

Fast Simulation of Steady-State Availability in Non-Markovian Highly Dependable Systems

Victor F. Nicola*
Perwez Shahabuddin
Philip Heidelberger

Peter W. Glynn

IBM T.J.Watson Research Center
Yorktown Heights, NY 10598

Department of Operations Research
Stanford University, CA 94305

Abstract

This paper considers efficient simulation techniques for estimating steady-state quantities in models of highly dependable computing systems with general component failure and repair time distributions. Earlier approaches in this application setting for steady-state estimation rely on the regenerative method of simulation, which can be used when the failure time distributions are exponentially distributed. However, when the failure times are generally distributed the regenerative structure is lost and a new approach must be taken. The approach we take is to exploit a ratio representation for steady-state quantities in terms of cycles that are no longer independent and identically distributed. A "splitting" technique is used in which importance sampling is used to speed up the simulation of rare system failure events during a cycle, and standard simulation is used to estimate the expected cycle length. Experimental results show that the method is effective in practice.

1 Introduction

This paper is concerned with the efficient simulation of certain steady-state quantities in models of highly dependable computing systems. (By "simulation" we mean a discrete event stochastic simulation.) In particular, we will consider techniques for accurately estimating the steady-state unavailability, U , for models in which the failure and repair time distributions are generally distributed. Because the system being modeled is assumed to be highly dependable, system failure events are rare and therefore $U \approx 0$. Standard simulation of such systems require enormous sample sizes in order to accurately estimate U ; typically, the closer U is to zero, the longer a standard simulation needs to be run.

We seek to avoid such long run lengths by using the technique of importance sampling [13], [16]. In importance sampling, the system is simulated using a new set of input distributions (e.g., failure distributions) that are chosen in such a way as to make the rare event much more likely to occur. An unbiased estimate is then obtained by multiplying the output of the simulation experiment by the likelihood ratio. If the new method of sampling is chosen properly, then the variance of the new estimator will be much less than the variance of the standard estimator. Importance sampling has been effectively employed in a variety of situations, including queueing models (see, e.g., [9], [26] and [27]). Another approach to variance reduction when estimating long-run averages in models of communication networks is considered in [22].

Its use in simulating highly dependable systems of the type described in [14] has been studied in [3], [4], [10],

[15], [17], [18], [20], [21], [23], [25], [24], [29] and [30]. A number of these references prove that when the importance sampling distribution is chosen according to a certain heuristic, then the resulting estimator satisfies the "bounded relative error" property. For example, if \hat{U} is an estimator of U , then \hat{U} has bounded relative error if Standard Deviation $[\hat{U}]/U$ remains bounded even as $U \rightarrow 0$. In practice, bounded relative error implies that only a fixed number of samples are required to accurately estimate U , no matter how small U is, i.e., no matter how rare system failure events become. The above papers have considered two distinct situations:

1. Estimation of steady-state quantities such as the long run unavailability U .
2. Estimation of transient quantities such as the reliability $R(t)$ which is defined as the probability that the system does not fail before some fixed time t .

Results on steady-state estimation have basically been restricted to cases in which the component failure time distributions are exponentially distributed. This restriction is required because the steady-state estimators exploit the regenerative structure of such models (see, e.g., [6]). In this case, because of the exponential assumption, regenerations occur whenever the model enters the state in which all components are operational. This permits a ratio representation of steady-state quantities, e.g., $U = E[D_i]/E[C_i]$ where D_i is the total amount of downtime during the i -th regenerative cycle and C_i is the length of the i -th regenerative cycle. (The time between regenerations is called a cycle.) In addition, different regenerative cycles are i.i.d. (independent and identically distributed) thereby permitting straightforward variance estimation and formation of confidence intervals.

While efficient simulation techniques for estimating transient quantities in models with generally distributed failure and repair times have been considered in [17], [18], [25] and [24], until now, these techniques have not been successfully applied to estimating steady-state quantities due to the fact that an entrance to the "all components up" state no longer constitutes a regeneration. In this paper, we show how these techniques can be extended to estimating steady-state measures in systems without regenerative structure. The approach makes use of a ratio representation for steady-state quantities in terms of "A-cycles" that is similar to that obtained in a regenerative setting, but which is more generally applicable. Here, a new A-cycle is defined to start whenever the process enters some set of states A. (In our setting, $A \equiv$ all components operational.) However, the A-cycles are no longer i.i.d., which somewhat complicates the use of importance sampling, variance estimation, and the formation of con-

*Current address: Department of Computer Science, University of Twente, P.O. Box 217, 7500 AE Enschede, The Netherlands.

confidence intervals. Approaches for effectively dealing with these problems will be described in the paper. In particular, we employ a “splitting” technique which uses importance sampling to estimate the expected downtime during an A – cycle and uses standard simulation to estimate the expected length of an A – cycle.

While we have not (yet) proven that this approach gives rise to estimates with bounded relative error, the method is shown to be highly effective in practice. In addition, if the failure and repair distributions are exponential, then the approach is closely related to using “balanced failure biasing” ([29], [15]) and “measure specific dynamic importance sampling” (MSDIS) [15] which, for Markovian models, was shown to have bounded relative error for estimating the steady-state unavailability in [29]. Also, when the failure and repair times are generally distributed, then the importance sampling heuristic used is known to produce an estimate, of the transient reliability $R(t)$, having bounded relative error [18]. For these reasons, we conjecture that (under appropriate technical conditions) the method does produce estimates of the steady-state unavailability with bounded relative error when failure and repair times are generally distributed.

The rest of the paper is organized as follows. In Section 2 the ratio representation is discussed. The issues of how to effectively combine importance sampling with this ratio representation for general rare event estimation, and how to estimate the variance are also described in Section 2. The application of these results to estimating the steady-state unavailability and our particular importance sampling heuristic are described in Section 3. Experimental results are presented in Section 4, and the results are summarized in Section 5.

2 Estimation of Steady-State Measures in Non-Regenerative Systems

Consider systems that can be modelled as (time homogeneous) Generalized Semi-Markov Processes (GSMP's). For a detailed exposition on GSMP's the reader is referred to [11] (see also [25] for an alternative description). Roughly speaking, these are systems which are characterized by an output state vector, say $X(t)$, that takes value in \mathbf{Z}^l , and an internal state vector, say $S(t)$, that takes values in \mathbf{R}^m ($\mathbf{Z} \equiv$ set of integers, l is a positive integer and m is a non-negative integer). The choice of the output state vector depends on the application at hand and the desired level of detail. The $S(t)$ is defined such that it has enough information about the history of the system, so that $\{(X(s), S(s)) : s > t\}$ depends on $\{(X(s), S(s)) : 0 \leq s \leq t\}$ only through $(X(t), S(t))$. Let $f(X(t))$ be some bounded real valued function on the output state space \mathcal{S} . For example, consider a system with N different components. Each component has generally distributed failure and repair times. The age of an operational component is the time since it is last became operational. The system has many repairmen (each component is assigned a particular repairman) which repair components using the FCFS (first come first served) service discipline, and components interact by sharing repairmen. In this case we may define $X(t)$ to be an N -dimensional vector, the i – th element of which is 1 if the i – th component is up and 0 otherwise. The internal state vector $S(t)$ may be defined to include the ages of the components that are operational, the repair queue and the elapsed repair time (if any) at each repairman. For the purposes of availability estimation, the function $f(X(t))$ may be given a value 1 if in the state $X(t)$ enough components are operational for the system to be considered up; otherwise, it is given a value of 0.

