

FAST SIMULATION METHODS FOR HIGHLY DEPENDABLE SYSTEMS

Marvin K. Nakayama

Department of Computer and Information Science
New Jersey Institute of Technology
Newark, NJ 07102, U.S.A.

ABSTRACT

We review fast simulation techniques for estimating various performance measures and their derivatives of highly dependable systems. We consider methods for steady-state and transient performance measures of Markovian and non-Markovian systems.

1 INTRODUCTION

Many systems arising in technological fields need to be highly dependable. For example, consider communications systems and transaction processing systems. Before actually constructing such a system, a designer usually will build a mathematical model to determine if the system will have acceptable performance. Performance measures of interest for these types of systems include the steady-state availability (which is the long-run fraction of time that the system is available), the mean time to failure, or the system reliability (which is the probability that the system does not fail before some fixed time).

Because of the size and complexity of most real-world systems, analytic methods typically are not applicable for analyzing the models. Thus, the designer must resort to simulation. However, standard simulation (i.e., without the use of variance reduction techniques) is not efficient for these problems because of the rareness of system failures, and so some variance reduction technique must be employed.

One such method which has proven useful is importance sampling; e.g., see Hammersley and Handscorn (1964) or Glynn and Iglehart (1989). The basic idea of the technique is to change the underlying dynamics (i.e., the probability distributions) of the system so as to cause the "important" events (in our case, system failures) to occur more frequently. This is called a "change of measure." To recover unbiased estimates, we multiply the estimates by a correction factor known as the likelihood ratio.

The main focus of the research in importance sampling is in determining an appropriate change of measure. This is somewhat problem-specific, and if done properly, then variance reductions of orders of magnitude can result. On the other hand, the variance may increase (or even become infinite) if one is not careful. In the context of simulating highly dependable systems, the approach for importance sampling is known as "failure biasing," which we review in this paper. Importance sampling has also been applied to the simulation of other types of systems, most notably in the estimation of probabilities of rare events arising in queueing models. These importance sampling methods typically are based on the theory of large deviations; see Heidelberger (1993) for a review.

The rest of the paper is organized as follows. In Section 2 we present our model of a highly dependable system. Section 3 reviews the basic ideas of importance sampling and why it is needed in our setting. We review the work on fast simulation methods for Markovian dependability models in Section 4 and do the same for non-Markovian systems in Section 5. We present importance sampling schemes for estimating steady-state and transient performance measures. Finally, it should be mentioned that in this paper we do not attempt to list all of the references for each technique presented; for a more complete list, see Nicola, Shahabuddin, and Heidelberger (1993).

2 MODEL OF HIGHLY DEPENDABLE SYSTEMS

The types of system we consider are a class of generalized machine repairman models which can be described by the System Availability Estimator (SAVE) package developed at IBM Research (see Goyal and Lavenberg 1987). We assume the system consists of a finite number of components (e.g., processors or disks) and some number of repairmen divided into classes. As time evolves, the components randomly

fail, where the time to failure of the different components may follow different distributions. After a component fails, it may be replaced by a spare, if available. Each repairman class fixes certain failed components according to some repair strategy. We assume that a repairman is always busy if there are any components failed.

A component can fail in several different failure modes, where each mode has associated with it a repairman class, a repair time distribution, and a (possibly empty) set of affected components. For example, a processor may have two failure modes, corresponding to hardware and software failures. When a component fails, all of the components in its list of affected components also fail with some probability. This is known as "failure propagation." For example, if a processor fails, it may then contaminate some data, thereby causing a disk to fail.

We allow for operational and repair dependencies between the components. Operational dependencies specify how one component's operating behavior depends on the states of other components in the system. For example, if there is a processor which has a power supply that is currently failed, then the processor cannot operate. However, the processor itself is not considered to be failed but is rather in the "dormant" state, which correspondingly has a (possibly) different failure time distribution than the operational state. Similarly, a repairman may not be able to repair a component unless some other components are in certain states. In our previous example, if the power supply and processor are both currently failed, then the repairman may not be able to fix the processor until its power supply is operational.

A component may be in any one of a number of states (operational, failed, spare, or dormant), which defines its current operating and repair characteristics (e.g., failure rate or repair rate). The states of the individual components, along with any relevant information on the queueing of failed components, define the state of the system. We decompose the state space S of the stochastic process into two sets: U , the set of system configurations (i.e., states) in which the system is considered available, and F , the set of states in which the system is considered to be failed. For example, in a computer system consisting of a number of processors and disks, the system must have at least one processor and disk operational to be available. Also, we let 0 denote the state in which all components are operational, and we assume that the system always starts in this state.

