Proceedings of the 1993 Winter Simulation Conference
G.W. Evans, M. Mollaghasemi, E.C. Russell, W.E. Biles (eds.)

ESTIMATION OF RELIABILITY AND ITS DERIVATIVES FOR LARGE TIME HORIZONS IN MARKOVIAN SYSTEMS

Perwez Shahabuddin
Marvin K. Nakayama

IBM T. J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598

ABSTRACT

A number of importance sampling methods have been previously proposed for estimating the system unreliability of highly reliable Markovian systems. These techniques are effective when the time horizon of interest is small. However, for large time horizons, these methods are no longer efficient. We describe a technique in which instead of estimating the actual measure, we estimate bounds on the measure. The bounds can be estimated efficiently, and for large time horizons, they are close to the actual measure. Similar techniques for derivative estimation are also presented.
Keywords: Variance reduction, Markov chains, importance sampling, reliability, derivative estimation, regenerative systems, likelihood ratios.

1 INTRODUCTION

Many mission-oriented systems need to be available during some fixed time interval. For example, consider computers used in space missions. For these types of systems, a performance measure of interest is the system unreliability, which is the probability that the system fails before some time horizon. Because of the typically huge state spaces of mathematical models of such systems, analytical methods may not be feasible for solving for performance measures, and so we have to resort to simulation. Using naive simulation (i.e., without the use of any variance reduction techniques) to estimate performance measures of highly reliable systems is inefficient because of the rarity of system failures. Thus, we must apply variance reduction techniques to obtain better estimators. One such method is importance sampling. The basic idea behind this approach is to simulate the system under a probability measure different from the original one, where the new measure is chosen so as to make the important events (in our case, system failures) occur more frequently.

A number of importance sampling methods have been previously proposed for estimating the system unreliability of highly reliable systems; e.g., see Lewis and Böhm (1984), Goyal et al. (1992), Nicola et al. (1991), and Nicola et al. (1992). Empirical results show that these methods are effective when the time horizon is small; i.e., it is less than one over the sum of all of the component failure rates, which is the expected time of the first component failure. However, for large time horizons, the importance sampling schemes do not perform as well and/or are sensitive to the values of the parameters used in the importance sampling. Carrasco (1991) developed a method which works for large time horizons, but the technique involves estimating a Laplace transform and inverting it, making implementation difficult. Thus, the need arises for simple methods to estimate the system unreliability that are efficient in the large time horizon setting.

In this paper, we examine the problem of estimating transient measures for large time horizons for highly reliable Markovian systems through estimating close upper and lower bounds. Preliminary work in this direction was done by Shahabuddin (1993), who investigated the problem of estimating the interval unavailability, which is the expected fraction of time that a system is failed during some fixed time interval. We now extend these ideas to estimate the unreliability. Moreover, we investigate the estimation of partial derivatives of the unreliability with respect to component failure rates.

Working within the mathematical framework developed by Shahabuddin (1990) (see Shahabuddin 1991 for a refined version), we study the asymptotic properties of the estimator of the unreliability obtained using an importance sampling scheme which combines forcing (Lewis and Böhm 1984) and balanced failure biasing (Shahabuddin 1990 and Goyal et al. 1992). We first present a theorem which establishes that for any given (fixed) time horizon, the estimator has bounded relative error. (An estimator is said to have bounded relative error if the expected width of the confidence interval for a fixed number of samples over the quantity to be estimated remains bounded as the component failure rates tend to zero with the repair rates fixed. This notion was introduced in Shahabuddin 1990.) We then give an intuitive explanation as to why the estimator is no longer efficient when the time horizon is large. In general, this is due to the fact that the likelihood ratio (which arises from importance sampling) has a variance which grows (approximately) exponentially with the time horizon; see Glynn (1992).

Because of this, we develop upper and lower bounds for the unreliability which are applicable for any time horizon. The bounds are in terms of expectations of random quantities defined over regenerative cycles. Since regenerative cycles of highly reliable Markovian systems are typically small, we can estimate these expectations, and thus the bounds, efficiently using importance sampling. Moreover, these bounds converge to the unreliability as the failure rates vanish. We present similar bounds for the partial derivatives of the unreliability with respect to the component failure rates and show that these bounds also converge (under certain assumptions) to the partial derivatives. We provide some experimental results showing that the bounds are typically quite accurate in practice.

The rest of the paper is organized as follows. Section 2 contains a brief description of the type of highly reliable Markovian systems we are considering and the importance sampling techniques used for such systems. Previous bounded relative error results in the estimation of steady state measures (using balanced failure biasing) are also reviewed in this section. Bounded relative error results in the estimation of unreliability and its derivatives (using balanced failure biasing and forcing) are presented in Section 3. Section 4 describes unreliability estimation using bounds (for the case of large time horizons), and Section 5 describes the same for derivatives. Experimental results are presented in Section 6. For formal proofs of results in this paper, the reader is referred to Shahabuddin and Nakayama (1993).

2 BACKGROUND

2.1 CTMC MODELS OF RELIABILITY SYSTEMS

We consider Markovian models of reliability systems of the type in Goyal and Lavenberg (1987). These systems consist of a number of different types of components, each having a certain redundancy. All components have exponentially distributed failure and repair times and components of the same type have the same failure rate and the same repair rate. Component interdependencies arise due to the sharing of a limited number of repairpersons, operational and repair dependencies (the operation/repair of a component depends on other components being up) and failure propagation (the failure of a component causes some other components to fail with certain probabilities). The system is considered to be up whenever certain combinations of components are up. In Goyal and Lavenberg (1987) only the preemptive random order service discipline was considered. However, our simulation analysis applies to systems with any preemptive or non-preemptive repair discipline.