Under fairly general conditions that ensure ergodicity (which are analogous to recurrence type properties in Markov chains with countable state space), as $t \rightarrow \infty$, the quantity $\int_{s=0}^t f(X(s))ds/t$ converges to a constant with probability 1. Let us denote this constant by U . Our goal is to estimate this steady-state measure U .

2.1 Standard Estimation Procedures for Steady-State Measures

Let Δ denote a length of time. Define $D_i = \int_{s=(i-1)\Delta}^{i\Delta} f(X(s))ds/\Delta$ and form the estimate $\hat{U} = \sum_{i=1}^n D_i/n$. Then by definition, as $n \rightarrow \infty$, $\hat{U} \rightarrow U$ with probability 1. Also, under some more regularity conditions, we have the central limit theorem (CLT): $\sqrt{n}(\hat{U} - U) \Rightarrow N(0, \sigma^2)$, as $n \rightarrow \infty$, where σ^2 is a variance constant. In that case, if we knew σ^2 , we could construct a $(1 - \delta)\%$ confidence interval (CI) for U . The half width (HW) of this CI will be given by $z_{\delta/2}\sigma/\sqrt{n}$, where $z_{\delta/2}$ is the $100(1 - \delta/2)$ percentile point of the standard normal distribution (see [1]). If the D_i 's were i.i.d., then $\sigma^2 = Var(D_i)$ which would be easy to estimate. If we discard some of the initial D_i 's (i.e., allow the system to reach steady-state), then the D_i 's from then on are (approximately) identically distributed but still not independent. Hence the method of batch means (see [1]) can be used to estimate σ^2 . We now briefly review this procedure.

In the method of batch means we group the D_i 's into batches, each batch having k successive D_i 's, i.e., if b is the number of batches, then $kb = n$. Let $\delta_j = \sum_{i=(j-1)k+1}^{jk} D_i/k$ for $1 \leq j \leq b$. Then we form the estimate $\hat{U} = \sum_{j=1}^b \delta_j/b = \sum_{i=1}^n D_i/n$ which is the same estimator as before. The method of batch means make use of the assumption that for sufficiently large k the δ_j 's are (approximately) normally distributed and uncorrelated. If we discard the first few δ_j 's to allow for the system to reach steady-state, then the δ_j 's are also identically distributed. In that case we again have the CLT, where now $k \rightarrow \infty$. If, in addition, b is large, then $\sigma^2 = Var(\delta_j)$ can easily be estimated.

Consider $f(X(s))$ of the form $1_{\{X(s) \in \mathcal{F}\}}$ (the indicator function), where $\mathcal{F} \subset \mathcal{S}$ is a rare set of states (but with non-zero probability). In this case, most of the D_i 's and thus the δ_j 's will be zero, and therefore, as mentioned in the introduction, it is hard to get accurate estimates. So techniques like importance sampling have to be used. In the following sections we develop a method based on a representation of steady-state measures as a ratio of expectations. The method uses the techniques of batch means, splitting and importance sampling to efficiently estimate the ratio.

2.2 A Ratio Representation of Steady-State Measures

Let $A \subset \mathcal{S}$. As in the introduction, an A – cycle is defined to start whenever the $\{X(t) : t \geq 0\}$ enters A . Let C be the length of an A – cycle. Let ϕ denote the probability dynamics governing the realizations of $\{(X(t), S(t)) : t \geq 0\}$. Let π be the steady-state distribution of $(X(t), S(t))$ at the times when $\{X(t) : t \geq 0\}$ enters A . Then under fairly general recurrence type conditions (which also ensure that the system returns to state A infinitely often) we have that (see [5] and Section 6.9 of [2] for details; see also [7])

$$U = \frac{E_{\pi, \phi}(D)}{E_{\pi, \phi}(C)}, \quad (1)$$

where now $D = \int_{s=0}^C 1_{\{X(s) \in \mathcal{F}\}} ds$ and the subscripts in the expectation denote the initial distribution and the probability dynamics governing the realizations of $\{(X(t), S(t)) : t \geq 0\}$.

To make use of this ratio, we can first run the process for some time so that it reaches steady-state and then run n A -cycles to get n samples of D and C . However, as in the standard estimation procedure described above, there are two problems. First, because the set \mathcal{F} is rare, most of the samples of D will be zero, leading to an inaccurate estimate of $E(D)$. To overcome this we use importance sampling. Second, the samples of D (and C) are identically distributed but not independent. This can be handled using the method of batch means.

2.3 Efficient Simulation of Rare Events Using Importance Sampling

Let ϕ' denote a new probability dynamics that now governs the realizations of $\{(X(t), S(t)) : t \geq 0\}$, such that the probability on the sample paths of $\{(X(t), S(t)) : t \geq 0\}$ using ϕ is absolutely continuous with respect to the probability on the sample paths using ϕ' (absolutely continuous essentially means that any event that had a positive probability of occurrence under the original dynamics also has a positive probability of occurrence under the new dynamics). When importance sampling is used, we can write $E_{\pi, \phi}(D) = E_{\pi, \phi'}(DL)$ (the second subscript in the second expectation indicates that we are using the new probability dynamics ϕ' to generate D and L), where L is the likelihood ratio. The main problem in importance sampling is to choose an easily implementable ϕ' so that $Var_{\pi, \phi'}(DL) \ll Var_{\pi, \phi}(D)$. Then for estimating the ratio we can write Equation 1 as

$$U = \frac{E_{\pi, \phi'}(DL)}{E_{\pi, \phi}(C)} \quad (2)$$

and use the following simulation scheme. We first run a few A -cycles using ϕ so that the system enters steady-state and we are assured of $(X(t), S(t))$ being distributed sufficiently close to π at the start of A -cycles. Then we do a splitting technique (see [16]) in each of the A -cycles, where we do one run using the dynamics ϕ' to get samples of D and L and a second run with the original dynamics ϕ to get a sample of C . The second run also ensures that we again get the distribution π when the system re-enters state A , so that we can start another A -cycle at that point in time. We repeat this procedure to get the samples D_i , L_i and C_i , $1 \leq i \leq n$, of D , L and C , respectively. This approach is very analogous to "measure specific dynamic importance sampling" (MSDIS) [15] for estimating the steady-state unavailability in Markovian systems; importance sampling is used to estimate the expected downtime in a cycle and standard simulation is used to estimate the expected cycle time.