We consider systems which achieve high dependability through (small) redundancies of highly reliable components. By a "highly reliable component,"

we mean one with an expected lifetime which is orders of magnitude smaller than its expected repair time. In the Markovian case (in which all failure and repair times are exponentially distributed), this means that the failure rate λ_i of each component i is much smaller than its repair rate μ_i . Shahabuddin (1994) modeled this mathematically by introducing a "rarity parameter" ϵ and letting

$$\lambda_i = a_i \epsilon^{b_i}, \quad (1)$$

where $a_i > 0$ and $b_i \geq 1$ are independent of ϵ and ϵ is small. We assume that the repair rate μ_i of each component i is independent of ϵ . By letting the b_i be different for different i , we can model systems in which certain components are much more reliable than others. For non-Markovian systems (i.e., systems with some failure or repair times which are non-exponential), Heidelberger, Shahabuddin, and Nicola (1994) proposed the natural generalization of (1) by assuming that for each component i , the hazard rate $h_i(s)$ of its failure time distribution is small (i.e., a function of ϵ) and that of its repair distribution is independent of the rarity parameter. More precisely, recall that $h_i(s) = f_i(s)/(1 - F_i(s))$, where f_i is the density function of the failure time distribution F_i of component i . Then we assume that

$$h_i(s) \leq a_i \epsilon^{b_i}, \quad (2)$$

for some constants $a_i > 0$ and $b_i \geq 1$. Many distributions satisfy (2); e.g., hyperexponential, Erlang, Weibull with increasing failure rate over a finite interval, and more general Markovian phase-type distributions.

3 IMPORTANCE SAMPLING

In this section we review the basic ideas of importance sampling and why we need to use it in our setting. For the moment, consider a real-valued random variable X which has probability density function f . Suppose our goal is to estimate $\gamma = E_f[1\{X \in A\}] = \int_A f(x)dx$, where $1\{\cdot\}$ is the indicator function, A is some set of real numbers, and the subscript f on the expectation operator denotes that the expectation is computed using the density f . Note that γ is the probability that X lies in the set A . To estimate γ using standard simulation, we collect i.i.d. samples X_1, X_2, \dots, X_n , where each X_k is generated using the density f . Then, letting $I_k = 1\{X_k \in A\}$, our estimator of γ is $\hat{\gamma}(n) = \frac{1}{n} \sum_{k=1}^n I_k$.

Let us now see why the above method of standard simulation is not efficient in our context of highly dependable systems (or, more generally, rare event

simulation). Suppose the event $\{X \in A\}$ is rare under the density f ; i.e., γ is close to zero. For example, if the random variable X denotes the time the system fails and the set A is the interval $(0, t)$, then γ is the unreliability for time t , the probability that the system fails before time t , which should be small. Our estimator $\gamma(n)$ is unbiased and has variance $\text{Var}_f[1\{X \in A\}]/n = \gamma(1 - \gamma)/n$, where the subscript f denotes the variance is computed using density f . Thus, an approximate $100(1 - \eta)\%$ confidence interval for γ is $\hat{\gamma}(n) \pm z_\eta \sqrt{\gamma(1 - \gamma)/n}$, where z_η is the $1 - \eta/2$ quantile point of the standard normal distribution. In many instances, we want a sample size n which will result in our estimator having a pre-specified relative precision δ , for example, 10%. To achieve this, we need $n \approx z_\eta^2 \text{Var}_f[1\{X \in A\}]/(\delta\gamma)^2 = z_\eta^2(1 - \gamma)/(\delta^2\gamma)$, which grows as γ decreases. In particular, if we define the "relative error" of our estimator to be the standard deviation of $1\{X \in A\}$ over its mean (i.e., $\sqrt{\gamma(1 - \gamma)}/\gamma$), then for a fixed relative precision and confidence level, the sample size grows as the square of the relative error, which diverges to infinity as $\gamma \rightarrow 0$. In this case, we say that our estimator has "unbounded relative error." Practically speaking, as the event of interest becomes rarer (i.e., as $\gamma \rightarrow 0$), it gets more difficult to estimate its probability using standard simulation.

Now consider another density function g for which $g(x) > 0$ whenever $f(x) > 0$. Then

$$\begin{aligned} \gamma &= \int_A f(x) dx = \int_A \frac{f(x)}{g(x)} g(x) dx \\ &= \int_A L(x) g(x) dx = E_g[1\{X \in A\}L], \quad (3) \end{aligned}$$

where $L(x) \equiv f(x)/g(x)$ is known as the likelihood ratio (or Radon-Nykodym derivative). Note that the expectation on the right side of (3) is computed using the density g . This is known as a "change of measure." We form a new estimator of γ based on (3) by collecting n i.i.d. pairs of samples $(X_1, L_1), (X_2, L_2), \dots, (X_n, L_n)$, where each X_k is generated using density g and $L_k = f(X_k)/g(X_k)$. Letting $I_k = 1\{X_k \in A\}$, we obtain the new estimator $\tilde{\gamma}(n) = \frac{1}{n} \sum_{k=1}^n I_k L_k$, which is unbiased. This is how importance sampling is implemented.

Now we observe that $\text{Var}_g[1\{X \in A\}L] = E_g[(1\{X \in A\}L)^2] - \gamma^2 = E_f[1\{X \in A\}L] - \gamma^2$. Hence, to obtain a variance reduction over standard simulation, we need $E_f[1\{X \in A\}L] < E_f[1\{X \in A\}]$. A sufficient condition to ensure this is that for all $x \in A$ with $f(x) > 0$, the likelihood ratio $L(x) < 1$, or equivalently, $g(x) > f(x)$. Thus, we want to choose the new density g so that it increases the probability of the important events (i.e., $x \in A$).