Let $\{\mathbf{X}(t) : t \geq 0\}$ be the continuous time Markov chain (CTMC) model of the Markovian reliability system described above, and let $\{\mathbf{Y}_n : n \geq 0\}$ denote its embedded discrete time Markov chain (DTMC). For systems in Goyal and Lavenberg (1987), $\mathbf{X}(t) = (X_1(t), X_2(t), \ldots, X_R(t))$, where $X_i(t)$ is the number of components of type i that are up at time t and R is the total number of component types. For models with the general repair disciplines which we consider, we may have to add an ordered list of components waiting to be repaired at each repairperson.

Let n_i be the total number of type i components in the system. We label the state with all components up as $\mathbf{1}$. (We also use $\mathbf{1}$ to denote the set of states which contains only the single state $\mathbf{1}$.) We will assume that $\mathbf{X}(0) = \mathbf{Y}_0 = \mathbf{1}$. Let S be the set of states that are accessible from $\mathbf{1}$ and we assume that the CTMC is irreducible over the set S. We partition S into two subsets: $S = U \cup F$, where U is the set of up states (i.e., the set of states in which the system is considered to be up) and F is the set of down states. An important property of systems in Goyal and Lavenberg (1987) is that

A: All states in S except state $\mathbf{1}$ have at least one repair transition.

We will be considering system regenerations that occur when the system enters state $\mathbf{1}$. In any regenerative cycle let Z be the random variable denoting the holding time in state $\mathbf{1}$ and W be the sum of the holding times in all other states.

Let Φ denote the probability measure on the sample paths of this CTMC. Let $\mathbf{Q} = (q(\mathbf{x}, \mathbf{y}) : \mathbf{x}, \mathbf{y} \in S)$ be the rate matrix of $\{\mathbf{X}(t) : t \geq 0\}$, and let $\mathbf{P} = (P(\mathbf{x}, \mathbf{y}) : \mathbf{x}, \mathbf{y} \in S)$ be the transition matrix of $\{\mathbf{Y}_n : n \geq 0\}$. One can simulate a CTMC by progressively generating state transitions using $\mathbf{P}$, and generating the random holding times in each state. Let $q(\mathbf{x})$ denote the total rate out of state $\mathbf{x}$. For any set of states $\mathcal{E} \subset S$, let $T_{\mathcal{E}} = \inf\{t > 0 : \mathbf{X}(t-) \notin \mathcal{E}, \mathbf{X}(t) \in \mathcal{E}\}$, and $\tau_{\mathcal{E}} = \inf\{n \geq 1 : \mathbf{Y}_n \in \mathcal{E}\}$. Of particular interest are $T_{\mathbf{1}}$, T_F, $\tau_{\mathbf{1}}$ and τ_F.

We will be interested in estimating the unreliability. Given a finite time horizon t, the unreliability, $U(t)$, is defined to be the probability that the system fails before time t given that it starts in state $\mathbf{1}$; i.e., $U(t) = P(T_F < t)$.

2.2 IMPORTANCE SAMPLING FOR HIGHLY RELIABLE MARKOVIAN SYSTEMS

Consider the problem of estimating $\mu = E_f(X)$, where the subscript indicates that X is sampled from the density $f(\cdot)$. Let $g(\cdot)$ be another density where $g(x) > 0$ whenever $xf(x) > 0$. Then we can write

$$\begin{aligned} \mu &= E_f(X) = \int_{\forall x} xf(x)dx \\ &= \int_{\forall x} xL(x)g(x)dx = E_g(XL(X)), \quad (1) \end{aligned}$$

where $L(x) \equiv f(x)/g(x)$ whenever $g(x) > 0$ and $L(x) \equiv 0$ otherwise. In importance sampling, to estimate μ, instead of sampling X from $f(\cdot)$, we sample $XL(X)$ from $g(\cdot)$. The main problem in importance sampling is to select g so that $E_g(X^2L^2(X)) \ll E_f(X^2)$, thereby substantially reducing the variance over naive simulation.

One change of measure often used for highly reliable Markovian systems is called failure biasing (Lewis and Böhm 1984). The idea behind failure biasing is to make component failure transitions of the embedded DTMC occur with a probability that is much higher than in the original system. In state $\mathbf{1}$ there are no repair transitions. Therefore, we do not need to use failure biasing in this state. However, in states that have both failure and repair transitions, the total probability of repair transitions is ≈ 1 and the total probability of failure transitions is ≈ 0. In such states we increase the total probability of failure transitions to ρ (where ρ is some constant between 0 and 1 that is significantly larger than the failure probabilities; in practice ρ is typically taken to be 0.5) and correspondingly decrease the probability of repair transitions to $1-\rho$. In the version of failure biasing called balanced failure biasing, the probabilities of particular failure transitions, given that a failure transition has occurred, are all made the same (this is also done in state $\mathbf{1}$).

Let $\mathbf{P}' = (P'(\mathbf{x}, \mathbf{y}) : \mathbf{x}, \mathbf{y} \in S))$ denote the transition matrix corresponding to balanced failure biasing. Let $\mathbf{\Phi}'$ be the new probability measure on the sample paths of $\{\mathbf{X}(t) : t \geq 0\}$ in which we use $\mathbf{P}'$ instead of $\mathbf{P}$ until the system fails and $\mathbf{P}$ from then on. Note that for the unreliability, for each sample, we only need to simulate the CTMC until the system fails.