2.4 Variance Estimation Using the Method of Batch Means

Since the A -cycles are not independent, we use the method of batch means to estimate the variance. As before let b be the number of batches and let k be the batch size. Let $\delta_j = \sum_{i=(j-1)k+1}^{jk} D_i L_i / k$ and $\gamma_j = \sum_{i=(j-1)k+1}^{jk} C_i / k$. Then we form the estimate

$$\hat{U} = \frac{\hat{\delta}}{\hat{\gamma}}, \quad (3)$$

where $\hat{\delta} = \sum_{j=1}^b \delta_j / b = \sum_{i=1}^n D_i L_i / n$, and $\hat{\gamma} = \sum_{j=1}^b \gamma_j / b = \sum_{i=1}^n C_i / n$.

In the steady-state, for sufficiently large k , the δ_j 's (and γ_j 's) are uncorrelated. We introduce the generic random variables δ and γ having the same distributions as δ_j and γ_j , respectively. It follows that $E(\hat{\delta}) = E_{\pi, \phi'}(\delta)$ and $E(\hat{\gamma}) = E_{\pi, \phi}(\gamma)$. Also, $Var(\hat{\delta}) = Var_{\pi, \phi'}(\delta)/b$, $Var(\hat{\gamma}) = Var_{\pi, \phi}(\gamma)/b$ and $Cov(\hat{\delta}, \hat{\gamma}) = Cov_{\pi, \phi'}(\delta, \gamma)/b$. The subscripts in the covariance term indicate that for each i , (D_i, L_i) (used in the δ_j 's) and C_i (used in the γ_j 's) are sampled using the same starting state (i.e., distributed according to π), and using ϕ' and ϕ , respectively. Finally, from the CLT we have $\sqrt{b}(\hat{U} - U) \approx N(0, \sigma^2)$ for large k and b , where (analogous to the regenerative method)

$$\sigma^2 = \frac{Var_{\pi, \phi'}(\delta) + U^2 Var_{\pi, \phi}(\gamma) - 2UCov_{\pi, \phi'}(\delta, \gamma)}{E_{\pi, \phi}^2(\gamma)} \quad (4)$$

3 Estimation of Unavailability in Highly Dependable Systems

The class of highly dependable systems considered in this paper is that composed of highly reliable components, i.e., the mean time between failures (MTBF) of its components is orders of magnitude larger than their repair time. High dependability can also be achieved by increasing the redundancy level of less reliable components; however, here we are not concerned with this type of systems.

Without loss of generality, we consider models of highly dependable systems in which a component may be in one of only two states: operational and failed. When a component fails in a given mode, it may cause other components to fail with some probability (failure propagation). Different sets of components may be affected at different failure modes. Failed components are repaired by one or more repair facilities according to some arbitrary service discipline. Basically, these models are similar to those that can be handled using the SAVE package [14]. However, unlike SAVE, we allow general distributions for failure and repair times, and repairs are assumed not to be instantaneous. Let $G_i(x)$ denote the failure distribution of component i . Then its hazard rate function is given by $h_i(x) = g_i(x)/\bar{G}_i(x)$, where $g_i(x)$ is the probability density function corresponding to $G_i(x)$ and $\bar{G}_i(x) = 1 - G_i(x)$. We further parameterize the hazard rate function in terms of a small (but positive) parameter ϵ and assume it is bounded, such that $h_i(x) \leq \lambda_i \epsilon^{b_i}$, $x \geq 0$, where $0 < \lambda_i < \infty$ and $b_i \geq 1$. As will be further discussed in Section 3.1, the assumption of bounded failure hazard rate functions (which holds for many, including phase-type, distributions) is necessary in order to use the uniformization approach. Weibull is not included in the class of bounded hazard rate distributions, however, it can be arbitrarily well approximated by appropriately bounding its hazard rate function; this will enable us to experiment with an increasing failure rate Weibull distribution, as will be described in Section 4.3.

3.1 A Uniformization-Based Importance Sampling Approach

The uniformization technique (also known as randomization [19]) can be used to sample from general distributions (e.g., nonhomogeneous Poisson processes) with bounded hazard rate functions. Such distributions include Markovian phase-type, but exclude discrete, uniform and Weibull distributions. To illustrate, consider simulating the nonhomogeneous Poisson process with a bounded hazard rate $h(t) \leq \beta$, where β is a constant rate. We generate the event times $\{T_k\}$, $k = 1, 2, \dots$, of a homogeneous Poisson process at rate β (β is called the uniformization rate.)

T_k is accepted as an actual event of the simulated process with probability $h(T_k)/\beta$ (real event), otherwise, it is rejected (pseudo event). The acceptance/rejection test is performed at consecutive uniformization events until an event is accepted, in which case the same procedure is repeated to generate the next (real) event of the simulated nonhomogeneous Poisson process.

In the context of reliability estimation, uniformization can be used to implement importance sampling in Non-Markovian models as described in [24]. In a similar way, it can also be used for unavailability estimation as we will describe in this section.

Consider a system with N components. At any time t , let $O(t)$ be the set of operational components, and denote by $a_i(t), i \in O(t)$, the age of component i (i.e., the time since it last became operational). Define the failure rate of component i at time t to be $\lambda_i(t)$, then

$$\lambda_i(t) = \begin{cases} h_i(a_i(t)), & \text{if } i \in O(t) \\ 0, & \text{otherwise} \end{cases}$$

The total failure rate at time t is given by $\lambda_F(t) = \sum_{i=1}^N \lambda_i(t)$.

Let $\{\tau_n\}, n = 1, 2, \dots$, corresponds to the real event times in the system (i.e., failure and repair events). At real event τ_n , let $\lambda_F(t), t \geq \tau_n$, be bounded by a constant rate, say, β_n (i.e., $\lambda_F(t) \leq \beta_n, t \geq \tau_n$). Then the time to next failure event in the system can be sampled by successively generating the Poisson (uniformization) event times $\{T_{nk}\}, k = 1, 2, \dots$, at rate β_n , and performing the acceptance/rejection test until an event is accepted. Event T_{nk} is rejected with probability $1 - \lambda_F(T_{nk})/\beta_n$, in which case the next uniformization event is generated. Otherwise, T_{nk} is accepted as the next failure event time. This procedure is repeated at every real event (failure or repair) to generate the time to next failure event. In the original system, the expected time to next failure event (β_n^{-1}) is much larger than the expected time to next repair event, if any. Therefore, a system failure is very unlikely to occur.

As we simulate the system, repairs are sampled from their original distributions, which, therefore, are not restricted to the class of bounded hazard rate functions. In particular, this allows arbitrary repair time distributions, including general discrete and uniform distributions. It is assumed that failure propagation probabilities are not changed in the simulated system. It follows that uniformization (with importance sampling) is used only to simulate the time of failure events.