Now suppose that we can choose g such that $E_f[1\{X \in A\}L] \sim \gamma$, where we use the notation that $a \sim b$ if $a/b \rightarrow c$ as $b \rightarrow 0$, where $c < \infty$ is some constant. Then, $\text{Var}_g[1\{X \in A\}L] \sim \gamma^2$ and the relative error of our new estimator is $\sqrt{\text{Var}_g[1\{X \in A\}L]}/\gamma \sim \gamma/\gamma \sim 1$, which remains bounded as $\gamma \rightarrow 0$. We then say that the estimator has "bounded relative error." From a practical standpoint, this means that we only need a fixed number of samples to obtain an estimator with a fixed relative precision and confidence level, independent of the how small γ is.

Although we have thus far only considered applying importance sampling to a single random variable which has a density, the technique can be applied much more generally to stochastic processes. In particular, many importance sampling schemes have been developed to simulate highly dependable systems. These typically fall into the class of "failure biasing" schemes. The basic idea behind these methods is to first fix $p_* > 0$, which is known as the "biasing parameter." (Experimental results of Goyal et al. 1992 suggest that we should select $0.25 \leq p_* \leq 0.9$.) Then from any state in which both component failures and repairs are possible, we increase the total probability that some component fails to p_* (and correspondingly decrease the total probability that some component completes its repair to $1 - p_*$). This makes the system fail more frequently. The various failure biasing schemes differ in the way they increase the probabilities of the individual component failures.

4 MARKOVIAN SYSTEMS

If we assume all component lifetimes and repair times are exponentially distributed, the resulting stochastic process $X = \{X_t : t \geq 0\}$ representing the evolution of the system is a continuous-time Markov chain (CTMC) on some state space S . We denote its generator matrix by $Q = \{q(x, y) : x, y \in S\}$, and define $q(x) = \sum_{z \neq x} q(x, z)$. Let $Y = \{Y_n : n \geq 0\}$ be the embedded discrete-time Markov chain (DTMC) of X , and let $P = \{P(x, y) : x, y \in S\}$ be its transition matrix. Recall that $P(x, y) = q(x, y)/q(x)$ if $y \neq x$, and $P(x, x) = 0$. We simulate the CTMC by first generating a transition of the embedded DTMC, and then given that the system is currently in some state Y_n , we generate the holding time in that state as an exponential random variable with (conditional) mean $1/q(Y_n)$.

We assume that each transition (x, y) with $P(x, y) > 0$ corresponds to either a failure transition (i.e., some components failing) or a repair transition (i.e., some components completing repair). For each state $x \in S$, let $\Lambda_F(x)$ denote the set of states y

for which (x, y) is a failure transition, and let $\Lambda_R(x)$ denote the set of states y for which (x, y) is a repair transition. Now consider any state x for which $\Lambda_F(x) \neq \emptyset$ and $\Lambda_R(x) \neq \emptyset$. If $(x, y) \in \Lambda_F(x)$, then by (1), $q(x, y) \sim \epsilon^{b(x,y)}$ for some $b(x, y) \geq 1$. Also, if $(x, y) \in \Lambda_R(x)$, then $q(x, y)$ is independent of ϵ . Thus, since both failure and repair transitions are possible from state x , $q(x) \sim \epsilon^0$ (i.e., its order of magnitude is independent of ϵ). This implies that if $(x, y) \in \Lambda_F(x)$, $P(x, y) \sim \epsilon^{b(x,y)}$. Also, $P(x, y) \sim \epsilon^0$ if $(x, y) \in \Lambda_R(x)$. Hence, from state x , the DTMC is very unlikely to take a failure transition when ϵ is small. Furthermore, we assume that for any failure transition $(0, y)$ with $y \in F$, $P(0, y) \sim \epsilon^{b(0,y)}$, where $b(0, y) \geq 1$. Therefore, the system fails infrequently.

4.1 Steady-State Performance Measures

We first study the estimation of steady-state performance measures by considering the mean time to failure (MTTF). (Although the MTTF is typically thought of as a transient measure since it depends on the initial state of the system, for our purposes, we will consider it as a steady-state measure since we only need to simulate the embedded DTMC to estimate it.) Let $\tau_A = \inf\{n > 0 : Y_n \in A\}$ for some set of states A , which is the hitting time of the DTMC to the set A , and let $\alpha_A = \sum_{k=0}^{\tau_A-1} 1/q(Y_k)$. Then the MTTF is $E_P[\alpha_F]$, where E_P denotes the expectation operator under P . It can be shown (e.g., see Goyal et al. 1992) that the MTTF satisfies $E_P[\alpha_F] = E_P[\alpha_0 \wedge \alpha_F] / E_P[1\{\alpha_F < \alpha_0\}]$, where $a \wedge b = \min(a, b)$. It turns out that we can efficiently estimate the numerator $E_P[\alpha_0 \wedge \alpha_F]$ using standard simulation. Hence, we now focus our attention on estimating the denominator $\gamma \equiv E_P[1\{\alpha_F < \alpha_0\}]$. If we use standard simulation to estimate this quantity, then generate n i.i.d. samples of the sample path (Y_0, \dots, Y_τ) using the transition matrix P , where $Y_0 = 0$ and $\tau = \tau_0 \wedge \tau_F$. If $\tau_F < \tau_0$ on the k -th sample, then we generate a statistic $I_k = 1$; otherwise, $I_k = 0$. Our standard simulation estimator of γ is $\frac{1}{n} \sum_{k=1}^n I_k$. To study the efficiency of this approach, we need an asymptotic expression for γ in terms of the rarity parameter ϵ . Shahabuddin (1994) showed that under certain conditions,