If the time horizon t is orders of magnitude less than $E_{\mathbf{\Phi}}(Z)$, then the system will fail very rarely in $[0, t]$ even though we failure bias. To avoid this, we sample the time of the first event from the distribution of Z conditioned on it being less than t. Since Z is exponentially distributed with rate $q(\mathbf{1})$, the time of the first event which we use in the simulation is sampled from the distribution function given by $F(s) = (1-e^{-q(\mathbf{1})s})/(1-e^{-q(\mathbf{1})t})$, where $0 \leq s \leq t$. This is called forcing (Lewis and Böhm 1984). Let $\mathbf{\Phi}'_F$ be the new probability measure on the sample paths corresponding to both balanced failure biasing and forcing.

2.3 MODELLING SYSTEMS WITH HIGHLY RELIABLE COMPONENTS

In mathematical models of highly reliable Markovian systems, the failure rate of any component type, say component type i, is given by $\lambda_i = \tilde{\lambda}_i \epsilon^{b_i}$ where $\tilde{\lambda}_i$ and b_i are positive constants and ϵ is a small parameter called the rarity parameter. The repair rates are represented by constants $\mu_i > 0$. The aim is to study the variance of the estimator of the given dependability measure, with and without importance sampling, for small ϵ. Let $b_0 = \min\{b_1, b_2, \ldots, b_N\} > 0$. Then $E_{\mathbf{\Phi}}(Z) = 1/q(\mathbf{1}) = 1/\sum_{i=1}^{N} n_i \lambda_i \epsilon^{b_i} = \Omega(\epsilon^{-b_0})$.[1] It was shown in Shahabuddin (1991) that $E_{\mathbf{\Phi}}(W) = \Omega(1)$ and thus both $E_{\mathbf{\Phi}}(T_{\mathbf{1}}) = E_{\mathbf{\Phi}}(Z) + E_{\mathbf{\Phi}}(W)$ and $E_{\mathbf{\Phi}}(\min(T_{\mathbf{1}}, T_F))$ are $\Omega(\epsilon^{-b_0})$.

We will now briefly review the order of magnitude results for the variance associated with importance sampling in the estimation of steady state measures (see Shahabuddin 1991). A crucial quantity to be estimated in the context of steady state measures is $\gamma_\epsilon = P_{\mathbf{\Phi}}\{T_F < T_{\mathbf{1}}\} = E_{\mathbf{\Phi}}(I(T_F < T_{\mathbf{1}}))$. Note that to obtain a sample of $I(T_F < T_{\mathbf{1}})$, we need only simulate the CTMC until time $T_{\min} \equiv \min(T_F, T_{\mathbf{1}})$. Let $\mathbf{\Phi}_s$ be a change of measure on the sample paths of the CTMC in which we use $\mathbf{P}'$ until time $T_{\min}$ and then $\mathbf{P}$ after that. For any stopping time τ of the embedded DTMC, define

$$L_\tau = \prod_{i=0}^{\tau-1} \frac{P(\mathbf{Y}_i, \mathbf{Y}_{i+1})}{P'(\mathbf{Y}_i, \mathbf{Y}_{i+1})}.$$

[1] A function $f(\epsilon)$ is said to be $\Omega(\epsilon^c)$ if there exist constants K_1 and K_2 such that for all sufficiently small ϵ, $K_1\epsilon^c < f(\epsilon) < K_2\epsilon^c$.

Let $\tau_{\min} = \min(\tau_F, \tau_1)$. Then in the same spirit as Equation 1 we can write

$$E_\Phi(I(T_F < T_1)) = E_{\Phi_s}(I(T_F < T_1)L_{\tau_{\min}}).$$

Let $\sigma^2_{\gamma_\epsilon}(\cdot)$ denote the variance associated with the estimation of γ_ϵ using the probability measure within the parentheses. For example, $\sigma^2_{\gamma_\epsilon}(\Phi_s) \equiv Var_{\Phi_s}(I(T_F < T_1)L_{\tau_{\min}}) = E_{\Phi_s}(I(T_F < T_1)L^2_{\tau_{\min}}) - \gamma^2_\epsilon$. Then we have the following theorem:

Theorem 1 (Shahabuddin (1990)) *Both γ_ϵ and $\sigma^2_{\gamma_\epsilon}(\Phi)$ are $\Omega(\epsilon^r)$, where r is a non-negative constant depending on the model. Also, $E_{\Phi_s}(I(T_F < T_1)L^2_{\tau_{\min}})$ is $\Omega(\epsilon^{2r})$, and so $\sigma^2_{\gamma_\epsilon}(\Phi_s)$ is $O(\epsilon^{2r})$.*

Let $RE'_{\gamma_\epsilon}(\cdot)$ denote the relative error (i.e., the expected confidence interval half width divided by the quantity to be estimated) of the estimator of γ_ϵ using the probability measure within the parentheses. Then, we get the following corollary:

Corollary 1 *For a fixed number n of regenerative cycles (and corresponding to a fixed $100(1-\delta)\%$ level of confidence), $RE'_{\gamma_\epsilon}(\Phi) = \Omega(\epsilon^{-r/2})$ and $RE'_{\gamma_\epsilon}(\Phi_s) = O(1)$.*

Proof. Let $z_{\delta/2}$ be the $\delta/2$ percentile point of the standard normal distribution. Then

$$RE'_{\gamma_\epsilon}(\Phi) = \frac{z_{\delta/2}}{\sqrt{n}}\frac{\sigma_{\gamma_\epsilon}(\Phi)}{\gamma_\epsilon} = \frac{z_{\delta/2}}{\sqrt{n}}\frac{\sqrt{O(\epsilon^r)}}{\Omega(\epsilon^r)} = O(\epsilon^{-r/2}).$$

A similar calculation can be used for $RE'_{\gamma_\epsilon}(\Phi_s)$. ∎

The corollary states that the relative error using naive simulation tends to infinity as component failure events get rarer, whereas using balanced failure biasing, it remains bounded. Another way of looking at it is that to achieve a given level of relative accuracy, the amount of computational effort required using naive simulation tends to infinity (as component failure events get rarer), whereas it remains bounded using balanced failure biasing.

