

Likelihood Ratio Sensitivity Analysis for Markovian Models of Highly Dependable Systems

Marvin K. Nakayama
IBM Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598

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

Peter W. Glynn
Department of Operations Research
Stanford University
Stanford, California 94305

Abstract

This paper discusses the application of the likelihood ratio gradient estimator to simulations of large Markovian models of highly dependable systems. Extensive empirical work, as well as some mathematical analysis of small dependability models, suggests that (in this model setting) the gradient estimators are not significantly more noisy than the estimates of the performance measures themselves. The paper also discusses implementation issues associated with likelihood ratio gradient estimation, as well as some theoretical complements associated with application of the technique to continuous-time Markov chains.

KEYWORDS: highly dependable systems, likelihood ratios, importance sampling, gradient estimators.

This paper discusses the application of the likelihood ratio gradient estimator to simulations of highly dependable systems. This paper makes the following contributions to the existing literature:

- 1 While the basis of the likelihood ratio gradient estimation algorithm has been known for some time (see, for example, Glynn 1986, 1987, and 1990, Reiman and Weiss 1989, and Rubinstein 1986 and 1989), much less is known about the empirical behavior of the estimator in practical problem settings. In this paper, we show, through extensive experimentation (see Section 5), that the likelihood ratio gradient estimator is an effective tool for measuring parameter sensitivity in the context of Markovian models of highly dependable systems. Both steady-state and terminating performance measures were studied. The positive results that we obtained for the steady-state gradient estimation problem are of particular interest, in light of the somewhat pessimistic conclusions reached in previous theoretical and empirical work (see, for example, Glynn 1987, L'Ecuyer, Giroux, and Glynn 1987, and Reiman and Weiss 1989). Thus, the results obtained here suggest that the steady-state likelihood ratio gradient estimator can be quite efficient when implemented in an appropriate problem setting.
- 2 The likelihood ratio gradient estimation algorithm, along with sophisticated variance reduction techniques, has been implemented in a widely distributed software package, namely the System Availability Estimator (SAVE) (see Goyal and Lavenberg 1987 and Goyal et al. 1992) developed within IBM. Because of the high degree of dependability of the systems typically simulated by SAVE, rare event simulation techniques (specifically, importance sampling) are used extensively within the package (so that failures can be observed). This paper describes how to combine likelihood ratio gradient estimation and importance sampling.
- 3 This paper shows how “discrete-time conversion” can be applied to the steady-state likelihood ratio gradient estimator (see also Glasserman 1992 and Reiman and Weiss 1989). This method reduces variance by removing variability due to the exponential holding time variates associated with the continuous-time Markov chain that is being simulated.
- 4 The computational burden imposed upon SAVE by the variance reduction techniques and likelihood ratio gradient estimator can be significant. For example, the numerical function evaluations required to compute the analytically-derived partial derivatives associated with the gradient estimator are time-consuming. Section 4 describes various ideas used within SAVE to improve the computational efficiency of the estimator.
- 5 Certain theoretical loose-ends concerned with the likelihood ratio gradient estimation technique are addressed within the paper. In particular, it is shown that for finite-state continuous-time Markov chains, the “amiability” assumption described in Reiman and Weiss (1989) and used in Glasserman (1992) is essentially always valid for reasonable performance measures (see the Appendix to this paper). Also, it is shown that discrete-time conversion applied to our steady-state gradient estimators is guaranteed to give a variance reduction.

An alternative approach for estimating derivatives is the infinitesimal perturbation analysis method; see Ho (1987) and Suri (1989) for overviews and extensive references. A drawback of the scheme is that it is typically more difficult to apply to a given application than the likelihood ratio approach (see Glynn 1990). In addition the perturbation analysis algorithm seems to be somewhat problem specific, thereby limiting its applicability. In particular the standard approach of implementing perturbation analysis does not work in our setting. However, Glasserman (1992) has developed an infinitesimal perturbation analysis methodology that can be used on a subset of the models which we consider. One difficulty is that the approach does not appear to extend easily to the setting in which failure propagation (see Section 1.1) is permitted.

This paper is organized as follows. Section 1 describes the basic mathematical model that is simulated by SAVE. In Sections 2 and 3, respectively, likelihood ratio gradient estimation for transient and steady-state performance measures is discussed. Section 3.4 also discusses certain insights that were obtained by analytically analyzing the behavior of the likelihood ratio gradient estimator for some (very) small models. In Section 4, implementation issues are discussed. Section 5 is devoted to a description and discussion of the experimental results obtained through extensive simulations of several large models having close to a million states. Section 6 discusses future research directions. The concluding Appendix contains most of the theoretical material alluded to in Item 5 above.

1 Problem Setting

In this section we briefly discuss the modeling problems being addressed by the SAVE package (see Goyal and Lavenberg 1987) and describe the basic mathematical model being simulated. We also describe various performance measures associated with the models which are considered.

1.1 Modeling Highly Dependable Systems

SAVE has been designed to construct and solve stochastic models of fault-tolerant computers. Fault-tolerant computing has been applied to two fundamentally different classes of applications. One deals with mission oriented systems with high reliability requirements, such as space computers, avionics systems, and ballistic missile defense computers (see Geist and Trivedi 1983). For the mission to succeed, the system must not fail during the mission time. Hence, the probability that the system does not fail during the mission time (i.e., the system reliability) is a measure of interest. Mean time to system failure is another measure that is used to evaluate such systems. The other class of applications deals with continuously operating systems with high availability requirements, such as telephone switching systems, general purpose computer systems, transaction processing systems (e.g., airline reservation systems), and communication network computers. For such systems, system failures can be tolerated if they occur infrequently and they result in short system down times. For such systems, the expected fraction of time the system is operational (i.e., the system availability) is a measure of interest.

From the modeling point of view, a system consists of a finite collection of hardware and software components, each of which may be subject to failure, recovery, and repair. Software components in operation can also be modeled with constant failure rates (see Laprie 1984). Component interactions often have a substantial effect on system availability and must therefore be considered in addition to the individual component behaviors. Thus, we have allowed for the possibility of component failure propagation; i.e., the failure of one component causes another component to fail immediately with some given probability. The state space size of such models grows (often exponentially) with the number of components being modeled. Therefore, SAVE provides a high level modeling language containing constructs which aid in representing the failure, recovery and repair behavior of components in the system as well as important component interactions.

If time independent failure and repair rates are assumed, then a finite state space, time homogeneous continuous-time Markov chain can be constructed automatically from the modeling constructs used to describe the system. Since the size of Markov chains grows exponentially with the number of components modeled, simulation appears to be a practical way for solving models of large systems. However, standard simulation takes very long simulation runs to estimate availability and reliability measures because the system failure event is a rare event. Therefore, variance reduction techniques which can aid in estimating rare-event probabilities quickly are of interest. Specifically, the importance sampling technique has been found to be most useful to estimate the various dependability measures (see Goyal et al. 1992). In this paper, we show that importance sampling also can be effective when combined with the likelihood ratio gradient estimation method to estimate derivatives. One change of measure is applied to compute the gradient using the likelihood ratio gradient estimation technique, and we employ another change of measure (importance sampling) to estimate these gradients quickly.

1.2 Markovian Model

Suppose $\mathbf{Y} = \{Y_s : s \geq 0\}$ is an irreducible, continuous-time Markov chain with finite state space E and infinitesimal generator $Q(\theta) = \{q(\theta, x, y) : x, y \in E\}$, where θ is a d -dimensional vector-valued parameter lying in some closed set Θ . We use the notation that P_θ and E_θ represent the probability measure and expectation, respectively, induced by the generator matrix $Q(\theta)$ for some value of θ . We assume that the state space E can be partitioned into two subsets: $E = O \cup F$, where O is the set of up states (i.e., the set of states for which the system is operational), and F is the set of down, or failed, states. We assume that the system starts out in the state in which all components are operational; we label this state as state 0. (A more detailed description of the CTMC \mathbf{Y} and its associated generator matrix $Q(\theta)$ is given in the Appendix.)

Let $\mathbf{X} = \{X_n : n \geq 0\}$ be the sequence of states visited by the chain and t_n be the time spent in each state, where $n \geq 0$. Also, we define $\mathbf{X}_n \equiv (X_0, X_1, \dots, X_n)$. Recall \mathbf{X} is a discrete time Markov chain (DTMC) with transition matrix $P(\theta)$ defined by $P(\theta, x, y) = q(\theta, x, y)/q(\theta, x)$ for $x \neq y$ and $P(\theta, x, x) = 0$, where $q(\theta, x) = -q(\theta, x, x)$. Furthermore, conditional on \mathbf{X} , the t_n 's are independent exponential random variables for which the (conditional) mean of t_n is $1/q(\theta, X_n)$.

