] J. C. S. Lui and R. R. Muntz, “Evaluating bounds on steady-state availability of repairable systems from Markov models,” in W. J. Stewart, editor, *Proc. 1st Int. Conf. on Numerical Solution of Markov Chains*, pp. 435–454, Marcel Dekker, 1991.

[15] R. R. Muntz, E. de Souza e Silva, and A. Goyal, “Bounding availability of repairable computer systems,” *IEEE Trans. on Computers*, vol. 38, no. 12, pp. 1714–1723, December 1989.

[16] R. A. Sahner and K. S. Trivedi, “Reliability modeling using SHARPE,” *IEEE Trans. on Reliability*, vol. R-36, no. 2, pp. 186–193, June 1987.

[17] E. de Souza e Silva and P. M. Ochoa, “State space exploration in Markov models,” *Performance Evaluation Review*, vol. 20, no. 1, pp. 152–166, June 1992.