$$\gamma \sim \epsilon^r \tag{4}$$

for some constant $r \geq 1$, which depends on the system. Also, $\text{Var}_P[1\{\alpha_F < \alpha_0\}] = \gamma - \gamma^2 \approx \gamma \sim \epsilon^r$, and so the relative error of our standard simulation estimator of γ is $\sqrt{\text{Var}_P[1\{\alpha_F < \alpha_0\}]} / \gamma \sim \epsilon^{-r/2} \rightarrow \infty$ as $\epsilon \rightarrow 0$. Thus, the standard simulation estimator of γ has unbounded relative error.

Now we examine how to use importance sampling to estimate γ . First, define a new transition matrix \hat{P} for our DTMC for which $\hat{P}(x, y) > 0$ whenever $P(x, y) > 0$. Then, using \hat{P} , we generate n i.i.d. samples of the sample path $\omega = (Y_0, \dots, Y_\tau)$, which we use to evaluate the likelihood ratio

$$L(\omega) = \prod_{j=0}^{\tau-1} \frac{P(Y_j, Y_{j+1})}{\hat{P}(Y_j, Y_{j+1})}. \tag{5}$$

This yields $(I_1, L_1), \dots, (I_n, L_n)$, and our estimator is $\frac{1}{n} \sum_{k=1}^n I_k L_k$. (Note that we are using standard simulation to estimate the numerator of our ratio expression for the MTTF, and importance sampling to estimate the denominator. This is known as “measure specific importance sampling”; see Goyal et al. 1992.)

We now describe some different approaches for constructing the new transition matrix \hat{P} . The first method is known as simple failure biasing, which was developed by Lewis and Böhm (1984). Since this is a failure biasing method, from any operational state from which there are both failure and repair transitions possible, we increase the total probability of a failure transition to p_* (and correspondingly decrease the total probability of a repair transition to $1 - p_*$). We allocate the p_* to the individual failure transitions in proportion to their original transition probabilities and do similarly for the repair transitions. More precisely, for any state x , let $p_F(x) \equiv \sum_{y \in \Lambda_F(x)} P(x, y)$ be the total probability of taking a failure transition from x , and let $p_R(x) \equiv \sum_{y \in \Lambda_R(x)} P(x, y)$ be the total probability of taking a repair transition from x . If $x \in U$ and $p_F(x) > 0$ and $p_R(x) > 0$, then we let

$$\hat{P}(x, y) = \begin{cases} p_* P(x, y) / p_F(x) & \text{if } (x, y) \in \Lambda_F(x), \\ (1 - p_*) P(x, y) / p_R(x) & \text{if } (x, y) \in \Lambda_R(x), \\ 0 & \text{otherwise.} \end{cases}$$

For any other state $x \in S$, let $\hat{P}(x, y) = P(x, y)$ for all $y \in S$. In particular, $\hat{P}(0, y) = P(0, y)$ for all $y \in S$ since there are no repairs taking place in state 0.

Now we examine the efficacy of simple failure biasing on the following example.

Example 1. Consider a system which consists of two different types of components. The first type of component has a redundancy of two and failure rate $\lambda_1 = \epsilon^2$. The second type has a redundancy of one and failure rate $\lambda_2 = \epsilon^3$. There is a single repairman who fixes all failed components at rate 1 using a processor sharing discipline. The system is available if and only if at least one component of each type is operational. Consider the sample path ω_0 of the DTMC which starts out in state 0 and has a single failure transition in which the component of type 2 fails.

This sample path satisfies $\tau_F < \tau_0$ and has probability $\epsilon^3/(2\epsilon^2 + \epsilon^3) \sim \epsilon$ under the original measure P . All other paths for which $\tau_F < \tau_0$ have probability $o(\epsilon)$ under P . We can show that the order of magnitude of γ is determined by the most likely path to failure (ω_0), so $\gamma \sim \epsilon$ and $r = 1$. The probability of the path ω_0 does not change under simple failure biasing, which implies that $L(\omega_0) = 1$. Consequently, $E_{\hat{P}}[(1\{\alpha_F < \alpha_0\}L)^2] = E_P[1\{\alpha_F < \alpha_0\}L] \geq L(\omega_0)P(\omega_0) \sim \epsilon$, so the variance of our estimator is at least of order ϵ . The relative error of the simple failure biasing estimator is then at least of order $\epsilon^{-1/2}$, which diverges as $\epsilon \rightarrow 0$. Thus, as noted by Shahabuddin (1994), simple failure biasing may not always result in bounded relative error. ■