Define $\{T_n : n \geq 0\}$ as the transition times of \mathbf{Y} ; i.e., $T_0 = 0$, and $T_n = t_0 + t_1 + \cdots + t_{n-1}$ for $n \geq 1$. Then define $N(t) = \sup\{n \geq 0 : T_n \leq t\}$.

For any set of states $A \subset E$, we let α_A denote the time the CTMC first enters the set A ; i.e., $\alpha_A = \inf\{s > 0 : Y_{s-} \notin A, Y_s \in A\}$. Of particular interest are α_0 , which is the first return time to state 0, and α_F , which is the first entrance time into the subset F of failed states. Our goals are to estimate (1) some performance measure $r(\theta) = E_\theta[Z(\theta)]$, where $Z(\theta)$ is some (measurable) function of \mathbf{Y} and (possibly) θ , and (2) its gradient $r'(\theta) = \frac{\partial}{\partial \theta} r(\theta)$. By varying our choice of the function Z , we can compute many different performance measures.

1.3 Performance Measures

The following development is based directly on material in Section 2.2 of Goyal et al. (1992). We will be interested in two types of dependability measures associated with the CTMC \mathbf{Y} : transient measures and so-called steady-state measures. Considering the transient measures first, the interval availability, $A(t)$, is defined by

$$A(t) = \frac{1}{t} \int_{s=0}^t 1_{\{Y_s \in O\}} ds.$$

This is the fraction of time that the system is operational in the time interval $(0, t)$. We let

$$I(t) = E_\theta[A(t)]$$

be the *expected interval availability*. Also let

$$F(t, x) = P_\theta\{A(t) \leq x\}$$

denote the *distribution of availability*. The *reliability* of the system is defined to be the probability that the system does not fail in the interval $(0, t)$:

$$R(t) = P_\theta\{\alpha_F > t\} = E_\theta[1_{\{\alpha_F > t\}}].$$

Since it was assumed that \mathbf{Y} is irreducible, we have that $Y_s \Rightarrow Y$ as $s \rightarrow \infty$, where \Rightarrow denotes convergence in distribution and Y is a random variable having the steady-state distribution $\pi = \{\pi_x, x \in E\}$ (π solves the equations $\pi Q = 0$). Notice that steady-state measures are independent of the starting state of the system; however, we will choose the fully operational state (i.e., state 0) to define a regenerative state for the system; i.e., the successive times at which the process makes a transition to state 0 form a sequence of regeneration points. Also, we assume that, when computing steady-state measures, we can express $Z(\theta)$ as $Z(\theta) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_{s=0}^t f(\theta, Y_s) ds$, where $f(\theta, \cdot)$ is a real-valued function on E for which $f(\cdot, x)$ is continuously differentiable in θ for each $x \in E$. By regenerative process theory (see Crane and Iglehart 1974), our steady-state measures take the form of a ratio of two expected values:

$$r(\theta) = E_\theta[Z(\theta)] = \frac{E_\theta[\int_{s=0}^{\alpha_0} f(\theta, Y_s) ds]}{E_\theta[\alpha_0]}.$$

If $f(\theta, x) = 1_{\{x \in O\}}$, then $E_\theta[Z(\theta)]$ is the long run fraction of time the system is operational and is called the *steady-state availability*. We denote this by $A = \lim_{t \rightarrow \infty} E_\theta[A(t)]$. We will sometimes find it convenient to consider the *expected interval unavailability* $U(t) = 1 - I(t) = 1 - E_\theta[A(t)]$ and the *steady-state unavailability*, $U = 1 - A$. The problem of steady-state estimation thus reduces to one of estimating the ratio of two expected values.

The *mean time to failure* (MTTF), $E_\theta[\alpha_F]$, is typically thought of as a transient measure, since it depends on the starting state of the system (state 0), which is assumed to be the fully operational state. A ratio representation for $E_\theta[\alpha_F]$ is found to be particularly useful and is given by

$$E_\theta[\alpha_F] = \frac{E_\theta[\min(\alpha_F, \alpha_0)]}{P_\theta\{\alpha_F < \alpha_0\}}.$$

The derivation of this formula is given in Goyal et al. (1992). Thus, we can view estimating $E_\theta[\alpha_F]$ as a ratio estimation problem, where both the numerator and the denominator are estimated using a regenerative simulation. Therefore, in Section 3 we consider the estimation of the mean time to failure (MTTF) together with steady-state measures, which are also (and more commonly) estimated using regenerative simulations.

2 Estimating Transient Performance Measures

In this section we discuss the estimation of transient performance measures and their derivatives. We present an overview of the likelihood ratio method as applied to the estimation of derivatives of transient performance measures. Also, we show how importance sampling can be applied to estimate more efficiently the performance measures and derivatives.

2.1 Using the Likelihood Ratio Method to Estimate Gradients of Transient Performance Measures

Recall that our goals are to estimate $r(\theta) = E_\theta[Z(\theta)]$ and its derivative $r'(\theta)$ with respect to the parameter θ . In the case of transient performance measures, we assume that our function $Z(\theta)$ has one of the two following forms:

- 1 $Z(\theta) = 1_S$, where S is some event determined by the process \mathbf{Y} up to time T , where T is a stopping time satisfying Assumption A5 in the Appendix.
- 2 $Z(\theta) = \int_0^T f(\theta, Y_s) ds$, where T is some stopping time satisfying Assumption A5 given in the Appendix and $f(\theta, \cdot)$ is a real-valued function on E for which $f(\cdot, x)$ is continuously differentiable in θ for each $x \in E$.

Let T be a stopping time satisfying Assumption A5 in the Appendix. We define the “likelihood” of a sample path up to time T under parameter θ as

$$d\mu(T, \theta) = \left[\prod_{k=0}^{N(T)-1} q(\theta, X_k) \exp\{-q(\theta, X_k)t_k\} P(\theta, X_k, X_{k+1}) \right] \cdot \exp\{-q(\theta, X_{N(T)})(T - T_{N(T)})\}, \quad (1)$$

and the likelihood ratio is given by

$$L(T, \theta, \theta_0) = d\mu(T, \theta)/d\mu(T, \theta_0), \quad (2)$$

where θ_0 is some fixed value of θ lying in the interior of Θ .

Thus, we express our performance measure as

$$r(\theta) = E_\theta[Z(\theta)] = E_{\theta_0}[Z(\theta)L(T, \theta, \theta_0)].$$

We call this transformation a “change of measure” since we are now computing the expectation based on a different probability measure. (In this case, the new probability measure has a different parameter value.) Billingsley (1961) and Goyal et al. (1992) discuss the validity of the change of measure. By performing the change of measure, the expectation operator is now independent of the parameter θ .

If we formally differentiate this expression, assuming that we can interchange the derivative and expectation operators, we have that by applying the product rule of differentiation,

$$r'(\theta) = E_{\theta_0}[Z'(\theta)L(T, \theta, \theta_0)] + E_{\theta_0}[Z(\theta)L'(T, \theta, \theta_0)],$$

where $Z'(\theta) = 0$ if $Z(\theta)$ has form 1 above, and $Z'(\theta) = \int_0^T f'(\theta, Y_s)ds$ if $Z(\theta)$ has form 2 above, and

$$L'(T, \theta, \theta_0) = \left[\sum_{k=0}^{N(T)-1} \left\{ \frac{q'(\theta, X_k)}{q(\theta, X_k)} - q'(\theta, X_k)t_k + \frac{P'(\theta, X_k, X_{k+1})}{P(\theta, X_k, X_{k+1})} \right\} - q'(\theta, X_{N(T)})(T - T_{N(T)}) \right] L(T, \theta, \theta_0). \quad (3)$$

The proof of the validity of the interchange of derivative and expectation is given in Theorem 1 in the Appendix.

The terms simplify when we evaluate $r'(\theta)$ at the point $\theta = \theta_0$. In this case, since $L(T, \theta_0, \theta_0) = 1$, we have

$$r'(\theta_0) = E_{\theta_0}[Z'(\theta_0)] + E_{\theta_0}[Z(\theta_0)L'(T, \theta_0, \theta_0)] \quad (4)$$

and

$$L'(T, \theta_0, \theta_0) = \left[\sum_{k=0}^{N(T)-1} \left\{ \frac{q'(\theta_0, X_k)}{q(\theta_0, X_k)} - q'(\theta_0, X_k)t_k + \frac{P'(\theta_0, X_k, X_{k+1})}{P(\theta_0, X_k, X_{k+1})} \right\} - q'(\theta_0, X_{N(T)})(T - T_{N(T)}) \right]. \quad (5)$$

Note that if T is either the time of the first transition after a deterministic time t or a hitting time to a set F , the $-q(\theta_0, X_{N(T)})(T - T_{N(T)})$ term drops out.