3 BOUNDED RELATIVE ERROR IN THE ESTIMATION OF RELIABILITY AND ITS DERIVATIVES

In this section we present bounded relative error results in the estimation of the unreliability. We can express the unreliability as $U_\epsilon(t) = E_\Phi(I(T_F < t))$. Let $\sigma^2_{U_\epsilon(t)}(\cdot)$ be the variance associated with the estimation of the $U_\epsilon(t)$ using the probability measure within the parenthesis.

Theorem 2 *For fixed t, both $U_\epsilon(t)$ and $\sigma^2_{U_\epsilon(t)}(\Phi)$ are $\Omega(\epsilon^{r+b_0})$. Also, $\sigma^2_{U_\epsilon(t)}(\Phi') = O(\epsilon^{2r+b_0})$ and $\sigma^2_{U_\epsilon(t)}(\Phi'_F) = O(\epsilon^{2(r+b_0)})$. Hence, for a fixed number n of replications, $RE'_{U_\epsilon(t)}(\Phi) \to \infty$ and $RE'_{U_\epsilon(t)}(\Phi') \to \infty$ as $\epsilon \to 0$, whereas $RE'_{U_\epsilon(t)}(\Phi'_F) = O(1)$.*

By "fixed t," we mean that t is independent of ϵ. From a modelling point of view, this means that the time horizon is small as compared to the expected failure time of components. In such cases, Theorem 2 implies that we can expect efficient simulation estimates using a combination of balanced failure biasing and forcing (though not through naive simulation or balanced failure biasing alone). Similar results hold for the estimation of the derivatives (using the likelihood ratio method) of the unreliability with respect to the component failure rates.

4 RELIABILITY ESTIMATION FOR LARGE TIME HORIZONS

To study the behavior of the unreliability estimator for large t, we will now consider the case where t is not fixed but also depends on ϵ. In particular, we consider the case when t is $\Omega(\epsilon^{-r_t})$, where $r_t \geq 0$. For $r_t = 0$ (i.e., t is fixed and we will get the results of the last section), t is of the same order as the expected repair times, and for $r_t = b_0$, t is of the same order as the expected first component failure time in the system (which is of the same order as the expected regenerative cycle time). For $r_t = b_0 + r$, t is of the same order as the mean time to system failure (MTTF). (The MTTF for regenerative systems may be expressed as $E_\Phi(\min(T_F, T_1))/\gamma_\epsilon$; hence, the MTTF is $\Omega(\epsilon^{-(b_0+r)})$.)

Theorem 3 *Consider the case of large time horizons where $t = \Omega(\epsilon^{-r_t})$ and $r_t \geq 0$. If $r_t \leq b_0$, then both $U_\epsilon(t)$ and $\sigma^2_{U_\epsilon(t)}(\Phi)$ are $\Omega(\epsilon^{r+b_0-r_t})$ and $\sigma^2_{U_\epsilon(t)}(\Phi'_F) = O(\epsilon^{2(r+b_0-r_t)})$.*

We saw in the previous section that for fixed t, for the bounded relative error property to hold, the order of the variance using importance sampling had to be twice the order as the unreliability. In Theorem 3, as long as $r_t \leq b_0$, the order of the variance is twice the order of the unreliability. Hence, for $r_t \leq b_0$, we can expect the simulation using Φ'_F to be efficient.

Now let us examine what happens when $r_t > r_q$. To do so, let $C_\epsilon = E_{\Phi'_s}(I(T_F < T_1)L^2_{\tau_{\min}})$ and $B_\epsilon = E_{\Phi'_s}(I(T_F > T_1)L^2_{\tau_{\min}})$. Also, let

$$\bar{V}_\epsilon(t) = (1 - e^{-q(1)t})\frac{C_\epsilon}{B_\epsilon - 1}(e^{(B_\epsilon - 1)q(1)t} - 1).$$

It is easy to show that for $0 < r_t \le b_0$, $E_{\Phi'_F}(I(T_F < t)L^2_{\tau_F})/\bar{V}_\epsilon(t) \to 1$ as $\epsilon \to 0$. We conjecture that this is true even for $b_0 < r_t < r + b_0$. If this conjecture holds, then using the fact from Shahabuddin (1991) that $C_\epsilon = \Omega(\epsilon^{2r})$ and $B_\epsilon - 1 = \Omega(1)$, we get that for $r_t > b_0$, $E_{\Phi'_F}(I(T_F < t)L^2_{\tau_F})$ is $\Omega(\epsilon^{2r}e^{\epsilon^{b_0 - r_t}})$ (since $1 - e^{-q(\mathbf{1})t} = \Omega(1)$ and $e^{(B_\epsilon - 1)q(\mathbf{1})t} \gg 1$). Hence for the case of large time horizons where $r_t > b_0$, we should not expect efficient estimates of $U_\epsilon(t)$ by using Φ'_F. This is also what we see in practice; see Goyal et al. (1992).