We now describe another importance sampling scheme known as balanced failure biasing, which was developed by Goyal et al. (1992) and Shahabuddin (1994). We construct the transition matrix \hat{P} for this method as follows. As in simple failure biasing, we increase the total probability of a failure transition to p_* and decrease the total probability of a repair transition to $1 - p_*$. However, we now allocate the p_* equally to each of the individual failure transitions. More precisely, for any transition (x, y) , define $I(x, y) = 1$ if $P(x, y) > 0$, and $I(x, y) = 0$ if $P(x, y) = 0$. For any state x , define $n_F(x)$ to be the number of failure transitions possible (under P) from x . If $x \in U$ and $p_R(x) > 0$, then we let

$$\hat{P}(x, y) = \begin{cases} p_* I(x, y)/n_F(x) & \text{if } (x, y) \in \Lambda_F(x), \\ (1 - p_*)P(x, y)/p_R(x) & \text{if } (x, y) \in \Lambda_R(x), \\ 0 & \text{otherwise.} \end{cases}$$

If $x \in U$ and $p_R(x) = 0$, let $P(x, y) = I(x, y)/n_F(x)$ if $(x, y) \in \Lambda_F(x)$, and $P(x, y) = 0$ otherwise. If $x \in F$, let $\hat{P}(x, y) = P(x, y)$ for all $y \in S$.

Shahabuddin (1994) proved that the estimator of γ obtained using balanced failure biasing will have variance $\text{Var}_{\hat{P}}[1\{\alpha_F < \alpha_0\}L] \sim \epsilon^{2r}$, where $\text{Var}_{\hat{P}}$ is the variance operator under \hat{P} and r is as defined in (4). Thus, the estimator will have relative error $\sqrt{\text{Var}_{\hat{P}}[1\{\alpha_F < \alpha_0\}L]}/\gamma \sim \epsilon^0$, which remains bounded as $\epsilon \rightarrow 0$. Shahabuddin also showed that we can estimate the steady-state unavailability with bounded relative error using balanced failure biasing. (When estimating this performance measure, we should use balanced failure biasing until a failed state is hit; from this point, we revert back to using the original transition matrix P until the system returns back to state 0, and then we repeat the process. This is known as “dynamic importance sampling.”) Experimental results of Goyal et al. (1992) show that

balanced failure biasing works well in practice, reducing the variance by orders of magnitude.

We now want to gain more insight into why balanced failure biasing always leads to bounded relative error. First, let Δ be the set of sample paths of the embedded DTMC for which $\tau_F < \tau_0$, and let Δ_m be the set of paths in Δ which have probability of the order ϵ^m under P , where $m \geq r$ and r is defined in (4). Now consider any path $\omega \in \Delta_m$. For any transition (x, y) on the path ω , $\hat{P}(x, y) \sim \epsilon^0$, so the probability of the entire path ω is of order ϵ^0 under balanced failure biasing. Thus, $L(\omega) \sim \epsilon^m$ by (5). It then can be shown that $E_P[1\{\alpha_F < \alpha_0\}L] = \sum_{m=r}^{\infty} \sum_{\omega \in \Delta_m} L(\omega)P(\omega) \sim \sum_{m=r}^{\infty} \sum_{\omega \in \Delta_m} \epsilon^m \epsilon^m \sim \epsilon^{2r}$, so $\text{Var}_{\hat{P}}[1\{\alpha_F < \alpha_0\}L] \sim \epsilon^{2r}$. Hence, the key to why balanced failure biasing works is the following. First, note that any path in Δ must have probability of order no greater than ϵ^r . Also, recall that for importance sampling to work well, the new probability measure should increase the probability of each of the important events, which, in our case, are the sample paths in Δ . Balanced failure biasing does this by making *all* sample paths (regardless if they are in Δ or not) have probability of order ϵ^0 .

As we previously saw in Example 1, the most likely paths to failure plays a key role in the study of highly dependable systems and importance sampling schemes for them. By “most likely paths to failure” we mean the set of paths in Δ_r . Nakayama (1991) formalized the notion of most likely paths to failure by establishing a conditional limit theorem for the sample paths of the DTMC, given that $\tau_F < \tau_0$.

However, the most likely paths to failure do not tell the whole story, as we cannot ignore the other paths in Δ . Specifically, consider any valid change of measure \hat{P} ; i.e., $\hat{P}(\omega) > 0$ whenever $P(\omega) > 0$. Nakayama (1993) established that a necessary condition for \hat{P} to have the bounded relative error property is that for each $\omega \in \Delta_m$ and $m \geq r$, $\hat{P}(\omega) \sim \epsilon^{d(\omega)}$, where $d(\omega) \leq 2m - 2r$. This basically ensures that all sample paths in Δ are sampled enough under \hat{P} . Moreover, if we assume more structure on \hat{P} (viz., it is a failure biasing scheme), then the condition is only needed for sample paths $\omega \in \Delta_m$ with $r \leq m \leq 2r - 1$. Also, this condition is then both necessary and sufficient. For example, if $r = 2$ (and so $\gamma \sim \epsilon^2$), a necessary and sufficient condition for \hat{P} corresponding to some failure biasing method to give rise to an estimator of γ having bounded relative error is that $\hat{P}(\omega) \sim \epsilon^{d(\omega)}$ with $d(\omega) = 0$ for each $\omega \in \Delta_2$ and $d(\omega) \leq 2$ for each $\omega \in \Delta_3$. Although these results may be difficult to apply in practice, they demonstrate that secondary paths to failure (i.e., paths in Δ_m with $m > r$) also play a role in determining the

effectiveness of an importance sampling methodology.