These results are similar to those derived in Reiman and Weiss (1989); however, there, only the specialization to Poisson processes is discussed. Also, the “amiability” assumption employed in Glasserman (1992) and Reiman and Weiss (1989) holds in the current context, and examples of performance measures satisfying this condition were discussed in Section 1.3 and in the Appendix. L’Ecuyer (1990) also gives a condition which can be used to justify the validity of the interchange of the derivative and expectation operators in a more general setting than the one which we are considering. However, we prove more directly that the result holds in our setting; see Theorem 1 in the Appendix.

2.2 Importance Sampling to Reduce Variance

We now apply another change of measure to implement the importance sampling. As a generalization of (1), let $\mu^*(T, \theta_0)$ be a probability measure such that $\mu^*(T, \theta_0)(B) > 0$ whenever $\mu(T, \theta_0)(B) > 0$, where B is a (measurable) set of sample paths up to time T . Then by applying a change of measure to the right hand side of (4), we obtain

$$r'(\theta_0) = E_{\theta_0}^*[Z'(\theta_0)L_*(T, \theta_0)] + E_{\theta_0}^*[Z(\theta_0)L'(T, \theta_0, \theta_0)L_*(T, \theta_0)], \quad (6)$$

where $E_{\theta_0}^*$ is the expectation operator under the measure $\mu^*(T, \theta_0)$ and

$$L_*(T, \theta_0) = \frac{d\mu(T, \theta_0)}{d\mu^*(T, \theta_0)}.$$

Rather than using the same measure $\mu^*(T, \theta_0)$ to evaluate both expectations on the right hand side of (6), we could employ two different changes of measure $\hat{\mu}^*(T, \theta_0)$ and $\tilde{\mu}^*(T, \theta_0)$ to estimate the two expectations. This is known as “measure-specific importance sampling”; see Goyal, Heidelberger, and Shahabuddin (1987). In this paper we will only use one change of measure to evaluate all expectations. Also, we can define the new measure $\mu^*(T, \theta_0)$ such that it depends on \mathbf{X}_k , the entire sequence of states visited up to that point, and also on T_k , the time of the k^{th} transition. We call this method “dynamic importance sampling” (DIS). These ideas are discussed in Section 4 and also in Goyal, Heidelberger, and Shahabuddin (1987).

We choose $\mu^*(T, \theta_0)$ such that $\mu^*(T, \theta_0) = \mu_1^*(T, \theta_0)\mu_2^*(T, \theta_0)$, where $\mu_1^*(T, \theta_0)$ is the probability measure of the state transitions and $\mu_2^*(T, \theta_0)$ is the probability measure of the holding times in each state visited (conditional on the states visited). Thus, we can separate the likelihood ratio into two different components. The first component includes only the transition probabilities, and the second incorporates only the random holding times; i.e.,

$$L_*(T, \theta_0) = L_{1*}(T, \theta_0)L_{2*}(T, \theta_0),$$

where

$$L_{1*}(T, \theta_0) = \prod_{k=0}^{N(T)-1} \frac{P(\theta_0, X_k, X_{k+1})}{P_{\theta_0}^*(X_k | \mathbf{X}_{k-1})}, \quad (7)$$

$P_{\theta_0}^*(X_k | \mathbf{X}_{k-1})$ is the transition probability (conditional on \mathbf{X}_{k-1}) under importance sampling,

$$L_{2*}(T, \theta_0) = \prod_{k=0}^{N(T)-1} \frac{q(\theta_0, X_k) \exp\{-q(\theta_0, X_k)t_k\}}{\mathbf{h}_{\theta_0}(t_k | X_k, T_k)} \cdot \frac{\exp\{-q(\theta_0, X_{N(T)})(T - T_{N(T)})\}}{\int_0^{T-T_{N(T)}} \mathbf{h}_{\theta_0}(s | X_{N(T)}, T_{N(T)}) ds}, \quad (8)$$

and $\mathbf{h}_{\theta_0}(t_k | X_k, T_k)$ is the density of the holding time in state X_k (conditional on X_k and T_k) under importance sampling. Decomposing the likelihood ratio in this manner allows us to tailor one change of measure for the transition probabilities and another for the holding times.

Lewis and Böhm (1984) presented an importance sampling technique for estimating transient measures. They apply “failure biasing” to the embedded DTMC; this causes failures to occur with higher probability and therefore quickly moves (biases) the DTMC towards the set of failed states. This change of measure is incorporated in the first component of the likelihood ratio L_{1*} . They also apply “forced transitions” to the holding time in state 0 (the state with all components operational) to estimate the reliability. This forces the next component failure to occur before time t . Specifically, if $X_n = 0$ and $T_n < t$, then the next holding time, t_n , is forced to be between zero and $t - T_n$ by selecting t_n from the conditional density given by

$$\mathbf{h}_{\theta_0}(t_n|X_n, T_n) = \frac{\lambda_0 e^{-\lambda_0 t_n}}{1 - e^{-\lambda_0(t-T_n)}},$$

where $0 \leq t_n \leq t - T_n$ and λ_0 is the total failure rate in state 0. This change of measure is incorporated into the second part of the likelihood ratio, L_{2*} . The simulation continues until time $T = \min(\alpha_F, T_{N(t)+1})$.

Note that $\mathbf{h}_{\theta_0}(t_n|X_n, T_n)$ is not positive whenever the exponential density is positive. According to the standard theory of importance sampling, this is not a legitimate change of measure. However, in this case, $\mathbf{h}_{\theta_0}(t_n|X_n, T_n)$ is positive over that part of the sample space ($\{\omega : \alpha_F < t\}$) that counts, which is sufficient (see Glynn and Iglehart 1989).

3 Estimating Steady-State Performance Measures

In this section we discuss the estimation of steady-state performance measures and their derivatives. We show how the likelihood ratio method can be applied to the estimation of derivatives of steady-state performance measures. We also describe some variance reduction techniques that can be applied to steady-state derivative estimation. Finally, we investigate analytically the behavior of the estimators of the steady-state unavailability and its derivatives in some small models.

3.1 Using the Likelihood Ratio Method to Estimate Gradients of Steady-State Performance Measures

Recall that our goal is to estimate $r(\theta) = E_{\theta}[Z(\theta)]$ and its derivative. For steady-state performance measures, we consider $Z(\theta)$ of the form

$$Z(\theta) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t f(\theta, Y_s) ds,$$

where $f(\theta, \cdot)$ is a real-valued function on E for which $f(\cdot, x)$ is continuously differentiable for each $x \in E$. Since we assumed that the CTMC \mathbf{Y} has finite state space E and the transition matrix P for the embedded DTMC is irreducible, we can define our stopping time T to be the time of the first return to the initial state 0; i.e., $T = \alpha_0$. Let τ_0 be the first return time of the embedded DTMC to state 0; i.e., $\tau_0 = \inf\{n \geq 1 : X_n = 0\}$. Since \mathbf{Y} is a CTMC, T is a regeneration time. Hence, assuming that $E_{\theta}[|Z(\theta)|] < \infty$, we express $r(\theta)$ using the ratio formula

$$r(\theta) = \frac{E_{\theta}[Z_T(\theta)]}{E_{\theta}[T]}, \tag{9}$$

where

$$Z_T(\theta) = \int_0^T f(\theta, Y_s) ds = \sum_{k=0}^{\tau_0-1} f(\theta, X_k) t_k$$

and

$$T = \alpha_0 = \sum_{k=0}^{\tau_0-1} t_k.$$

Let $(\hat{Z}_{T,j}(\theta), \hat{T}_j)$, $j = 1, 2, \dots, m$, denote independent, identically distributed observations of $(Z_T(\theta), T)$. Then the ratio estimator satisfies the following central limit theorem:

$$\sqrt{m}(\hat{r}_m - r) \Rightarrow N(0, \sigma_r^2)$$

as $m \rightarrow \infty$, where $\hat{r}_m = \sum_{j=1}^m \hat{Z}_{T,j}(\theta) / \sum_{j=1}^m \hat{T}_j$ and $\sigma_r^2 = \text{Var}[Z_{T,j}(\theta) - rT_j] / E[T_j]^2$. See Crane and Iglehart (1974) for further details.