So what do we do for the case where $r_t > b_0$? In this case it is best to estimate bounds on the unreliability.

Theorem 4 *Let $\bar{U}_\epsilon(t) = 1 - e^{-\gamma_\epsilon q(\mathbf{1})t}$. Then $U_\epsilon(t) \le \bar{U}_\epsilon(t)$ for all ϵ and t. If $t = \Omega(\epsilon^{r_t})$ where $0 < r_t < r + b_0$, then $U_\epsilon(t)/\bar{U}_\epsilon(t) \to 1$ as $\epsilon \to 0$.*

Note that the upper bound is in terms of the regenerative cycle based measure γ_ϵ, which can be efficiently estimated. Hence, the upper bound estimate, which we form by replacing the γ_ϵ term in the expression for $\bar{U}_\epsilon(t)$ by the estimate of γ_ϵ, can also be estimated efficiently. To see this, note that for $r_t < r + b_0$, $\bar{U}_\epsilon(t) = (1 - e^{-q(\mathbf{1})\gamma_\epsilon t}) \approx q(\mathbf{1})t\gamma_\epsilon = \Omega(\epsilon^{r + b_0 - r_t})$. From Theorem 1 we see that $\sigma^2_{\gamma_\epsilon}(\Phi_s) = O(\epsilon^{2r})$ and hence $\sigma^2_{\bar{U}_\epsilon(t)}(\Phi_s) \approx q(\mathbf{1})^2 t^2 \sigma^2_{\gamma_\epsilon}(\Phi_s) = O(\epsilon^{2r + 2b_0 - 2r_t})$. Thus the order of the variance is again twice the order of the upper bound, and so the simulation is efficient. Also note that for small ϵ, the upper bound is close to the actual measure.

It is also possible to obtain a close lower bound on the unreliability. Let $l \equiv l(t,q) = \max(\sqrt{t}, t\sqrt{q})$. Then we have the following theorem:

Theorem 5 *Let*

$$\underline{U}_\epsilon(t) \equiv \bar{U}_\epsilon(t) - \Bigg(e^{-\gamma_\epsilon q(\mathbf{1})(t-l)} - e^{-\gamma_\epsilon q(\mathbf{1})t} + \frac{E_\Phi(W)}{\gamma_\epsilon l}(1 - e^{-\gamma_\epsilon q(\mathbf{1})(t-l)}) - \frac{q(\mathbf{1})(t-l)E_\Phi(WI(T_{\mathbf{1}} < T_F))}{l} e^{-\gamma_\epsilon q(\mathbf{1})(t-l)} \Bigg). \quad (2)$$

Then $U_\epsilon(t) \ge \underline{U}_\epsilon(t)$ for all ϵ and t. If $t = \Omega(\epsilon^{-r_t})$ then for $0 < r_t < r + b_0$, $\lim_{\epsilon \to 0} U_\epsilon(t)/\underline{U}_\epsilon(t) = 1$.

Note that this lower bound is also in terms of regenerative cycle based measures, which can be estimated efficiently using importance sampling.

We mention at this point that other bounds on the unreliability, such as the ones in Brown (1990) and Kalashnikov (1989), can also be used. However, most of these bounds converge to the actual value only for $b_0 < r_t < r + b_0$, whereas ours converge for $0 < r_t < r + b_0$.

5 DERIVATIVE ESTIMATION FOR LARGE TIME HORIZONS

In this section, we derive bounds for the partial derivatives of the unreliability with respect to the failure rates of the components. For the sake of notational simplicity, we will now assume that there is no failure propagation. (For results in which failure propagation is allowed, see Shahabuddin and Nakayama 1993.) Also, we use the notation $\partial_i A(\lambda_1, \ldots, \lambda_R) = \frac{\partial}{\partial \lambda_i} A(\lambda_1, \ldots, \lambda_R)$ for some function $A(\lambda_1, \ldots, \lambda_R)$.

Now we make some additional definitions. Let

$$M_\epsilon(s) = 1 - e^{-\gamma_\epsilon q(\mathbf{1})s} - \gamma_\epsilon q(\mathbf{1})\, s\, e^{-\gamma_\epsilon q(\mathbf{1})s},$$

$$N_\epsilon(s) = M_\epsilon(s) - \frac{1}{2}\left(\gamma_\epsilon q(\mathbf{1})s\right)^2 e^{-\gamma_\epsilon q(\mathbf{1})s},$$

and

$$K_\epsilon = \left(e^{\gamma_\epsilon q(\mathbf{1})l}\left(1 + \gamma_\epsilon q(\mathbf{1})(t-l)\right) - 1 - \gamma_\epsilon q(\mathbf{1})t\right).$$

Also, consider component type i, and define τ_i to be the first (failure) transition of the DTMC $\{\mathbf{Y}_n, n \ge 0\}$ in which a component of type i fails. Nakayama (1991) showed that there exists constants $r_i \ge r$, $i = 1, \ldots, R$, (which depend on the model) such that

$$P_\Phi\{\tau_i \le \tau_F < \tau_{\mathbf{1}}\} = \Omega(\epsilon^{r_i}).$$

We denote the partial derivative of the unreliability by $U_\epsilon(t,i) \equiv \partial_i E_\Phi[I(T_F < t)]$. Now define

$$\bar{U}_\epsilon(t,i) = \bar{g}_\epsilon(t,i) - \underline{h}_\epsilon(t,i),$$

where

$$\bar{g}_\epsilon(t,i) = \frac{\alpha_\epsilon}{\gamma_\epsilon} M_\epsilon(t) + \frac{\beta_\epsilon}{\gamma_\epsilon} \bar{U}_\epsilon(t),$$