Nakayama (1991) showed that balanced failure biasing also is effective in reducing the variance of estimates of derivatives of γ with respect to failure rates of components when using the likelihood ratio derivative method (see Glynn 1990). In particular, if balanced failure biasing is employed, the derivative estimates will have bounded relative error. For empirical results, see Nakayama, Goyal, and Glynn (1994).

There are other importance sampling methods available for simulating the embedded DTMC of Markovian dependability systems. These include bias2 failure biasing (Goyal et al. 1992) and failure distance biasing (Carrasco 1991). Nakayama (1993) showed that neither of these is guaranteed to yield estimators of γ having bounded relative error. However, for a given system, one of these methods (or simple failure biasing) may yield estimators having (slightly) smaller variance than balanced failure biasing; see Nakayama (1994). Also, see Juneja and Shahabuddin (1992) and Strickland (1993) for other importance sampling schemes.

4.2 Transient Performance Measures

Now we consider estimating transient performance measures such as the unreliability or the interval unavailability, which is the expected fraction of time that the system is not available during the interval $(0, t)$. For estimating these quantities, we no longer can consider only the embedded DTMC and must take into account the random holding times in each of the states.

Recall that the system starts out in state 0, the state with no failed components. Given that the system is in state 0, the holding time in that state is exponentially distributed with rate $q(0)$. However, since no components are failed in state 0, there are no repair transitions possible from this state, so $q(0) \sim e^{b(0)t}$, where $b(0) \geq 1$. Thus, the probability that the system makes at least one transition before the time horizon t expires is $1 - e^{-q(0)t} \approx q(0)t \rightarrow 0$ as $\epsilon \rightarrow 0$. Hence, even if we use failure biasing to move the embedded DTMC towards a failed state, the probability that the CTMC fails before time t is still small. This leads us to use another change of measure known as "forcing" for generating the holding times in certain states; see Lewis and Böhm (1984). In this method, we sample the initial holding time from the density $f_0(s) = q(0)e^{-q(0)s}/(1 - e^{-q(0)t})$, which is the (conditional) exponential distribution with rate $q(0)$, given that it is less than t . (The distribution of the holding times of subsequent visits to state 0 similarly can be changed.)

Shahabuddin (1993) proved that combining forcing and balanced failure biasing leads to estimates of the interval unavailability which have bounded relative error. Shahabuddin and Nakayama (1993) established similar results for the unreliability and its derivatives with respect to component failure rates. Empirical work in Goyal et al. (1992) also shows that this way of importance sampling leads to large reductions in variance over standard simulation and balanced failure biasing alone when the time horizon t is fairly small; i.e., when $t < 1/q(0)$. However, when t is large, then the method does not work well. Thus, Shahabuddin (1993) and Shahabuddin and Nakayama (1993) establish upper and lower bounds, $U(t)$ and $L(t)$, for the interval unavailability and unreliability and its derivatives which apply for all time horizons. Also, they model the large time horizon scenario by letting $t = \epsilon^{-b_0}$, where $b_0 > 0$, and show that in this case, the upper and lower bounds are tight (i.e., $L(t)/U(t) \rightarrow 1$) as $\epsilon \rightarrow 0$.

5 NON-MARKOVIAN SYSTEMS

In this section, we no longer assume that all of the component lifetimes and repair times are exponentially distributed. Thus, the resulting stochastic process is not a CTMC. It turns out that we can apply a generalization of balanced failure biasing to simulate these systems efficiently. The basic idea is to use "thinning" (Lewis and Shedler 1979).

First, we consider the estimation of the system unreliability for time horizon t . To simplify the notation, we will assume that each component only can be in one of two states, operational or failed, and that there is no failure propagation. (These assumptions can be eliminated.) We assume that the system consists of N components, where component i has failure time distribution F_i and repair time distribution G_i . Let $h_i(\cdot)$ and $r_i(\cdot)$ be the hazard rates of distributions $F_i(\cdot)$ and $G_i(\cdot)$, respectively. After a component is repaired, it is considered to be new. Let $O(s)$ denote the set of components at time s which are operational. For each component $i \in O(s)$, let $A_i(s)$ denote the "age" of the component; i.e., the amount of time which has elapsed since it was last new. Let $\lambda_i(s) \equiv h_i(A_i(s))$ be the failure rate of component i at time s if $i \in O(s)$, and 0 otherwise. Also, let $R(s)$ be the set of components at time s which are currently undergoing repair. For a component $i \in R(s)$, let $B_i(s)$ be the "age" of the repair process of component i at time s ; i.e., the amount of time which has elapsed since the repair began on the component. Let $\mu_i(s) \equiv r_i(B_i(s))$ be the repair rate of component i at time s if $i \in R(s)$, and 0 otherwise. Recall

that the failure rates of the components are assumed to be small; i.e., (2) is in force for $0 \leq s \leq t$. Let $\lambda_F(s) \equiv \sum_{i \in O(s)} \lambda_i(s)$ and $\mu_R(s) \equiv \sum_{i \in R(s)} \mu_i(s)$ be the total failure and repair rates, respectively, at time s . The total event rate at time s is then $\lambda(s) \equiv \lambda_F(s) + \mu_R(s)$.