Formally differentiating the expression for $r(\theta)$ by interchanging the derivative and expectation, we obtain

$$r'(\theta_0) = \frac{u'(\theta_0)l(\theta_0) - l'(\theta_0)u(\theta_0)}{l^2(\theta_0)}, \quad (10)$$

where

$$\begin{aligned} u(\theta_0) &= E_{\theta_0}[Z_T(\theta_0)] \\ u'(\theta_0) &= E_{\theta_0}[Z_T'(\theta_0)] + E_{\theta_0}[Z_T(\theta_0)L'(T, \theta_0, \theta_0)] \\ l(\theta_0) &= E_{\theta_0}[T] \\ l'(\theta_0) &= E_{\theta_0}[TL'(T, \theta_0, \theta_0)] \end{aligned}$$

and

$$\begin{aligned} Z_T'(\theta_0) &= \sum_{k=0}^{\tau_0-1} f'(\theta_0, X_k) t_k \\ L'(T, \theta_0, \theta_0) &= \sum_{k=0}^{\tau_0-1} \left[\frac{q'(\theta_0, X_k)}{q(\theta_0, X_k)} - q'(\theta_0, X_k) t_k + \frac{P'(\theta_0, X_k, X_{k+1})}{P(\theta_0, X_k, X_{k+1})} \right]. \end{aligned}$$

The proof of the validity of the interchange of the operators is given in the Appendix.

To construct confidence intervals for our estimate of $r'(\theta)$, we need an expression for its asymptotic variance. This is given by

$$\begin{aligned} & \frac{1}{\beta^2} \sigma_A^2 + \left[\frac{2\gamma\delta - \alpha\beta}{\beta^3} \right]^2 \sigma_B^2 + \frac{\delta^2}{\beta^4} \sigma_C^2 + \frac{\gamma^2}{\beta^4} \sigma_D^2 \\ & + 2 \left[\frac{2\gamma\delta - \alpha\beta}{\beta^4} \sigma_{AB} - \frac{\delta}{\beta^3} \sigma_{AC} - \frac{\gamma}{\beta^3} \sigma_{AD} - \frac{2\gamma\delta^2 - \alpha\beta\delta}{\beta^5} \sigma_{BC} - \frac{2\gamma^2\delta - \alpha\beta\gamma}{\beta^5} \sigma_{BD} + \frac{\gamma\delta}{\beta^4} \sigma_{CD} \right], \end{aligned} \quad (11)$$

where $\sigma_X^2 = \text{Var}(X)$, $\sigma_{XY} = \text{Cov}(X, Y)$,

$$A = Z_T'(\theta_0) + Z_T(\theta_0)L'(T, \theta_0, \theta_0) \quad (12)$$

$$B = T \quad (13)$$

$$C = Z_T(\theta_0) \quad (14)$$

$$D = TL'(T, \theta_0, \theta_0), \quad (15)$$

and

$$\begin{aligned}\alpha &= E_{\theta_0}[A] \\ \beta &= E_{\theta_0}[B] \\ \gamma &= E_{\theta_0}[C] \\ \delta &= E_{\theta_0}[D].\end{aligned}$$

A proof of the validity of the expression for the variance is given in Reiman and Weiss (1989). However, we give a simpler proof in the Appendix.

3.2 Conditioning to Reduce Variance

Conditional Monte Carlo is a technique which reduces the variance in simulations of CTMCs (see Fox and Glynn 1986 and Hordijk, Iglehart, and Shassberger 1976). By conditioning on the embedded DTMC \mathbf{X} , we arrive at what is known as the discrete time method, in which the holding times, t_n , are replaced by their (conditional) means, $1/q(\theta, X_n)$. There are two advantages of using this approach. First, since we replace the random holding times t_n with their (conditional) means, we do not have to generate exponential variates. Thus, there is an increase in computational efficiency. Also, as discussed in Fox and Glynn (1986) and Hordijk, Iglehart, and Shassberger (1976), this transformation is guaranteed to give a reduction in the variance of the estimate of $r(\theta)$. We also show that the transformation is guaranteed to reduce the variance of the estimate of $r'(\theta)$. Hence, the statistical efficiency is improved.

Using conditional Monte Carlo, we obtain another ratio formula

$$r(\theta) = \frac{E_{\theta}[E_{\theta}[Z_T(\theta)|\mathbf{X}]]}{E_{\theta}[E_{\theta}[T|\mathbf{X}]]} = \frac{E_{\theta}[G(\theta)]}{E_{\theta}[H(\theta)]}, \quad (16)$$

where a straightforward calculation shows that

$$G(\theta) = \sum_{k=0}^{\tau_0-1} g(\theta, X_k) \quad (17)$$

$$H(\theta) = \sum_{k=0}^{\tau_0-1} h(\theta, X_k) \quad (18)$$

$$g(\theta, x) = f(\theta, x)/q(\theta, x), \quad x \in E$$

$$h(\theta, x) = 1/q(\theta, x), \quad x \in E$$

and τ_0 is the first return time of the DTMC to state 0; see Goyal et al. (1992) for further details. Note that (9) and (16) are equivalent representations of $r(\theta)$; however, the estimators which arise from the two expressions are different. To form an estimator of $r(\theta)$ based on (9), we collect samples of $Z_T(\theta)$ and T , whereas we accumulate samples of $G(\theta)$ and $H(\theta)$ when basing an estimator on (16).

In Glynn (1986), it is shown that under certain conditions (viz., Assumption A4 given in the Appendix)

$$r(\theta) = \frac{E_{\theta_0}[G(\theta)\hat{L}(\tau_0, \theta, \theta_0)]}{E_{\theta_0}[H(\theta)\hat{L}(\tau_0, \theta, \theta_0)]}, \quad (19)$$

where $\hat{L}(\tau_0, \theta, \theta_0)$ is the DTMC likelihood ratio, which is defined as

$$\hat{L}(\tau_0, \theta, \theta_0) = \prod_{k=0}^{\tau_0-1} \frac{P(\theta, X_k, X_{k+1})}{P(\theta_0, X_k, X_{k+1})}. \quad (20)$$

A simple calculation shows that $\hat{L}(\tau_0, \theta, \theta_0) = E_{\theta}[L(\alpha_0, \theta, \theta_0)|\mathbf{X}]$. Formally differentiating the right hand side of (19) by interchanging the derivative and expectation, we obtain

$$r'(\theta_0) = \frac{\hat{u}'(\theta_0)\hat{l}(\theta_0) - \hat{l}'(\theta_0)\hat{u}(\theta_0)}{\hat{l}^2(\theta_0)}, \quad (21)$$

where

$$\begin{aligned} \hat{u}(\theta_0) &= E_{\theta_0}[G(\theta_0)] \\ \hat{u}'(\theta_0) &= E_{\theta_0}[G'(\theta_0)] + E_{\theta_0}[G(\theta_0)\hat{L}'(\tau_0, \theta_0, \theta_0)] \\ \hat{l}(\theta_0) &= E_{\theta_0}[H(\theta_0)] \\ \hat{l}'(\theta_0) &= E_{\theta_0}[H'(\theta_0)] + E_{\theta_0}[H(\theta_0)\hat{L}'(\tau_0, \theta_0, \theta_0)] \end{aligned}$$

and

$$G'(\theta_0) = \sum_{k=0}^{\tau_0-1} \frac{f'(\theta_0, X_k)q(\theta_0, X_k) - q'(\theta_0, X_k)f(\theta_0, X_k)}{q^2(\theta_0, X_k)} \quad (22)$$

$$H'(\theta_0) = \sum_{k=0}^{\tau_0-1} -\frac{q'(\theta_0, X_k)}{q^2(\theta_0, X_k)} \quad (23)$$

$$\hat{L}'(\tau_0, \theta_0, \theta_0) = \sum_{k=0}^{\tau_0-1} \frac{P'(\theta_0, X_k, X_{k+1})}{P(\theta_0, X_k, X_{k+1})}. \quad (24)$$

The proof of the validity of the interchange of the operators in this case is given in Glynn (1986).

Note that (10) and (21) are both exact representations of $r'(\theta_0)$ which give rise to different estimators. To form an estimator of $r'(\theta_0)$ based on (10), we collect samples of A , B , C , and D given in (12)–(15). We accumulate samples of $G'(\theta_0) + G(\theta_0)\hat{L}'(\tau_0, \theta_0, \theta_0)$, $H(\theta_0)$, $G(\theta_0)$, and $H'(\theta_0) + H(\theta_0)\hat{L}'(\tau_0, \theta_0, \theta_0)$, respectively, when basing an estimator on (21).