To implement importance sampling using uniformization we do the following. At real event τ_n (during a repair), we can simply increase the uniformization rate, β_n , and fix the acceptance probability at some level, say, p_n ,

such that $\alpha_n \stackrel{\text{def}}{=} \beta_n p_n$ (this is the effective rate at which the next failure event is generated; we call it the "biasing level") is of the same order as the repair "rate". This will increase the probability of subsequent failure events leading to a system failure.

Upon the occurrence of a failure event in the original system, say, at τ_n , component i is selected as the failed component with probability $p_{ni} = \lambda_i(\tau_n)/\lambda_F(\tau_n)$. However, with importance sampling, this probability could be changed. For example, in "balanced failure biasing" we equalize the failure probability for all operational components. (We are using the term "balanced failure biasing" in a generic sense. "Balanced failure biasing" was originally defined for the case of Markovian systems and consisted of making all failure transitions from a state occur with equal probabilities [29]. This is different from making all operational components fail with the same probability, as

in systems with failure propagation, the number of failure transitions possible from a state may be different from the number of operational components in that state. However, it can easily be shown using techniques in [29], that for Markovian systems in which the failure propagation probabilities are considerably higher than the failure transition probabilities, the version of balanced failure biasing based on the number of operational components is as effective. We are generalizing this version to non-Markovian systems, assuming that the failure propagation probabilities are not very small.) In this case, component $i \in O(\tau_n)$ is selected to fail with probability $p_{ni} = 1/|O(\tau_n)|$. In addition to being simple and robust, "balanced failure biasing" is known to be provably effective in the context of unreliability estimation [18], [29], [30].

The likelihood ratio is computed recursively, by updating it only at pseudo and failure event times as follows. Let $L_k, k = 0, 1, 2, \dots$, be the likelihood ratio at the k -th (pseudo or failure) event, at time t_k . Then $L_0 = 1$,

$$L_k = L_{k-1} \left(\frac{\lambda_F(t_k)/\beta_k}{p_k} \right) \left(\frac{\lambda_i(t_k)/\lambda_F(t_k)}{p_{ki}} \right) \quad (5)$$

if the k -th event is a type i failure event and

$$L_k = L_{k-1} \left(\frac{1 - \lambda_F(t_k)/\beta_k}{1 - p_k} \right) \quad (6)$$

if the k -th event is a pseudo event. In other words, the likelihood ratio is updated by a factor equal to the ratio of the probability of the k -th event in the original and simulated systems, respectively. Notice in the above equation that β_k, p_k and p_{ki} can be changed at pseudo events; however, in our implementation (as described above) they are changed only at real (failure or repair) events. Heuristics for choosing β_k, p_k and p_{ki} , as well as other practical considerations, are discussed in Section 3.2. For unreliability estimation [18], it is shown that under reasonable assumptions and appropriate heuristics, the above method is provably effective, i.e., it yields estimates with bounded relative error. In Section 4 we experimentally demonstrate the effectiveness of our method for unavailability estimation. However, theoretical results to establish the property of "bounded relative error" are not yet available.

3.2 Implementation Issues

In this section we consider specific implementation issues in the estimation of steady-state unavailability using uniformization and importance sampling as discussed in the previous sections. In our implementation we use CSIM [28], a process-oriented simulation language based on the C programming language.

Following our discussion in Section 2.3 as applied to highly dependable systems, we define an A -cycle to be a sample path between two successive entries into the fully operational state (in which all system components are operational). Specifically, in our context, the set A constitutes all possible components' ages upon entering the fully operational state. (Notice that upon entering A at least one component has an age identical to zero).

The ratio representation of the steady-state system unavailability U is given by Equation 2, where, C is the length of an A -cycle in the original system, D is the total downtime in an A -cycle in the simulated system (with importance sampling) and L is the corresponding likelihood ratio. An estimate of the steady-state unavailability is given by Equation 3, where $\hat{\delta}$ and $\hat{\gamma}$ are estimates of $E_{\pi, \phi'}(DL)$ and $E_{\pi, \phi}(C)$, respectively. Recall that the subscripts π and ϕ (ϕ') indicate that the expectation is taken over A -cycles

having the typical steady-state entry distribution π , with respect to the original (new) probability dynamics.

The system is simulated sufficiently long, under the original probability dynamics ϕ , until it (approximately) reaches the steady-state. From that point, let n be the number of A -cycles used to obtain the estimate $\hat{\gamma}$. Since, in general, successive A -cycles are not independent, we use the method of batch means to get an estimate $\widehat{Var}(\gamma)$ of the variance $Var_{\pi, \phi}(\gamma)$. Following the same notation as in Section 2.4, let b be the number of batches, each having $k(= n/b)$ A -cycles. Let C_i be the length of the i -th A -cycle, and for batch j , let $\gamma_j = \sum_{i=(j-1)k+1}^{jk} C_i/k$. Then we have

$$\hat{\gamma} = \sum_{j=1}^b \gamma_j/b = \sum_{i=1}^n C_i/n$$

and

$$\widehat{Var}(\gamma) = \sum_{j=1}^b (\gamma_j - \hat{\gamma})^2 / (b-1).$$

k should be sufficiently large, so as to eliminate dependence between successive batches. Our experimental results in Section 4.1 indicate that k need not be large.

To obtain the estimate $\hat{\delta}$ we do the following. For each A -cycle that we simulate under the original probability dynamics (we call this an "original" A -cycle), we simulate the same A -cycle (i.e., starting with the same initial components' ages) under the new probability dynamics, i.e., with importance sampling (we call this a "biased" A -cycle). Usually more effort is needed to estimate $E_{\pi, \phi'}(DL)$ than that to estimate $E_{\pi, \phi}(C)$ (to the same degree of relative error). This is because of three reasons: 1) there are typically more events in a biased cycle, 2) the L must be computed in a biased cycle and 3) the ratio of the variance to the quantity to be estimated is higher for a biased cycle. Therefore, we run several, say, m , biased A -cycles for each original A -cycle. In this way, we use more cycles, namely, $n' = mn$, to obtain the estimate $\hat{\delta}$. If we use the same number of batches, b , to get an estimate $\widehat{Var}(\delta)$ of the variance $Var_{\pi, \phi'}(\delta)$, then the number of biased A -cycles per batch is equal to $k' = mk$. Let $D_{i,s}$ be the total system downtime in the s -th run of the i -th biased A -cycle, and $L_{i,s}$ is the corresponding likelihood ratio. For batch j , let $\delta_j = \frac{1}{k'} \sum_{i=(j-1)k+1}^{jk} \sum_{s=1}^m D_{i,s} L_{i,s}$. It follows that

$$\hat{\delta} = \sum_{j=1}^b \delta_j/b = \frac{1}{n'} \sum_{i=1}^n \sum_{s=1}^m D_{i,s} L_{i,s}$$

and

$$\widehat{Var}(\delta) = \sum_{j=1}^b (\delta_j - \hat{\delta})^2 / (b-1).$$

Furthermore, an estimate of the covariance $Cov_{\pi, \phi'}(\delta, \gamma)$ is given by

$$\widehat{Cov}(\delta, \gamma) = \sum_{j=1}^b (\delta_j - \hat{\delta})(\gamma_j - \hat{\gamma}) / (b-1).$$

An estimate $\hat{\sigma}^2$ of the variance σ^2 of the estimator \hat{U} can now be obtained from Equation 4.