One approach for simulating the system without using importance sampling is as follows (see Nicola, Heidelberger, and Shahabuddin 1992). First, assume that $\lambda(s) \leq \beta$ for all $0 \leq s \leq t$ w.p. 1, where $\beta < \infty$ is some constant. Note that this implicitly assumes that the hazard rates of each of the failure and repair distributions are bounded on the interval $(0, t)$. Many distributions satisfy this condition (e.g., hyperexponential, Erlang, and Weibull with increasing failure rate), but others do not (e.g., constant and uniform). (We will discuss another technique later which relaxes this assumption to some degree.) Then, we can use a uniformization-based approach to simulate the system. More specifically, let $N_\beta = \{N_\beta(s) : s \geq 0\}$ be a Poisson process with rate β . We will simulate our dependability system by simulating the process N_β and categorizing each event of N_β as one of three types of events for our dependability system: failure event, repair event, or pseudo event. If an event of N_β occurs at time s , it is a failure event with probability $\lambda_F(s)/\beta$, a repair event with probability $\mu_R(s)/\beta$, and a pseudo event with probability $1 - \lambda(s)/\beta$. Given that an event at time s is a failure event, we choose component $i \in O(s)$ to be the one which failed with probability $\lambda_i(s)/\lambda_F(s)$. Similarly, given that an event at time s is a repair event, we choose component $i \in R(s)$ to be the one which completed repair with probability $\mu_i(s)/\mu_R(s)$. Pseudo events have no effect on the dependability system.

Nicola, Heidelberger, and Shahabuddin (1992) proposed the following importance sampling scheme to simulate the system. Again simulate the Poisson process N_β , but we now change the probabilities used to determine if a Poisson event corresponds to a failure, repair, or pseudo event of the dependability system. We accomplish this by changing the failure and repair rates of component i at time s to $\tilde{\lambda}_i(s)$ and $\tilde{\mu}_i(s)$, respectively. Thus, the new total failure, repair, and event rates at time s are $\tilde{\lambda}_F(s) \equiv \sum_{i \in O} \tilde{\lambda}_i(s)$, $\tilde{\mu}(s) \equiv \sum_{i \in R} \tilde{\mu}_i(s)$, and $\tilde{\lambda}(s) \equiv \tilde{\lambda}_F(s) + \tilde{\mu}(s)$, respectively. We assume that $\tilde{\lambda}_i$ and $\tilde{\mu}_i$ are chosen so that for all $0 \leq s \leq t$, $\tilde{\lambda}(s) \leq \beta$ w.p. 1 and the change of measure is valid (i.e., $\tilde{\lambda}_i(s) > 0$ whenever $\lambda_i(s) > 0$, $\tilde{\mu}_i(s) > 0$ whenever $\mu_i(s) > 0$, and $\beta - \tilde{\lambda}(s) > 0$ whenever $\beta - \lambda(s) > 0$). Under importance sampling, if an event of N_β occurs at time s , it is a failure of component i with probability $\tilde{\lambda}_i(s)/\beta$, a repair of compo-

nent i with probability $\tilde{\mu}_i(s)/\beta$, and a pseudo event with probability $1 - \tilde{\lambda}(s)/\beta$.

Now we show how to choose $\tilde{\lambda}_i$ and $\tilde{\mu}_i$ to approximate balanced failure biasing. Suppose that an event of N_β occurred at time s . Under the new thinning probabilities, it is not a pseudo event with probability $\tilde{\lambda}(s)/\beta$. Suppose that $\lambda_F(s) > 0$ and $\mu_R(s) > 0$ so the event can be either a failure or a repair in the original system, and we now have to decide which it is in the new system. Given that the event is not a pseudo event, to approximate failure biasing, we increase the probability that it is a failure event to p_* and correspondingly decrease the probability of a repair event to $1 - p_*$ by choosing $\tilde{\lambda}_F(s)$ and $\tilde{\mu}(s)$ such that $\tilde{\lambda}_F(s)/\tilde{\lambda}(s) = p_*$ and $\tilde{\mu}(s)/\tilde{\lambda}(s) = 1 - p_*$, where p_* is some constant. If the event is a failure, then we implement balancing by selecting component $i \in O(s)$ to be the one that failed with probability $1/O(s)$. If the event is a repair, then we select component $i \in R(s)$ to be the one to complete its repair with probability $\tilde{\mu}_i(s)/\tilde{\mu}(s)$, where we select $\tilde{\mu}_i(s)$ so that $\tilde{\mu}_i(s)/\tilde{\mu}(s) = \mu_i(s)/\mu_R(s)$.

But, as in the Markovian setting, even though we have now implemented a generalization of balanced failure biasing, the probability that any component fails before the time horizon t expires is still very small. Thus, we need to use some form of forcing. To do this, we choose $\tilde{\lambda}_F(s)$ to be significantly larger than $\lambda_F(s)$ when $\mu_R(s) = 0$, and let $\tilde{\lambda}_i(s) = \tilde{\lambda}_F(s)/N$. Hence, the probability that some component will fail before time t is now much larger. However, we are not guaranteed that some component will fail before time t as in the Markovian case; therefore, we call this "approximate forcing."