To obtain an expression for the asymptotic variance of the derivative estimator based on (21), we modify (11) by replacing A , B , C , and D given in (12)–(15) with their conditional expectations. Let σ_1^2 and σ_2^2 be the variances of the gradient estimators when using the ratio formula obtained without and with conditional Monte Carlo, respectively; i.e., σ_1^2 , which is given in (11), is the asymptotic variance of the estimator of derivative given in (10), and σ_2^2 is the asymptotic variance of the estimator of the derivative given in (21). Then, we have that $\sigma_2^2 \leq \sigma_1^2$, which states that when using the ratio formula, conditional Monte Carlo always gives rise to a lower asymptotic variance constant (see Proposition 1 in the Appendix).

3.3 Importance Sampling to Reduce Variance

As in Section 2.2, we can use importance sampling by applying another change of measure. However, since in this case we use conditional Monte Carlo to condition out the holding times in each state when estimating steady-state performance measures, the likelihood ratio only consists of its first component L_{1*} , given in (7).

3.4 Two Simple Examples

In this section, we consider two simple availability examples. The first is a one-dimensional birth and death process with three states, which is also analyzed in Goyal et al. (1992), and the second is a two-dimensional five state birth and death process. Because of their simple structure, we are able to extensively analyze these models.

Before analyzing the models, we need to make a definition. We define the *sensitivity* $r_i(\theta)$ of a performance measure $r(\theta)$ with respect to a certain parameter θ_i to be the product of the parameter itself multiplied by the partial derivative of the performance measure with respect to the parameter θ_i ; i.e.,

$$r_i(\theta) = \theta_i \cdot \frac{\partial}{\partial \theta_i} r(\theta).$$

Sensitivities measure the effects on the overall system performance of relative changes in the value of a parameter.

We will show that (for our two examples)

- 1 when one sensitivity is much larger in magnitude than another, the relative accuracy (as measured by the squared coefficient of variation) of the the estimate of the larger one is much better than that of the smaller one;
- 2 we can estimate the sensitivities with the largest orders of magnitude with about the same relative accuracies as the performance measure estimate, as long as each sample (e.g., a regenerative cycle in the case of steady-state estimation) consists of only a few transitions. This is true in the highly reliable component situation which we consider in this paper.

Much of the analysis was done using the symbolic manipulator Scratchpad (see Sutor 1989).

We define the vector of parameters θ with respect to which we compute sensitivities as the vector of all continuous-valued parameters of the model. Note that items 1 and 2 above depend on the parameterization of the model. However, in the reliability context which we are considering in this paper, there is a natural parameterization of the model, which is to have θ consist of the values of all of the component failure rates and repair rates. With θ defined in this manner, we will see that items 1 and 2 above hold for our models.

3.4.1 A Three State Example

The three state example, which is taken directly from Goyal et al. (1992), can be viewed as a reliability system in which there is one type of component with redundancy of two and the components

fail and are repaired. The components have failure rate λ and repair rate μ . The state space is $E = \{0, 1, 2\}$. We assume that births correspond to failures and deaths correspond to repairs so that state i corresponds to having i failed components. We consider the system to be operational in states 0 and 1 but failed in state 2. Hence, $O = \{0, 1\}$ and $F = \{2\}$.

The transition matrix P of the embedded DTMC has the following non-zero entries: $P(0, 1) = P(2, 1) = 1$, $P(1, 2) = \lambda/(\lambda + \mu)$, and $P(1, 0) = \mu/(\lambda + \mu)$. Using the method of conditional Monte Carlo, we let h_i be the mean holding time in state i . Thus, $h_0 = 1/(2\lambda)$, $h_1 = 1/(\lambda + \mu)$, and $h_2 = 1/\mu$. Since we are working with highly reliable systems, we assume that $\lambda \ll \mu$. We first allow μ to vary so that we are able to compute the partial derivative with respect to μ . However, after computing derivatives, we fix $\mu = 1$, which only fixes the time scale we are considering.

We are interested in the steady-state unavailability r , which is the steady-state probability of being in the failed state 2. Recall that we can estimate this quantity using the regenerative method and can express r as the ratio $E[G]/E[H]$, as in (16). In this example, we set $f(0) = f(1) = 0$ and $f(2) = 1$. We select state 0 as the regenerative state, and so $G = n_F h_2$ and $H = h_0 + h_1 + n_F(h_1 + h_2)$, where n_F is the number of times the failure state is reached in the regenerative cycle. Note that n_F has a geometric distribution, so that $E[n_F] = \lambda/\mu$ and $\text{Var}[n_F] = \lambda(\lambda + \mu)/\mu^2$. Thus, $E[G] = h_2 \lambda/\mu$ and $E[H] = (h_0 + h_1) + (h_1 + h_2)\lambda/\mu$. As shown in Goyal et al. (1992), $r = \Omega(\lambda^2)$ and

$$CV^2(r, m) = \frac{\text{Var}[G - rH]}{mr^2(E[H])^2} = \frac{1}{m}\Omega\left(\frac{1}{\lambda}\right),$$

where $CV^2(r, m)$ denotes the asymptotic squared coefficient of variation of our estimate of r after m regenerative cycles (we modify the notation of Knuth 1976 to mean $f(x) = \Omega(g(x))$ if there exist constants C_1 and C_2 such that for all x sufficiently small, $0 < C_1 g(x) < f(x) < C_2 g(x)$).

Straightforward calculations show that the sensitivities $r_\lambda = \Omega(\lambda^2)$ and $r_\mu = \Omega(\lambda^2)$, where we use the notation $r_\theta = \theta \cdot \partial r / \partial \theta$. Hence, r_θ is the sensitivity of the performance measure r with respect to the parameter θ . Using the asymptotic variance for estimates of derivatives that arise when using conditional Monte Carlo (see Section 3.2), we obtain expressions for the asymptotic squared coefficients of variation of our sensitivity estimates, which are given by $CV^2(r_\lambda, m) = \Omega(1/\lambda)/m$ and $CV^2(r_\mu, m) = \Omega(1/\lambda)/m$. All of the variance and covariance terms in the expression for the asymptotic variance of the gradients were used in the calculations in this example. It turns out that the dominant terms in the expression for the variance of the gradients are the ones involving the variances of the down time in a cycle G and its derivatives.

Thus, when $\lambda \ll \mu$, both of the sensitivities are of the same order as the performance measure estimate. In addition, the relative accuracies of the performance measure estimate and the sensitivity estimates are within a constant factor that is independent of λ .

3.4.2 A Five State Example

The five state example models a system with two types of components, each of which has a redundancy of two. There is also the added restriction that once a component of one type fails, the components of the other type cannot fail until the state with all components operational is reached.

Thus the state space of this example is $E = \{(0, 0), (1, 0), (2, 0), (0, 1), (0, 2)\}$, where in state (i, j) , i represents the number of failed components of type 1, and j is the number of failed components of type 2. We select the regenerative state as the state in which all components are operational, i.e., $(0, 0)$. We consider the system to be operational in states $(0, 0)$, $(1, 0)$, and $(0, 1)$, and failed in states $(2, 0)$ and $(0, 2)$. Hence, $O = \{(0, 0), (1, 0), (0, 1)\}$ and $F = \{(2, 0), (0, 2)\}$. Let λ_i denote the failure rate of component type i , and let μ denote the repair rate of both types of components. We assume that $\lambda_2 \ll \lambda_1 \ll \mu = 1$.

The transition matrix of the embedded DTMC has the following non-zero entries:

$$\begin{aligned}
P((0, 0), (1, 0)) &= \lambda_1/(\lambda_1 + \lambda_2) \\
P((1, 0), (2, 0)) &= \lambda_1/(\lambda_1 + \mu) \\
P((1, 0), (0, 0)) &= \mu/(\lambda_1 + \mu) \\
P((0, 0), (0, 1)) &= \lambda_2/(\lambda_1 + \lambda_2) \\
P((0, 1), (0, 2)) &= \lambda_2/(\lambda_2 + \mu) \\
P((0, 1), (0, 0)) &= \mu/(\lambda_2 + \mu) \\
P((2, 0), (1, 0)) &= P((0, 2), (0, 1)) = 1.
\end{aligned}$$

Let $h_{(i,j)}$ denote the mean holding time in state (i, j) . Thus, $h_{(0,0)} = 1/(2\lambda_1 + 2\lambda_2)$, $h_{(1,0)} = 1/(\lambda_1 + \mu)$, $h_{(0,1)} = 1/(\lambda_2 + \mu)$, and $h_{(2,0)} = h_{(0,2)} = 1/\mu$.