with $\alpha_\epsilon = E_\Phi[G\ 1\{\tau_{\mathbf{1}} < \tau_F\}]$, $\beta_\epsilon = E_\Phi[G\ 1\{\tau_F < \tau_{\mathbf{1}}\}]$, and $G = \sum_{l=0}^{\tau_{\min}-1} \frac{\partial_i q(\mathbf{Y}_l, \mathbf{Y}_{l+1})}{q(\mathbf{Y}_l, \mathbf{Y}_{l+1})}$, and

$$\underline{h}_\epsilon(t,i) = \frac{n_i}{q(\mathbf{1})\gamma_\epsilon} M_\epsilon(t) - \frac{n_i \xi_\epsilon}{l q(\mathbf{1})\gamma_\epsilon^2} M_\epsilon(t-l) - \frac{2 n_i \zeta_\epsilon}{l q(\mathbf{1})\gamma_\epsilon^2} N_\epsilon(t-l) - \frac{n_i e^{-q(\mathbf{1})\gamma_\epsilon t}}{q(\mathbf{1})\gamma_\epsilon} K_\epsilon,$$

where $\zeta_\epsilon = E_\Phi[W_{\min} 1\{\tau_{\mathbf{1}} < \tau_F\}]$, $\xi_\epsilon = E_\Phi[W_{\min} 1\{\tau_F < \tau_{\mathbf{1}}\}]$, and $W_{\min} = T_{\min} - Z$. (Recall that $l = \max(\sqrt{t}, t\sqrt{q})$ and n_i is the total number of components of type i.) Also, define

$$\underline{U}_\epsilon(t,i) = \underline{g}_\epsilon(t,i) - \bar{h}_\epsilon(t,i),$$

where

$$\underline{g}_\epsilon(t,i) = \bar{g}_\epsilon(t,i) - \frac{\nu_\epsilon}{l\gamma_\epsilon}\bar{U}_\epsilon(t-l) - \frac{1}{l\gamma_\epsilon}\left(\eta_\epsilon + \frac{\alpha_\epsilon\xi_\epsilon + \beta_\epsilon\zeta_\epsilon}{\gamma_\epsilon}\right) M_\epsilon(t-l) - \frac{2\alpha_\epsilon\zeta_\epsilon}{l\gamma_\epsilon^2} N_\epsilon(t-l) - \frac{e^{-q(\mathbf{1})\gamma_\epsilon t}}{\gamma_\epsilon}\left(\beta_\epsilon(e^{q(\mathbf{1})\gamma_\epsilon l}-1) + \alpha_\epsilon K_\epsilon\right),$$

with $\nu_\epsilon = E_\Phi[GW_{\min}1\{\tau_F < \tau_{\mathbf{1}}\}]$ and $\eta_\epsilon = E_\Phi[GW_{\min}1\{\tau_{\mathbf{1}} < \tau_F\}]$, and

$$\bar{h}_\epsilon(t,i) = \frac{n_i\xi_\epsilon}{\gamma_\epsilon}\bar{U}_\epsilon(t) + \frac{n_i}{\gamma_\epsilon}\left(\frac{1}{q(\mathbf{1})} + \zeta_\epsilon\right) M_\epsilon(t).$$

Note that α_ϵ, β_ϵ, ν_ϵ, η_ϵ, ζ_ϵ, and ξ_ϵ are all expectations of random quantities that are defined over a single regenerative cycle. Thus, we can efficiently estimate these quantities using importance sampling and form stable estimates of $\bar{U}_\epsilon(t,i)$ and $\underline{U}_\epsilon(t,i)$.

The following result shows that $\bar{U}_\epsilon(t,i)$ and $\underline{U}_\epsilon(t,i)$ are upper and lower bounds for the derivative of the unreliability with respect to the failure rate of component type i.

Theorem 6 *For all ϵ and t,*

$$\underline{U}_\epsilon(t,i) \le U_\epsilon(t,i) \le \bar{U}_\epsilon(t,i).$$

Now we determine when the upper bound is close to the actual value of the derivative. Let s_i be the state reached from state $\mathbf{1}$ from a component of type i failing. Then the following is true.

Theorem 7 *Suppose $t = \Omega(\epsilon^{-r_t})$ with $0 < r_t < r + r_q$. Then $U_\epsilon(t,i)/\bar{U}_\epsilon(t,i) \to 1$ as $\epsilon \to 0$ if one of the following holds:*

(i) $r_t < r - r_i + b_i$;

(ii) $r_t = r - r_i + b_i$ *and one of the following hold:*

- *(a) if* $s_i \in U$;
- *(b)* $s_i \in F$ *and* $t \ne 2\beta_\epsilon/(n_i\gamma_\epsilon) + o(\epsilon^{-r_t})$;

(iii) $r - r_i + b_i < r_t < \min\{2(r-r_i+b_i),\ r - r_i + b_i + r_q/2,\ r - r_i + b_i + 1\}$ *and* $s_i \in U$.

(iv) $r_t > r - r_i + b_i$ *and* $s_i \in F$.

Because of the subtraction of terms in the definition of $\bar{U}_\epsilon(t,i)$, we need the conditions in (ii), (iii), and (iv) to ensure that the highest order terms in the upper bound do not cancel. Note that one of the conditions of (ii) is almost always satisfied; they do not hold only for a small set of points. Thus, for smaller time horizons (i.e., when r_t satisfies either condition (i) or (ii)), the upper bound is typically good. However, for larger time horizons (i.e., when r_t satisfies condition (iii) or (iv)), the upper bound may not be as accurate.