With importance sampling, our goal is to increase the frequency of A -cycles that contain typical system failures. Our heuristic is similar to that for regenerative models in [25]. In a biased A -cycle, upon the occurrence of the first component failure, we activate failure biasing to accelerate subsequent component failure events relative to subsequent repair completion events. Failure biasing is continued until either system failure or the end of the current A -cycle. In doing so, we increase the probability of system failure in biased A -cycles. More specifically, let r_n be equal to the maximum of the expected repair times (different from the expected remaining repair times) of components undergoing repair at the time of the n -th real (failure or repair) event. As discussed in Section 3.1, if failure biasing is activated at the n -th real event, then the uniformization rate β_n and the acceptance probability p_n could be chosen such that α_n (recall that $\alpha_n = \beta_n p_n$) is equal to $1/r_n$. (For a single repairman and exponential repairs, this choice is equivalent to setting the probability of a failure before repair to 0.5.) This heuristic does not necessarily lead to the most variance reduction, however, it is quite effective and robust. In our CSIM implementation, we were not able to determine the components currently undergoing repair; therefore, we set r_n equal to the maximum of the expected repair times of all components that are either undergoing or waiting for repair. We did not anticipate the following problem, however. When the time to the next scheduled repair completion is much larger than r_n , the biasing level α_n (as determined by $1/r_n$) may become excessively high (relative to the time to the next scheduled repair completion), causing untypical failure sequences. This tends to significantly increase the variability of the likelihood ratio, leading to unstable estimates. We overcome this problem by setting r_n equal to the maximum of the actual repair times (in our CSIM implementation, as soon as a component fails, its actual repair time is generated; note the difference from the actual remaining repair time) of all components either undergoing or waiting for repair, whenever the latter exceeds the maximum of the expected repair times by several times (say, 5 times). The above heuristics have been shown to work well, as will be demonstrated in Section 4.

At every uniformization event, the ages of all operational components are adjusted and their hazard rates are determined. This is required to update the likelihood ratio depending on whether the event is accepted or rejected (see Equation 5 and 6). Since repair times are unchanged in the simulated system, no updating of the likelihood ratio is necessary at repair events.

With a given appropriate biasing level ($\beta_n p_n$), there is freedom in choosing the uniformization rate β_n (and hence the acceptance probability p_n). Experiments described in [24] show that higher β_n and lower p_n result in a less noisy estimation of the likelihood ratio, and hence somewhat lower variance of the resulting estimate. On the other hand, higher β_n and lower p_n results in an excessive number of rejected (pseudo) events, i.e., very inefficient generation of failure events. An appropriate uniformization rate should be chosen low enough to limit excessive generation of pseudo events, yet high enough to preserve accuracy. Experiments described in [24] show that $\beta_n = 5r_n$ is a good choice (p_n is then determined by the chosen biasing level).

Once a uniformization event is accepted as a failure, then one of the operational components is selected as the failing component. In our experiments in Section 4 we use "balanced failure biasing" (as described in Section 3.1). We also balance the first component failure in a biased A -cycle. This is quite important, particularly for "unbalanced" systems (e.g., when component reliabilities are

of different orders of magnitude).

4 Experimental Results

In this section we experiment with our method and demonstrate its effectiveness for the estimation of steady-state unavailability in highly dependable systems. We use small and large examples with general failure and repair time distributions. In special cases, some examples conform to the class of models having a “product form” solution. In these cases the results are invariant with respect to failure time distributions having the same mean. Therefore, we are able to validate with numerical (non-simulation) results obtained using SAVE. Assuming exponential failure and repair time distributions, models not having a product form solution can also be validated using SAVE.

Some of our design choices are based on earlier work (e.g., in [15], [25] and [24]). For example, when failure biasing is activated, we set the biasing level $\beta_n p_n$ equal to r_n (as defined in Section 3.2). An appropriate uniformization rate β_n should be chosen low enough to limit excessive generation of pseudo events, yet high enough to preserve accuracy. Experiments in [24] suggests that $\beta_n = 5r_n$ is a good choice (and hence the acceptance probability (p_n) is set to 0.2).

In Section 4.1 a small machine repairman model is used to experiment with the batch size used in the batch means method. The results indicate that the accuracy of the estimates (at least in this example) is insensitive to the batch size. In Section 4.2 another small example is used to study the efficiency of the method under a variety of circumstances. This is accomplished by varying the order of ϵ (used to parameterize the failure hazard rate functions, see Section 3) and by experimenting with “balanced” and “unbalanced” systems. In Section 4.3 a large example is used to demonstrate the effectiveness of our method when dealing with large and complex systems. In this example we experiment with different failure time distributions, namely, Erlang, Weibull, exponential and hyperexponential. With exponential repair times (at the same rate), FCFS (first come first served) repair discipline and no failure propagation, this example has a product form solution. In this case we can validate our simulation results with numerical solutions obtained using SAVE.

For each table entry in all experiments of this section, we run a total of $n = 64000$ original A – cycles, which are used to estimate the expected cycle length $\hat{\gamma}$. For each original A – cycle, we run $m = 4$ biased A – cycles. It follows that the total number of biased A – cycles used to estimate the expected downtime $\hat{\delta}$, is $n' = 256000$. For all, but the experiment in Section 4.1, we fix the number of batches b to 1000. Accordingly, the batch size k (k') is fixed at 64 (256) original (biased) A – cycles. In each table entry we display the estimate (from simulation) of the steady-state unavailability, along with its 99% half-width confidence interval as a percentage of the point estimate.

Except for special cases, the models considered in this section cannot be evaluated either analytically or numerically. Because of their high dependability feature, standard (naive) simulation is also not practical. As a result, effectiveness studies in this section demonstrate the usefulness of our method, as it considerably extends the class of models in which importance sampling can be used to evaluate various dependability measures.

4.1 Batch Size

As described in Sections 2.4 and 3.2, before collecting the samples C_i (from original A – cycles) and $D_i L_i$ (from biased A – cycles) for the method of batch means, the simulation should be run long enough to reach its steady-state dynamics. This can be accomplished by discarding the first

few batches of the simulation. Furthermore, the batch size, k , needs to be sufficiently large, so as to (approximately) eliminate dependence between successive batches. In this section we use a small example to experiment with the batch size.

We consider a machine repairman model with two component types and two components of each type. The system is considered operational as long as one component of each type is operational. All components have the same failure time distribution; namely, Erlang with two stages, each having a rate equal to 0.0002 per hour. Thus, the MTBF of individual components is 10000 hours. When components fail, they get repaired by a single repairman according to FCFS discipline. For all components, we assume that the repair time distribution is exponential with a mean equal to 1.0 hour. With exponential repairs, this model has a product form solution, which depends on the failure distribution only through its mean. Therefore, we can validate our results by solving the same example, with exponential failure times (having the same MTBF, i.e., 10000 hours). Using SAVE, a numerical estimate of the steady-state unavailability for this model is given by 4.0×10^{-8} .