In this example, we set $f(0, 0) = f(1, 0) = f(0, 1) = 0$ and $f(2, 0) = f(0, 2) = 1$. Using the ratio formula again to estimate the steady-state unavailability, we have that

$$G = 1_{\{X_1=(1,0)\}}n_1h_{(2,0)} + 1_{\{X_1=(0,1)\}}n_2h_{(0,2)},$$

and

$$\begin{aligned}
H &= h_{(0,0)} + 1_{\{X_1=(1,0)\}}(h_{(1,0)} + n_1(h_{(1,0)} + h_{(2,0)})) \\
&\quad + 1_{\{X_1=(0,1)\}}(h_{(0,1)} + n_2(h_{(0,1)} + h_{(0,2)})),
\end{aligned}$$

where X_1 is the first state visited by the embedded DTMC, n_1 is the number of times state $(2, 0)$ is visited in the regenerative cycle, and n_2 is the number of times state $(0, 2)$ is visited in the regenerative cycle. Note that, conditional on $X_1 = (1, 0)$, n_1 has a geometric distribution, and conditional on $X_1 = (0, 1)$, n_2 has a geometric distribution. Thus, $E[n_1|X_1 = (1, 0)] = \lambda_1/\mu$, $\text{Var}[n_1|X_1 = (1, 0)] = \lambda_1(\lambda_1 + \mu)/\mu^2$, $E[n_2|X_1 = (0, 1)] = \lambda_2/\mu$, and $\text{Var}[n_2|X_1 = (0, 1)] = \lambda_2(\lambda_2 + \mu)/\mu^2$. Thus,

$$E[G] = \frac{\lambda_1^2 + \lambda_2^2}{(\lambda_1 + \lambda_2)\mu^2}$$

and

$$E[H] = \frac{2(\lambda_1^2 + \lambda_2^2 + (\lambda_1 + \lambda_2)\mu) + \mu^2}{2(\lambda_1 + \lambda_2)\mu^2}.$$

Therefore, assuming that $\lambda_2 \ll \lambda_1 \ll \mu = 1$, we have that $r = \Omega(\lambda_1^2)$ and

$$CV^2(r, m) = \frac{\text{Var}[G - rH]}{mr^2(E[H])^2} = \frac{1}{m}\Omega\left(\frac{1}{\lambda_1}\right).$$

Straightforward calculations show that the sensitivities $r_{\lambda_1} = \Omega(\lambda_1^2)$ and $r_{\lambda_2} = \Omega(\lambda_1^2 \lambda_2 + \lambda_2^2)$. Using the asymptotic variance from the Central Limit Theorem for gradient estimators, we arrive at the asymptotic squared coefficient of variation of our sensitivity estimates:

$$CV^2(r_{\lambda_1}, m) = \frac{1}{m} \Omega \left(\frac{1}{\lambda_1} \right)$$

and

$$CV^2(r_{\lambda_2}, m) = \frac{1}{m} \Omega \left(\frac{\lambda_1^3}{\lambda_1^6 + \lambda_1^4 \lambda_2 + \lambda_1^2 \lambda_2^2} \right).$$

We only used the terms involving the variances of the down time in a cycle and its gradient in these calculations since, when $\lambda_2 \ll \lambda_1 \ll \mu = 1$, these turn out to be the dominant terms as they were in the three-state example.

Now observe that the sensitivity with respect to λ_1 is much larger in magnitude than the sensitivity with respect to λ_2 when $\lambda_2 \ll \lambda_1 \ll \mu = 1$. The relative accuracy of the estimate of the sensitivity with the larger magnitude is of the same order as that of the performance measure estimate. Furthermore, the relative accuracy of the estimate of the sensitivity with the larger magnitude is much better than that of the estimate of the sensitivity with the smaller magnitude.

Thus, in these two examples the relative accuracies of the estimate of the sensitivity with the larger magnitude is of the same order as the relative accuracy of the performance measure estimate. Although these results were derived for simple examples, we see that this is also true for the models used in experimentation in Section 5

4 Implementation Issues

In developing the SAVE package, various implementation techniques were used to generate more quickly the transitions of the simulated Markov chain. For a discussion of the methods, the reader is referred to Goyal et al. (1992).

In addition, the variance reduction schemes described in the previous sections have been implemented in the SAVE package so that large models can be analyzed efficiently. The importance sampling methods for the embedded Markov chain use the following heuristics. (The ensuing exposition is based directly on material in Goyal et al. 1992.) We need to move the system very quickly to the set of failed states F , and once F is entered, the importance sampling should be turned off so that the system quickly returns to state 0, the “all components operational” state. We achieve this by increasing the probability of failure transitions over repair transitions. Lewis and Böhm (1984) call this “failure biasing.” We assign a combined probability *bias1* to the failure transitions in all the states where both failure and repair transitions are feasible. Individual failure and repair transitions are selected in the ratio of their rates given that a failure or a repair is selected, respectively. We call this the *Bias1/Ratio* method, or simply the *Bias1* method. We have found two other methods useful for selecting individual failure transitions, given that a failure has occurred. The first is to use a uniform distribution on the failure transitions, which has very good performance for “unbalanced” systems, as shown in Section 5 and in Goyal et al. (1992). We call

this the *Bias1/Balancing* method. The second is to give higher combined probability, *bias2*, to those failure transitions which correspond to component types which have at least one component of their type already failed. This exhausts the redundancy quickly and has much better performance for “balanced” systems, as shown in Section 5 and in Goyal et al. (1992). We call this the *Bias1/Bias2* method.

In the SAVE package the user can estimate via simulation the derivative of any of the performance measures with respect to any continuous parameter of the model. From (5), (22), (23), and (24), we see that to construct our derivative estimators, we must compute the derivatives $q'(\theta, \cdot)$, $P'(\theta, \cdot, \cdot)$, and $f'(\theta, \cdot)$. Because of the typically huge state spaces of the reliability systems considered, SAVE does not compute and store all of these values prior to running the simulation. Instead, SAVE calculates these quantities as needed during the simulation. In its most general form, SAVE analytically computes the values using a symbolic differentiator. This allows the user, for example, to define the component failure and repair rates to be complicated (differentiable) functions of a parameter θ , and then estimate derivatives of performance measures with respect to θ . However, using the symbolic differentiator is somewhat slow, with the extra CPU time needed to estimate each additional derivative being about the same as the time spent just estimating the (non-derivative) performance measure itself. Therefore, we employed special techniques in the implementation of the SAVE package to allow the user to estimate derivatives in certain cases with little extra computational effort. We now describe the method.

Suppose the user wishes to estimate the derivative with respect to the failure rate λ_i of component type i . Consider the derivative $q'(\theta, x)$, which we must compute during the simulation whenever the system is in state x . Recall that $q(\theta, x) = \sum_{y \neq x} q(\theta, x, y)$, and so we can express $q'(\theta, x)$ in terms of the $q'(\theta, x, y)$. For now, suppose that there is no failure propagation. If the transition (x, y) corresponds to the failure of a component of type i , then $q(\theta, x, y) = n_i(x)\lambda_i$, where $n_i(x)$ is the number of components of type i that are operational in state x ; see (25) in the Appendix. We easily compute the derivative of $q(\theta, x, y)$ as $q'(\theta, x, y) = n_i$. If the transition (x, y) does not correspond to the failure of a component of type i , then $q(\theta, x, y)$ is not a function of λ_i , and so $q'(\theta, x, y) = 0$. (A similar situation occurs when there is failure propagation.) The other derivatives that we must compute during the simulation also have simple expressions when differentiating with respect to λ_i . Thus, when estimating derivatives with respect to component failure rates, we are able to compute the derivatives needed in the simulation without resorting to the symbolic differentiator. The same is true when we estimate derivatives with respect to component repair rates. By avoiding the symbolic differentiator in this approach, SAVE can estimate derivatives with respect to component failure and repair rates with little additional CPU time, as we shall see in Section 5.2.3.

5 Experimental Results

In this section, we discuss the results of simulations of two different models to analyze the behavior of gradient estimates via the likelihood ratio method and to demonstrate the effectiveness of different

variance reduction techniques. We compare the sensitivity estimates to the performance measure estimates in many cases to show empirically that certain partial derivatives can be estimated as accurately as the performance measure itself. All numerical (non-simulation) and simulation results were obtained using the SAVE package (see Goyal and Lavenberg 1987). Because all of the models considered were analytically tractable, we were able to compute either exact values (up to machine precision) or tight bounds for the performance measures with the numerical method; see Muntz, de Souza e Silva, and Goyal (1989). We estimated all quantities using two different simulation methods: standard simulation (i.e., without using importance sampling) and importance sampling. In the case of transient performance measures, we did not employ any variance reduction technique for standard simulation. For the steady-state performance measures, we used conditional Monte Carlo (see Section 3.2) for both standard simulation and importance sampling.