The following theorem shows when the lower bound is close to the actual value of the derivative.

Theorem 8 *Suppose $t = \Omega(\epsilon^{-r_t})$ with $0 < r_t < r + r_q$. Then $U_\epsilon(t,i)/\underline{U}_\epsilon(t,i) \to 1$ as $\epsilon \to 0$ if one of the following holds:*

(i) $r_t < r - r_i + b_i$;

(ii) $r_t = r - r_i + b_i$ *and one of the following hold:*

- *(a)* $s_i \in U$;
- *(b)* $s_i \in F$ *and* $t \ne 2\beta_\epsilon/(n_i\gamma_\epsilon) + o(\epsilon^{-r_t})$;

(iii) $r - r_i + b_i < r_t < \min\{2(r-r_i+b_i),\ r - r_i + b_i + r_q/2,\ r - r_i + b_i + 1,\ 2(r - r_i + b_i + r_q)/3\}$ *and both of the following hold:*

- *(a)* $s_i \in U$;
- *(b)* $r_i < r + b_i$;

(iv) $r_t > r - r_i + b_i$ *and* $s_i \in F$.

Note that conditions (i), (ii), and (iv) of Theorem 8 are the same as those of Theorem 7. However, condition (iii) of Theorem 8 is stronger than that of Theorem 7. Thus, the upper bound may be better over a larger range of time horizons than the lower bound.

6 EXPERIMENTAL RESULTS

In this section we discuss some experimental results obtained using the SAVE package (see Goyal and Lavenberg 1987). The results are for a large computing system, previously considered by Goyal et al. (1992) and others. The system consists of two types of processors, each having a spare; two sets of disk controllers, each with a redundancy of two; and six disk clusters, each consisting of four disks. The system is operational as long as at least one processor of each type is operational and at least one disk controller from each set is operational and at least three out of the four disks in each disk cluster are operational. When a processor of either type fails, it causes a processor of the other type to fail simultaneously with probability 0.01. The failure rates of the processors, disk controllers, and disks are 1/2000 per hour, 1/2000 per hour, and 1/6000 per hour, respectively. Each of the component types can fail in one of two modes, each with probability 0.5. The repair rate in the first mode is 1 per hour, and in the second mode it is 1/2 per hour. There is a single repairperson who fixes failed components using preemptive random order service.

Results for the unreliability are given in Table 1. To see the effect of our simulation schemes for varying lengths of the time horizon t, we simulated for different values of t. The time horizon is given in the first column. The second column contains results using numerical computation (i.e., non-simulation) with the SAVE package. The CPU times expended to numerically compute the quantities are in parentheses. The numerical computations were inefficient as it took significant amounts of computer time (for large time horizons) and memory (for all cases). The next four columns contain the results using simulation. We simulated each case for 400,000 events, which took approximately 90 CPU seconds. The third column gives the RE corresponding to 99% confidence intervals (CI) if we use naive simulation (i.e., without using any importance sampling). (Estimates lie within the CI 99% of the time). The fourth column gives the estimate of the unreliability and the RE using balanced failure biasing and forcing. In the last column we give estimates of the upper bound $\bar{U}(t) = 1 - e^{-q(1)\gamma t}$. This we do by estimating γ and putting it in the equation for $\bar{U}(t)$. The γ can be estimated using failure biasing as done in Goyal et al. (1992). Similarly, in the fifth column we present estimates of the lower bound given by Equation 2. Note that in this example $E(T_c) \approx 125$ and $E(T_F) \approx 152,240$.

As the time horizon gets larger, the RE using failure biasing and forcing increases. However the RE on the upper bound and the lower bound estimate remains bounded. Also, in the region where failure biasing and forcing do not work well (i.e., for larger time horizons), the estimate of the bounds are very close to the actual value. Naive simulation seems to give moderately good RE for larger time horizons. However it is still much larger than the RE of the upper bound estimates.

Table 2 contains the estimates for the derivative of the unreliability with respect to the processor failure rate. All of the naive simulation estimates are poor, especially for very small and very large time horizons t. When we used importance sampling, the relative error of the estimator of the derivative of the unreliability is small for small t, but it deteriorates as the time horizon grows. All of our estimates of the lower and upper bounds given in the last two columns have small relative errors. However, the lower bound estimates provide reasonably accurate approximations to the actual values of the derivatives only for the middle range of t, and they are not as good for values of t which are very large or very small. Also, the upper bound estimates are close to the actual derivatives for small and medium values of t but not for very large t.

Table 1: Estimating the Unreliability

t	Num. ($\times 10^{-3}$)	Naive Sim. RE	IS Est. & RE ($\times 10^{-3}$)	LB Est. & RE ($\times 10^{-3}$)	UB Est. & RE ($\times 10^{-3}$)
4	0.016 (26s)	183%	0.015 ±8%	0.0016 ±9.6%	0.025 ±5.3%
16	0.087 (27s)	60%	0.086 ±8%	0.041 ±5.3%	0.099 ±5.3%
64	0.381 (28s)	38%	0.381 ±7%	0.258 ±5.3%	0.397 ±5.3%
256	1.55 (29s)	26%	1.59 ±7%	1.249 ±5.3%	1.59 ±5.3%
512	3.11 (33s)	24%	3.44 ±14%	2.606 ±5.3%	3.18 ±5.3%
1024	6.23 (40s)	23%	7.58 ±31%	5.314 ±5.3%	6.34 ±5.3%
2048	12.4 (52s)	22%	15.6 ±71%	10.71 ±5.3%	12.6 ±5.2%
4096	24.9 (80s)	22%	39.8 ±151%	21.1 ±5.3%	25.1 ±5.2%
8192	47.8 (170s)	21%	39.8 ±151%	42.4 ±5.3%	49.6 ±5.2%
16384	95.3 (400s)	20%	39.8 ±151%	83.1 ±5.3%	96.8 ±5.2%

ACKNOWLEDGEMENTS

The authors would like to thank Philip Heidelberger for helpful discussions on the subject.