Now we examine the likelihood ratio L arising from this importance sampling scheme. We can decompose L into a product of terms: $L = L_F L_R L_P$, where L_F is the likelihood ratio corresponding to the failure events and L_R and L_P are the same for the repair and pseudo events. Let $T_{i,j}$ be the time of the j -th failure of component i , and $N_i(t)$ be the number of times component i fails before time t . Then $L_F = \prod_{i=1}^N \prod_{j=1}^{N_i(t)} \lambda_i(T_{i,j})/\tilde{\lambda}_i(T_{i,j})$. We can similarly express L_R and L_P .

Heidelberger, Shahabuddin, and Nicola (1994) established that if the above importance sampling method is used to estimate the system unreliability for some fixed time horizon t , then the resulting estimator will have bounded relative error. Nicola, Heidelberger, and Shahabuddin (1992) developed another importance sampling scheme called "exponential transformation" which allows repair time distributions (but not the failure time distributions) to have unbounded hazard rates on $(0, t)$. Heidelberger,

Shahabuddin, and Nicola (1994) proved that this approach also has bounded relative error when estimating the unreliability.

REFERENCES

- Carrasco, J. A. 1991. Failure distance-based simulation of repairable fault-tolerant systems. *Proceedings of the Fifth International Conference on Modeling Techniques and Tools for Computer Performance Evaluation*, 337-351.
- Glynn, P. W. 1990. Likelihood ratio derivative estimators for stochastic systems. *Communications of the ACM* 33:75-84.
- Glynn, P. W. and D. L. Iglehart. 1989. Importance sampling for stochastic simulations. *Management Science* 35:1367-1392.
- Goyal, A. and S. S. Lavenberg. 1987. Modeling and analysis of computer system availability. *IBM Journal of Research and Development* 31, 6:651-664.
- Goyal, A., P. Shahabuddin, P. Heidelberger, V. F. Nicola, and P. W. Glynn. 1992. A unified framework for simulation Markovian models of highly reliable systems. *IEEE Transactions on Computers* C-41:36-51.
- Hammersley, J. M. and D. C. Handscomb. 1964. *Monte Carlo Methods*. London: Methuen and Co., Ltd.
- Heidelberger, P. 1993. Fast simulation of rare events in queuing and reliability models. IBM Research Report RC 19028, Yorktown Heights, NY.
- Heidelberger, P., P. Shahabuddin, and V. F. Nicola. 1994. Bounded relative error in estimating transient measures of highly dependable non-Markovian systems. *ACM Transactions on Modeling and Computer Simulation*, 4:137-164.
- 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.
- Lewis, E. E. and F. Böhm. 1984. Monte Carlo simulation of Markov unreliability models. *Nuclear Engineering and Design* 77:49-62.
- Lewis, P. A. W. and G. S. Shedler. 1979. Simulation of nonhomogeneous Poisson processes by thinning. *Naval Research Logistics Quarterly* 26:403-413.
- Nakayama, M. K. 1991. Asymptotics of likelihood ratio derivative estimators in simulations of highly reliable Markovian systems. *Management Science*, to appear.
- Nakayama, M. K. 1993. General conditions for bounded relative error in simulations of highly reliable Markovian systems. IBM Research Report RC 15400, Yorktown Heights, NY.
- Nakayama, M. K. 1994. A characterization of the simple failure biasing method for simulations of highly reliable Markovian systems. *ACM Transactions on Modeling and Computer Simulation*, 4:52-88.
- Nakayama, M. K., A. Goyal, and P. W. Glynn. 1994. Likelihood ratio sensitivity analysis for Markovian models of highly dependable systems. *Operations Research* 42:137-157.
- 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.
- Nicola, V. F., P. Shahabuddin, and P. Heidelberger. 1993. Techniques for fast simulation of highly dependable systems. IBM Research Report RC 18956, Yorktown Heights, NY.
- Shahabuddin, P. 1993. Fast transient simulation of Markovian models of highly dependable systems. *Proceedings of the Performance '93 Conference*.
- Shahabuddin, P. 1994. Importance sampling for the simulation of highly reliable Markovian systems. *Management Science* 40:333-352.
- Shahabuddin, P. and M. K. Nakayama. 1993. Estimation of reliability and its derivatives for large time horizons in Markovian systems. *Proceedings of the 1993 Winter Simulation Conference*, IEEE Press, 422-429.
- Strickland, S. G. 1993. Optimal importance sampling for quick simulation of highly reliable Markovian systems. *Proceedings of the 1993 Winter Simulation Conference*, IEEE Press, 437-444.

AUTHOR'S BIOGRAPHY

MARVIN K. NAKAYAMA is an assistant professor in the Computer and Information Science Department at the New Jersey Institute of Technology. He previously held positions in the Management Science/Computer Information Systems Department of the Graduate School of Management at Rutgers University and at the IBM Thomas J. Watson Research Center in Yorktown Heights, New York. He received a B.A. in Mathematics/Computer Science from University of California, San Diego and an M.S. and Ph.D. in Operations Research from Stanford University. His research interests include simulation output analysis, gradient estimation, and rare event simulation.