5.1 An n -component Parallel System

The purpose of the first set of experiments is to examine the effect of the length of the regenerative cycle on the variability of likelihood ratio gradient estimators of steady-state performance measures. We will discuss the results of some simulations of an n -component parallel system, where the number of components n was varied from 2 to 12. For the system to be operational, there must be at least one functioning component. The repair rate μ was fixed at 1.0 for all values of n , and the values of the failure rate λ were varied so that the actual value of the steady-state unavailability remained fixed at 0.001. For each value of n , we simulated for 1,024,000 events, where an event is defined as any transition of the continuous-time Markov chain. From these runs, we formed estimates of the steady-state unavailability and the sensitivities of the steady-state unavailability with respect to λ . Table I contains the values for the steady-state unavailability and its sensitivity with respect to the failure rate λ obtained using the numerical method and standard simulation. Because of the small state spaces of these models, we were able to compute the values exactly (to within machine precision) using the numerical method of SAVE. All of the simulation estimates are given with the percentage relative half-widths of their 90% confidence intervals, which is defined to be 100% times the confidence interval half-width divided by the point estimate.

It is interesting to note that for small values of n , the relative size of the confidence intervals of the sensitivity estimates are close to the relative size of the confidence intervals of the estimates of the steady-state unavailability. However, as the number of components in the system increases, the relative accuracy of the sensitivity estimates degrades slightly as compared to that of the performance measure estimates. The reason for this is that the number of events per regenerative cycle is increasing as the number of components in the system grows since we have adjusted the failure rate so that the value of the steady-state unavailability remains constant. Since the derivative of the likelihood ratio turns out to be a sum of random variables where the number of summands equals the number of events in the regenerative cycle, as the regenerative cycles become longer we are summing up more random variables. This in turn leads to more variability.

Table I: Results for an n -component system with relative half-widths of the 90% confidence intervals for the simulation estimates.

No. of Comps. (n)	Failure Rate (λ)	Avg. No. of Events per Regen. Cycle	Steady-State Unavailability		Sensitivity w.r.t. λ	
			Numerical Result	Standard Simulation	Numerical Result	Standard Simulation
2	0.023	2.05	0.100×10^{-2}	0.099×10^{-2} $\pm 1.6\%$	0.195×10^{-2}	0.193×10^{-2} $\pm 1.6\%$
4	0.088	2.63	0.100×10^{-2}	0.101×10^{-2} $\pm 4.4\%$	0.359×10^{-2}	0.356×10^{-2} $\pm 4.7\%$
6	0.123	4.12	0.100×10^{-2}	0.100×10^{-2} $\pm 6.5\%$	0.490×10^{-2}	0.487×10^{-2} $\pm 7.3\%$
8	0.138	7.53	0.100×10^{-2}	0.101×10^{-2} $\pm 6.5\%$	0.586×10^{-2}	0.568×10^{-2} $\pm 9.0\%$
10	0.143	16.39	0.100×10^{-2}	0.101×10^{-2} $\pm 8.1\%$	0.646×10^{-2}	0.623×10^{-2} $\pm 10.4\%$
12	0.144	43.89	0.100×10^{-2}	0.102×10^{-2} $\pm 8.3\%$	0.676×10^{-2}	0.678×10^{-2} $\pm 13.1\%$

5.2 Balanced and Unbalanced Systems

The next model is a large computing system whose block diagram is shown in Figure 1. The model is also considered in Goyal et al. (1992) and Muntz, de Souza e Silva, and Goyal (1989). The following description of the model is based directly on material in Goyal et al. (1992). We use two different parameter sets to create a “balanced” and an “unbalanced” system. For a system to be considered balanced it must satisfy two criteria. First, each type of component has the same amount of redundancy, (i.e., the same number of components of a type must fail for the system to become nonoperational; e.g., 1-out-of-2 of a type has the same redundancy as 3-out-of-4 of another type). Also, the failure rates of all of the components must be of the same order of magnitude. A system that is not balanced is called unbalanced.

For a balanced system we select two sets of processors with two processors per set, two sets of controllers with two controllers per set, and six clusters of disks, each consisting of four 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 inaccessible without losing access to the data. The connectivity of the system is shown in Figure 1. We assume that when a processor of a given type fails, it has a 0.01 probability of causing the operating processor of the other type to fail. Each unit in the system has two failure modes which occur with equal probability. The failure rates of the processors, controllers, and disks are assumed to be 1/2000, 1/2000, 1/6000 per hour, respectively. The repair rates for all mode 1 and all mode 2 failures are 1 per hour and 1/2 per hour, respectively. Components are repaired by a single repairman who chooses components at random from the set of failed units. 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 six disk cluster are operational. We also assume

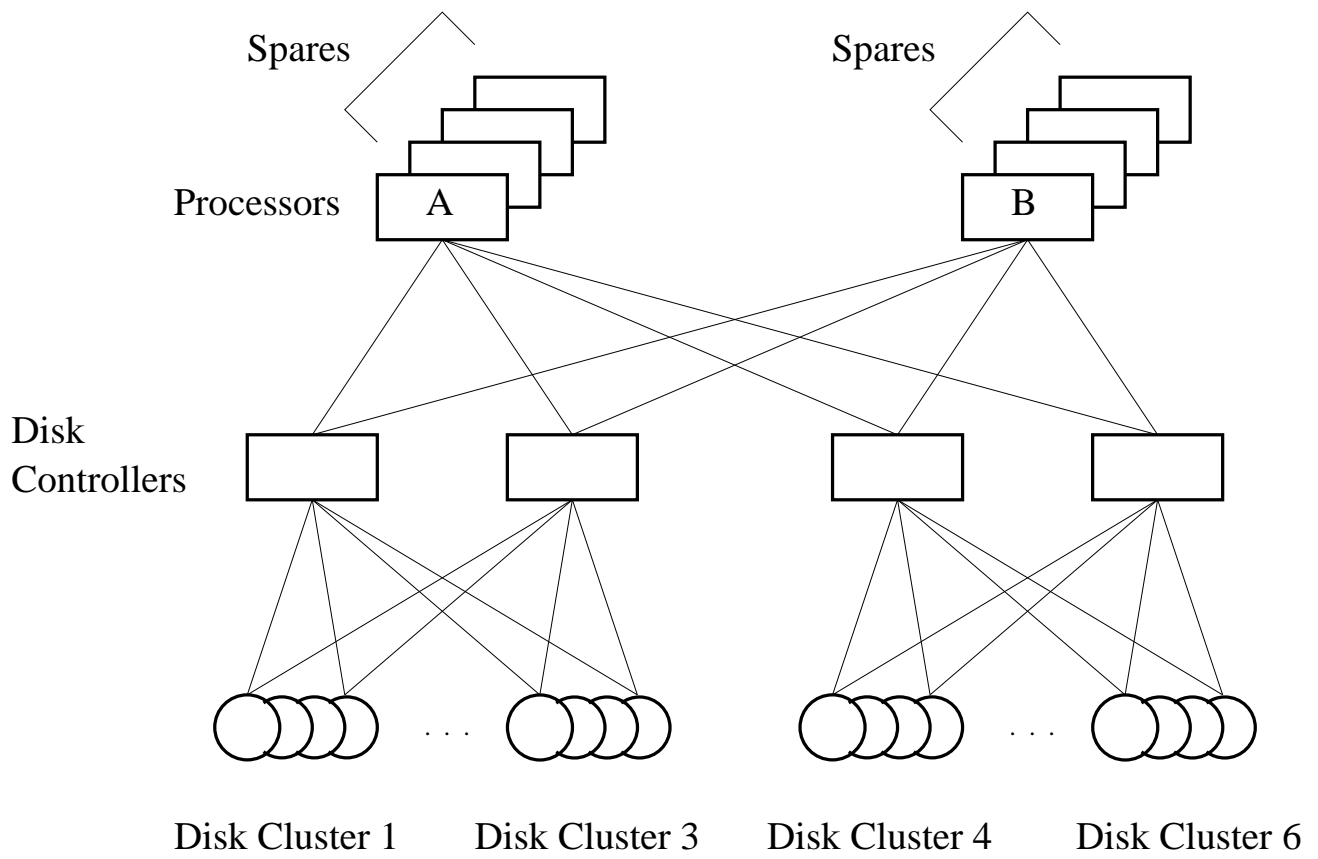


Figure 1: A block diagram of a reliability system.