REFERENCES

Brown, M. 1990. Error bounds for exponential approximations of geometric convolutions. *The Annals of Probability* 18:1388–1402.

Carrasco, J. A. 1991. Efficient transient simulation of failure/repair Markovian models. In *Proceedings of the Tenth Symposium on Reliable Distributed Systems*, IEEE Press, 152–161.

Glynn, P. 1992. Importance sampling for Markov chains: asymptotics for the variance. Technical Report, Department of Operations Research, Stanford University, California.

Goyal, A., and S. S. Lavenberg. 1987. Modelling and analysis of computer system availability. *IBM Journal of Research and Development* 31:651–664.

Table 2: Estimating the Sensitivity of the Unreliability w.r.t. Processor Failure Rate

t	Num. ($\times 10^{-4}$)	Naive Sim. RE	IS Est. & RE ($\times 10^{-4}$)	LB Est. & RE ($\times 10^{-4}$)	UB Est. & RE ($\times 10^{-4}$)
4	0.078	261%	0.082 ±13%	0.006 ±1.0%	0.13 ±9.0%
16	0.450	126%	0.440 ±13%	0.196 ±3.5%	0.53 ±8.9%
64	1.96	71%	1.95 ±12%	1.19 ±5.1%	2.2 ±8.7%
256	8.01	50%	8.41 ±16%	5.06 ±5.5%	9.6 ±8.3%
512	16.0	56%	19.8 ±35%	8.77 ±5.2%	21.3 ±8.1%
1024	32.0	67%	39.3 ±51%	10.4 ±4.8%	50.6 ±8.0%
2048	63.7	60%	32.2 ±122%	−9.06 ±5.2%	133 ±8.3%
4096	126	76%	28.4 ±156%	−138 ±6.5%	391 ±9.2%
8192	246	78%	28.4 ±156%	−739 ±7.9%	1267 ±10%
16384	468	174%	28.4 ±156%	−3215 ±8.9%	4348 ±11%

Goyal, A., P. Shahabuddin, P. Heidelberger, V. F. Nicola, and P. W. Glynn. 1992. A unified framework for simulating Markovian models of highly dependable systems. *IEEE Transactions on Computers* C-41:36–51.

Kalashnikov, V. V. 1989. Analytical and simulation estimates of reliability for regenerative models. *Syst. Anal. Model. Simul.* 6:833–851.

Lewis, E. E. and F. Bohm. 1984. Monte Carlo simulation of Markov unreliability models. *Nuclear Engineering and Design* 77:49–62.

Nakayama, M. K. 1991. Asymptotics for likelihood ratio derivative estimators in simulations of highly reliable Markovian systems. Research Report RC 17357, IBM T. J. Watson Research Center, Yorktown Heights, New York. Submitted for publication.

Nicola, V. F., P. Heidelberger, and P. Shahabuddin. 1992. Uniformization and exponential transformation: Techniques for fast simulation of highly dependable non-Markovian systems. In *Proceedings of the Twenty-Second Annual International Symposium on Fault Tolerant Computing*, IEEE Computer Society Press, 130–139.

Nicola, V. F., M. K. Nakayama, P. Heidelberger, and A. Goyal. 1991. Fast simulation of highly dependable systems with general failure and repair Processes. Research Report RC 16993, IBM T. J. Watson Research Center, Yorktown Heights, New York. To appear in *IEEE Transactions on Computers*.

Shahabuddin, P. 1990. Simulation and Analysis of Highly Reliable Systems. Ph.D. Thesis, Department of Operations Research, Stanford University, California.

Shahabuddin, P. 1991. Importance sampling for the simulation of highly reliable Markovian systems. Research Report RC 16729, IBM T. J. Watson Research Center, Yorktown Heights, New York. To appear in *Management Science*.

Shahabuddin, P. 1993. Fast transient simulation of Markovian models of highly dependable systems. To appear in *Proceedings of the Performance '93 Conference, Rome, Italy*.

Shahabuddin, P., and M. K. Nakayama. 1993. Fast simulation of transient measures and their derivatives in highly reliable Markovian systems. In preparation.

AUTHOR BIOGRAPHIES

PERWEZ SHAHABUDDIN received a B. Tech. in Mechanical Engineering from the Indian Institute of Technology, Delhi, in 1984. After working for a year in Engineers India Limited, India, he joined Stanford University from where he received a M.S. in Statistics in 1987 and a Ph.D. in Operations Research in 1990. Currently he is a Research Staff Member at the IBM T. J. Watson Research Center. His research interests include modelling and analysis of stochastic systems, with particular emphasis on communication and computer systems, simulation methodology and simulation output analysis. He won the first prize in the 1990 George E. Nicholson Student Paper Competition organised by the Operations Research Society of America.

MARVIN K. NAKAYAMA is currently a postdoctoral fellow at the IBM T. 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 a M.S. and Ph.D. in Operations Research from Stanford University. His research interests include simulation output analysis, gradient estimation, and rare event simulation. He won second prize in the 1992 George E. Nicholson Student Paper Competition organized by the Operations Research Society of America.