For the same total number $n = 64000$ ($n' = 256000$) of original (biased) A – cycles, in Table 1 we successively half the number of batches b from 64000 to 250. Accordingly, the batch size k is successively doubled from $k = 1$ ($k' = 4$) to $k = 256$ ($k' = 1024$). Note that the estimates in the table compare well with the above numerical result from SAVE. Observe that the confidence interval widths do not depend, in any significant way, on the batch size. (This being the case also for the smallest batch size of one original A – cycle.) This is an indication that, in the steady-state, consecutive A – cycles are almost uncorrelated. While it seems to be the case in this particular example, this is not generally true for other systems. However, additional experiments (not reported here) suggest that “near independence” of consecutive A – cycles (as defined in Section 3.2) may be a feature of highly dependable systems. In all subsequent experiments we set the batch size k (k') to 64 (256); this is large enough to achieve approximate independence between successive batches.

4.2 A Small Example

In this section we provide empirical results illustrating the desirable “bounded relative error” property of our method. We show that as system failure becomes rarer, we can still estimate the steady-state unavailability with the same accuracy; also, when the system is “unbalanced”. In Section 3 we parameterized the failure hazard rate functions in terms of ϵ . In the following example we parameterize the failure time distributions in terms of their inverse mean ($1/\text{MTBF}$), which we denote by ϵ . By varying ϵ we change component reliabilities (i.e., their MTBF), and, hence, the steady-state system unavailability. Also, by having different component types with different ϵ , we create examples of “unbalanced” systems.

Again, we consider a machine repairman model with two types of components; 3 components of Type I and 2 components of Type II. The system is considered operational as long as one component of each type is operational. Failure time distributions are either Erlang or hyperexponential and may be different for each component type (as specified below). Type II components have a higher (preemptive-resume) priority at the (single) repair facility. The repair time distribution is constant (deterministic), at 1.0 hour, for Type I components and uniform, between 0.0 and 1.0 hour, for Type II components. For the same example, we perform the same set of experiments with and without failure propagation. If failure propagation is con-

sidered, then with probability 0.25, a failure of Type II component causes two components of Type I to fail (all operational Type I components fail if they are equal or less than two).

We experiment with non-exponential failure time distributions. Specifically, *Erlang* $(2, \lambda)$ (two stages, each with a rate λ per hour) and *Hyperexponential* $(p, \lambda_1, \lambda_2)$ (two stages, with probabilities p and $1 - p$ at rates λ_1 per hour and λ_2 per hour, respectively). We parameterize these distributions in terms of their inverse mean (ϵ) as follows:

- $E_2(\epsilon) \stackrel{\text{def}}{=} \text{Erlang}(2, 2\epsilon)$, having a CV (coefficient of variation) = 0.707,
- $H_2(\epsilon) \stackrel{\text{def}}{=} \text{Hyperexponential}(0.2727, 0.3342\epsilon, 4.01\epsilon)$, having a CV = 2.0.

We simulate the system using our method with the batch means parameters given earlier in Section 4. For two values of ϵ , namely, 10^{-2} and 10^{-4} , in Table 2 we give estimates of the steady-state unavailability for the following 4 combinations of components' failure time distributions (without failure propagation):

- C1: $E_2(\epsilon)$ for Type I and $E_2(\epsilon)$ for Type II
- C2: $E_2(\epsilon)$ for Type I and $E_2(\epsilon^{1.5})$ for Type II
- C3: $H_2(\epsilon)$ for Type I and $E_2(\epsilon)$ for Type II
- C4: $H_2(\epsilon)$ for Type I and $E_2(\epsilon^{1.5})$ for Type II.

In C1 and C3 the system is "balanced". In C2 and C4 the system is "unbalanced". Notice that the relative error of the estimates is about the same for "balanced" and "unbalanced" systems and is independent of the value of ϵ . For the experiment in C1, we ran standard simulation with the same batch means parameters. For $\epsilon = 10^{-2}$, it produced an estimate having relative error $\pm 13.41\%$ (compared with $\pm 5.26\%$ using importance sampling). This means that standard simulation should be run 6.5 times longer to achieve about the same accuracy obtained with importance sampling. No failures were observed for $\epsilon = 10^{-4}$ using standard simulation.

With failure propagation, we ran the same experiments using our method. The resulting estimates are given in Table 3. Again, the accuracy of the estimates is quite consistent throughout the table.

4.3 A Large Example

It remains to show that the method described here is also feasible and effective when dealing with large and complex highly dependable systems. In this section we consider a large example with many types of components. We experiment with different failure time distributions, such as Erlang, Weibull, exponential and hyperexponential. Without failure propagation, the example falls within the class of product form models and validation with numerical (non-simulation) results is possible.

The system we consider is based on a model of a fairly complex computing system (also considered in [24]). The computing system is composed of two sets of processors with 2 processors per set, two sets of controllers with 2 controllers per set, and 6 clusters of disks, each consisting of 4 disk units. In a disk cluster, data is replicated so that one disk can fail without affecting the system. The "primary" data on a disk is replicated such that one third is on each of the other three disks in the same cluster. Thus, one disk in each cluster can be down without losing access to the data. Components are repaired by a single repairman according to a FCFS discipline. The system is defined to be operational if all data is accessible to both processor types, which means that at least one processor of each type, one controller in each set, and 3 out of 4

disk units in each of the 6 disk clusters are operational. Operational components continue to fail at the given rates when the system is failed. When failure propagation is considered, a failing processor in any of the two sets causes one processor in the other set to fail with probability 0.1.

All repair time distributions are exponential with mean 1 hour (however, any general distribution could be allowed). All of the component failure times follow the same distribution, with the same coefficient of variation (CV), but possibly with different means (for the different types of components). The MTBF for processors, controllers and disks are assumed to be 200000, 200000 and 600000 hours, respectively. We experiment with four failure time distributions; namely, Erlang with 2 stages (CV = 0.707), Weibull with a shape parameter equal to 1.25 (CV = 0.805), exponential (CV = 1.0), and hyperexponential with 2 stages (CV = 2.0). For the Weibull distribution, the scale parameters corresponding to the overall means 200000 and 600000 are equal to 2.1634×10^{-7} and 5.4793×10^{-8} , respectively. For the hyperexponential distribution, the parameters are as follows: a probability equals to 0.2727 of branching to the first stage with a mean 600000 (1800000) and a probability equals to 0.7273 of branching to the second stage with a mean 50000 (150000), corresponding to the overall mean 200000 (600000).

Notice that uniformization cannot be used to sample from a Weibull distribution, since its hazard rate function is not bounded. However, as we do in this example, random variates from an IFR (increasing failure rate) Weibull can be arbitrarily well approximated by sampling (using uniformization) from another distribution having a bounded hazard rate function. This approximation is obtained by simply bounding the hazard rate function $h(t)$ of the (IFR) Weibull at $\lambda_m = h(t_m)$ beyond a sufficiently large time t_m , such that $\bar{G}(t_m)$ is extremely small, say, 10^{-20} . If λ_m is not too high compared to other hazard rates in the system, then a reasonably efficient uniformization rate can be used to generate failure events.

We simulate the described system using our method (with the batch means parameters given in Section 4) to get estimates of the steady-state unavailability for two sets of failure distributions. In the first set (Set I) we use the components' MTBFs given above. In the second set (Set II) we reduce all components' MTBFs by a factor of 10 (i.e., we use less reliable components). Accordingly, all means in the Erlang, exponential and hyperexponential stages are also reduced by a factor of 10. For the Weibull distribution, the scale parameters corresponding to the overall means 20000 and 60000 are equal to 3.847×10^{-6} and 9.744×10^{-7} , respectively.

Without failure propagation, the above example has a product form solution, which depends on the failure time distributions only through their means. It follows that the steady-state unavailability is the same for all failure time distributions having the same mean. Furthermore, using SAVE, we can obtain numerical estimates, which are given by 4.0×10^{-10} and 4.0×10^{-8} for failure data sets, I and II, respectively.

In Table 4 we give estimates of the steady-state unavailability for the system without failure propagation, for both sets of failure data, I and II. All relative errors in this table are less than $\pm 10\%$. Notice the agreement among the estimates for different failure distributions, on one hand, and with the above results from SAVE, on the other hand. For the hyperexponential failure distribution, the estimates are slightly less accurate than those corresponding to failure distributions with a lower coefficient of variation. For this (hyperexponential) case, standard

simulation (with the same total number of cycles) produced unstable estimates having relative errors of $\pm 260\%$ and $\pm 111\%$ for failure data Sets I and II, respectively. In fact, these confidence intervals are meaningless, since they contain negative values.

In Table 5 we give estimates of the steady-state unavailability for the system with failure propagation, for both sets of failure data, I and II. A preemptive-resume discipline at the repair facility is now assumed, with processors having the highest priority and disks having the lowest priority. Notice that the steady-state unavailability is affected only a little by the failure time distribution. However, as expected, the point estimates are consistently higher than those in Table 4. Because this is a different system, the confidence intervals happen to be slightly wider than those in Table 4.

Again, in each of the Tables 4 and 5, the accuracies of the estimates are about the same, regardless of the failure distributions or their means. These empirical results are consistent with the conjecture that our method produces estimates of steady-state unavailability having bounded relative error.

5 Conclusions

This paper has considered the problem of estimating steady-state quantities for models of highly dependable computing systems in which the component failure and repair times have general distributions. Such models are analytically and numerically intractable; simulation is the only possible means of analysis. However, standard simulation is inefficient when system failure events are rare and importance sampling needs to be used to speedup the simulation. Earlier importance sampling approaches are effective for steady-state estimation only when the failure time distributions are exponentially distributed, in which case the regenerative structure of the model can be exploited. When failure times are generally distributed, no such regenerations exist. However, a ratio representation in terms of non-i.i.d. cycles still exists for steady-state quantities. Using this representation, a splitting technique can be devised in which importance sampling is used to estimate the expected downtime during a cycle and standard simulation is used to estimate the expected cycle length. The particular method of importance sampling that we use is based on uniformization, and is provably effective for estimating certain transient quantities within this class of models. Experiments showed the method to be effective in practice for estimating the steady-state unavailability. For the efficient estimation of the steady-state mean time between failures, the reader is referred to [12].

As a result of this work, the class of highly dependable systems that can be efficiently simulated to estimate steady-state measures is greatly broadened.

6 Acknowledgements

The work of Peter W. Glynn was supported by grants ARO DAAL03-91-G-0101 and NSF DDM-9101580.

References

- [1] Bratley, P., B.L. Fox and L.E. Schrage. 1987. *A Guide to Simulation*, Second Edition. New York: Springer Verlag.
- [2] Breiman, L. 1968. *Probability*. Reading, MA: Addison-Wesley.
- [3] Carrasco, J.A. 1991. Failure distance-based simulation of repairable fault-tolerant systems. *Proceedings of the fifth International Conference on Modelling Techniques and Tools for Computer Performance Evaluation*, 337-351.
- [4] Carrasco, J.A. 1991. Efficient transient simulation of failure/repair Markovian models. *Proceedings of the Tenth Symposium on Reliable and Distributed Computing*, IEEE Press, 152-161.
- [5] Cogburn, R. 1975. A uniform theory for sums of Markov chain transition probabilities. *The Annals of Probability* 3: 191-214
- [6] Crane, M.A. and D.L. Iglehart. 1975. Simulating stable stochastic systems, III: Regenerative processes and discrete event simulations. *Operations Research* 23: 33-45.
- [7] Fox, B.L. and P.W. Glynn. 1989. Estimating time averages via randomly-spaced observations. *Probability in the Engineering and Informational Sciences* 3: 299-318.
- [8] Fox, B.L. and P.W. Glynn. 1989. Replication schemes for limiting expectations. *SIAM Journal of Applied Mathematics* 47: 186-214.
- [9] Frater, M.R., T.M. Lennon, and B.D.O. Anderson. 1991. Optimally efficient estimation of the statistics of rare events in queueing networks. *IEEE Transactions on Automatic Control* 36: 1395-1405.
- [10] Geist, R.M. and M.K. Smotherman. 1989. Ultrahigh reliability estimates through simulation. *Proceedings of the Annual Reliability and Maintainability Symposium*, IEEE Press, 350-355.
- [11] Glynn, P.W. 1989. A GSMP formalism for discrete event systems. *Proceedings of the IEEE* 77: 14-23.
- [12] Glynn, P.W., P. Heidelberger, V.F. Nicola and P. Shahabuddin. 1993. Efficient estimation of steady state measures in non-regenerative dependability models. To appear in *Proceedings of the 1993 Winter Simulation Conference*, IEEE Press.
- [13] Glynn, P.W. and D.L. Iglehart. 1989. Importance sampling for stochastic simulations. *Management Science* 35: 1367-1392.
- [14] Goyal, A. and S.S. Lavenberg. 1987. Modeling and analysis of computer system availability. *IBM Journal of Research and Development* 31: 651-664.
- [15] Goyal, A., P. Shahabuddin, P. Heidelberger, V.F. Nicola and P.W. Glynn. 1992. A unified framework for simulating Markovian models of highly reliable systems. *IEEE Transactions on Computers* C-41: 36-51.
- [16] Hammersley, J.M. and D.C. Handscomb. 1964. *Monte Carlo Methods*. London: Methuen and Co., Ltd.
- [17] Heidelberger, P., V.F. Nicola and P. Shahabuddin. 1992. Simultaneous and efficient simulation of highly dependable systems with different underlying distributions. *Proceedings of the 1992 Winter Simulation Conference*, IEEE Press, 458-465.
- [18] Heidelberger, P., V.F. Nicola and P. Shahabuddin. 1993. Bounded relative error in estimating transient measures in highly dependable non-Markovian Systems. IBM Research Report RC 18794, Yorktown Heights, New York. Submitted.
- [19] Jensen, A. 1953. Markov chains as an aid in the study of Markov processes. *Skand. Aktuarietidskr.* 36: 87-91.
- [20] Juneja, S. and P. Shahabuddin. 1992. Fast simulation of Markovian reliability/availability models with general repair policies. *Proceedings of the Twenty-Second International Symposium on Fault-Tolerant Computing*, IEEE Computer Society Press, 150-159.
- [21] Lewis, E.E. and F. Bohm. 1984. Monte Carlo simulation of Markov unreliability models. *Nuclear Engineering and Design* 77: 49-62.

- [22] Moorsel, A.P.A. van, B.R. Haverkort and I.G. Niemegeers. 1991. Fault injection simulation: A variance reduction technique for systems with rare events. *Dependable Computing for Critical Applications 2*, Springer-Verlag, 115-134.
- [23] Nakayama, M.K. 1991. *Simulation of Highly Reliable Markovian and Non-Markovian Systems*. Ph.D. Thesis, Department of Operations Research, Stanford University, California.
- [24] Nicola, V.F., P. Heidelberger and P. Shahabuddin. 1992. Uniformization and exponential transformation: Techniques for fast simulation of highly dependable non-Markovian systems. *Proceedings of the Twenty-Second International Symposium on Fault-Tolerant Computing*, IEEE Computer Society Press, 130-139.
- [25] Nicola, V.F., M. Nakayama, P. Heidelberger and A. Goyal. 1990. Fast simulation of dependability models with general failure, repair and maintenance processes. *Proceedings of the Twentieth International Symposium on Fault-Tolerant Computing*, IEEE Computer Society Press, 491-498.
- [26] Parekh, S. and J. Walrand. 1989. A quick simulation method for excessive backlogs in networks of queues. *IEEE Transactions on Automatic Control* 34: 54-56.
- [27] Sadowsky, J.S. 1991. Large deviations and efficient simulation of excessive backlogs in a GI/G/m queue. *IEEE Transactions on Automatic Control* 36: 1383-1394.
- [28] Schwetman, H. 1988. Using CSIM to model complex systems. *Proceedings of the 1988 Winter Simulation Conference*, IEEE Press, 491-499.
- [29] Shahabuddin, P. 1991. Importance sampling for the simulation of highly reliable Markovian systems. IBM Research Report RC 16729, Yorktown Heights, New York. To appear in *Management Science*.
- [30] Shahabuddin, P. 1993. Fast transient simulation of Markovian models of highly dependable systems. IBM Research Report RC 18587, Yorktown Heights, New York. To appear in *Proceedings of the 16th International Symposium on Computer Performance Modelling, Measurement and Evaluation*.

k	1	2	4	8	16	32	64	128	256
k'	4	8	16	32	64	128	256	512	1024
b	64000	32000	16000	8000	4000	2000	1000	500	250
\hat{U}	3.962 ± 3.20%	3.959 ± 3.17%	3.961 ± 3.16%	3.969 ± 3.15%	3.953 ± 3.10%	3.952 ± 3.06%	3.956 ± 3.09%	3.942 ± 2.96%	3.967 ± 3.05%

Table 1: Estimates of steady-state unavailability ($\times 10^8$) in a machine repairman model (experiments with batch size).

Combination	C1	C2	C3	C4
Type I Fail. Dist.	$E_2(\epsilon)$	$E_2(\epsilon)$	$H_2(\epsilon)$	$H_2(\epsilon)$
Type II Fail. Dist.	$E_2(\epsilon)$	$E_2(\epsilon^{1.5})$	$E_2(\epsilon)$	$E_2(\epsilon^{1.5})$
$\epsilon = 10^{-2}$	3.399×10^{-5} ± 5.26%	1.317×10^{-6} ± 3.37%	3.387×10^{-5} ± 5.56%	1.401×10^{-6} ± 6.07%
$\epsilon = 10^{-4}$	3.422×10^{-9} ± 5.56%	1.336×10^{-12} ± 3.62%	3.282×10^{-9} ± 5.95%	1.379×10^{-12} ± 6.42%

Table 2: Estimates of steady-state unavailability in a machine repairman model (without failure propagation).

Combination	C1	C2	C3	C4
Type I Fail. Dist.	$E_2(\epsilon)$	$E_2(\epsilon)$	$H_2(\epsilon)$	$H_2(\epsilon)$
Type II Fail. Dist.	$E_2(\epsilon)$	$E_2(\epsilon^{1.5})$	$E_2(\epsilon)$	$E_2(\epsilon^{1.5})$
$\epsilon = 10^{-2}$	2.605×10^{-4} ± 2.30%	2.195×10^{-5} ± 2.70%	3.570×10^{-4} ± 2.74%	2.351×10^{-5} ± 3.35%
$\epsilon = 10^{-4}$	2.625×10^{-8} ± 2.45%	2.080×10^{-10} ± 3.57%	3.519×10^{-8} ± 2.97%	2.153×10^{-10} ± 4.10%

Table 3: Estimates of steady-state unavailability in a machine repairman model (with failure propagation).

Failure Data Set	Erlang 2 (CV = 0.707)	Weibull (CV = 0.805)	Exponential (CV = 1.0)	Hyperexponential 2 (CV = 2.0)
Set I	4.046×10^{-10} $\pm 7.41\%$	4.012×10^{-10} $\pm 6.32\%$	3.953×10^{-10} $\pm 5.68\%$	3.810×10^{-10} $\pm 8.75\%$
Set II	3.861×10^{-8} $\pm 6.15\%$	3.911×10^{-8} $\pm 6.07\%$	4.107×10^{-8} $\pm 5.82\%$	3.904×10^{-8} $\pm 9.90\%$

Table 4: Estimates of steady-state unavailability in a large example (without failure propagation).

Failure Data Set	Erlang 2 (CV = 0.707)	Weibull (CV = 0.805)	Exponential (CV = 1.0)	Hyperexponential 2 (CV = 2.0)
Set I	6.856×10^{-10} $\pm 9.87\%$	6.863×10^{-10} $\pm 8.93\%$	6.646×10^{-10} $\pm 8.83\%$	7.383×10^{-10} $\pm 13.77\%$
Set II	6.610×10^{-8} $\pm 9.18\%$	6.570×10^{-8} $\pm 9.53\%$	6.861×10^{-8} $\pm 8.17\%$	7.766×10^{-8} $\pm 12.53\%$

Table 5: Estimates of steady-state unavailability in a large example (with failure propagation).