Table II: Estimates of performance measures and their sensitivities for the balanced system with relative half-widths of the 90% confidence intervals for the simulation estimates.

Performance Measure	Estimate		
	Numerical Result	Standard Simulation	Bias 1/Bias 2 Method
Unavailability	0.931×10^{-5}	1.017×10^{-5} $\pm 27.1\%$	0.940×10^{-5} $\pm 2.7\%$
Sens. of Unavail. wrt <i>rr2</i>	-0.125×10^{-4}	-0.126×10^{-4} $\pm 33.0\%$	-0.127×10^{-4} $\pm 3.3\%$
Sens. of Unavail. wrt <i>c1fr</i>	0.232×10^{-5}	0.372×10^{-5} $\pm 64.1\%$	0.259×10^{-5} $\pm 6.3\%$
MTTF	$0.164 \times 10^{+6}$	$0.152 \times 10^{+6}$ $\pm 25.7\%$	$0.163 \times 10^{+6}$ $\pm 2.5\%$
Sens. of MTTF wrt <i>rr2</i>	$0.110 \times 10^{+6}$	$0.088 \times 10^{+6}$ $\pm 33.4\%$	$0.111 \times 10^{+6}$ $\pm 2.6\%$
Sens. of MTTF wrt <i>c1fr</i>	$-0.407 \times 10^{+5}$	$-0.589 \times 10^{+5}$ $\pm 58.6\%$	$-0.442 \times 10^{+5}$ $\pm 6.0\%$

that operational components continue to fail at the given rates when the system is failed.

We make minor changes to the above parameter settings to create an unbalanced system. We increase the number of processors of each type to 4, and double each processor’s failure rate to 1/1000 per hour. We decrease the failure rates of all other components by a factor of ten. In this system, although a processor failure is more likely to occur in a failure transition, it is less likely to cause a system failure due to the high processor redundancy. This is typical behavior for an unbalanced system.

5.2.1 Steady-State Measures

In this section we discuss the results of our experiments for estimating the steady-state unavailability and the mean time to failure and their sensitivities with respect to the parameters *rr2* (failure mode 2 repair rate) and *c1fr* (disk controller 1 failure rate). These two parameters were selected to demonstrate that we can estimate the sensitivities with the largest magnitude with about the same relative accuracy as the performance measure estimates, and the sensitivities of smaller magnitude are not estimated as precisely, as shown by the example in Section 3.4.2. We simulated both the balanced and the unbalanced systems. Tables II and III show the numerical and simulation results obtained for the balanced and the unbalanced systems, respectively.

The numerical results for the (non-derivative) performance measures are taken from Goyal et al. (1992). Since the balanced system has a few hundred thousand states and the unbalanced system has close to a million states, only bounds could be computed (see Muntz, de Souza e Silva, and Goyal 1989). These bounds are very tight and typically do not differ from the exact results significantly.

We see that the various variance reduction techniques have the same effect on the sensitivity estimates as they do on the performance measure estimates. Significant variance reductions can

Table III: Estimates of performance measures and their sensitivities for the unbalanced system with relative half-widths of the 90% confidence intervals for the simulation estimates.

Performance Measure	Performance Measure Estimate		
	Numerical Result	Standard Simulation	Bias 1/Balancing Method
Unavailability	0.697×10^{-7}	0.417×10^{-7} $\pm 164.5\%$	0.698×10^{-7} $\pm 2.4\%$
Sens. of Unavail. wrt $rr2$	-0.944×10^{-7}	-0.794×10^{-7} $\pm 164.5\%$	-0.938×10^{-7} $\pm 3.1\%$
Sens. of Unavail. wrt $c1fr$	0.232×10^{-7}	-0.343×10^{-10} $\pm 166.8\%$	0.236×10^{-7} $\pm 6.0\%$
MTTF	$0.219 \times 10^{+8}$	$0.470 \times 10^{+8}$ $\pm 164.5\%$	$0.218 \times 10^{+8}$ $\pm 2.3\%$
Sens. of MTTF wrt $rr2$	$0.148 \times 10^{+8}$	$0.456 \times 10^{+8}$ $\pm 164.5\%$	$0.147 \times 10^{+8}$ $\pm 2.4\%$
Sens. of MTTF wrt $c1fr$	$-0.730 \times 10^{+7}$	$0.110 \times 10^{+5}$ $\pm 191.0\%$	$-0.732 \times 10^{+7}$ $\pm 5.2\%$

be obtained using the *Bias1/Bias2* method for the balanced systems and *Bias1/Balancing* method for the unbalanced systems, as is shown in Goyal et al. (1992). These results hold for both the performance measure estimates and the sensitivity estimates.

We ran the simulations long enough so that the smallest entry in the tables for the percentage relative half-widths of the 90% confidence intervals was less than 5%. This corresponds to approximately 100,000 events for each entry in Table II and 1,000,000 events for each entry in Table III, respectively. Based on empirical results obtained in Goyal, Heidelberger, and Shahabuddin (1987), the values for $bias1 = 0.5$ and $bias2 = 0.5$ were selected for DIS.

There are several important points to note in the tables. For the balanced system, we used the *Bias1/Bias2* method, and *Bias1/Balancing* was used for the unbalanced system. As is shown in Goyal, Heidelberger, and Shahabuddin (1987), these methods are most effective for their respective models when estimating the (non-derivative) performance measures. We can see that this is also the case for the sensitivities since we obtain estimates of the largest sensitivities that are about as accurate as the performance measure estimate.

The relative precision of the performance measure estimates and the estimates of each of their respective sensitivities with respect to $rr2$ are approximately equal, which agrees with the analytic results we obtained from the simple examples in Section 3.4.1. Also, as claimed in Section 3.4.2, we do not obtain the same accuracy for the estimate of the sensitivity with respect to $c1fr$ since it is of smaller magnitude. It is also interesting to note that the amount of variance reduction from importance sampling over standard simulation in the sensitivity estimates is about the same as the variance reduction in the performance measure estimates. This is because the same likelihood ratio needed for importance sampling is used in both the performance measure estimate and the sensitivity estimates, and the likelihood ratio in both cases is multiplied by the accumulators at the end of each cycle.

Also note that the sensitivity estimates with respect to $c1fr$ in the unbalanced system using standard simulation given in Table III are very poor. This is because the value of $c1fr$ is much smaller than the value of parameter $procf_r$, the processors' failure rate, and so events corresponding to failures of disk controller 1 are somewhat rare compared to failures of one of the processors. Therefore, we are not able to obtain accurate results for both the point estimate of the sensitivity with respect to $c1fr$ and its variance. However, when using *Bias1/Balancing*, we obtain much better estimates of these quantities.

We next performed coverage experiments (e.g., see Lavenberg, Moeller, and Sauer 1979) to determine the validity of the confidence intervals that are formed based on the asymptotic central limit theorems described in Section 3. Such studies are important since certain variance reduction techniques sometimes do not produce valid confidence intervals, except for very long run-lengths (e.g., see Lavenberg, Moeller, and Sauer 1979).

We performed experiments on the estimates of the steady-state unavailability, U , and its sensitivities with respect to both $rr2$ and $c1fr$, denoted by U_{rr2} and U_{c1fr} , respectively, in the above described balanced system, as follows. We chose three run lengths corresponding to small, medium and large sample sizes, and we considered two ways of estimating U and its sensitivities: standard simulation and the *Bias1/Bias2* method with DIS. For each method and run length we ran $R = 100$ replications and formed point estimates $\hat{U}_1, \dots, \hat{U}_R$ of the performance measure estimate and $\hat{U}_{\theta,1}, \dots, \hat{U}_{\theta,R}$, for $\theta = rr2$ and $c1fr$, of the sensitivity estimates, and 90% confidence intervals for all of these estimates. We then calculated the mean percent relative bias ($= 100\% \cdot (1/R) \sum_{i=1}^R (\hat{U}_i - U)/U$ for the steady-state unavailability estimator, and likewise for the sensitivity estimators) and the standard deviation of this mean. Note that if an estimate is unbiased, then its mean percent relative bias should converge to zero as $R \rightarrow \infty$. We also calculated the 90% coverages, which is the percentage of the (computed) 90% confidence intervals that actually contain the true values of U , U_{rr2} , and U_{c1fr} , respectively. If the confidence interval is valid, then by definition, the 90% coverage should be equal to 90%.

We also calculated the mean percent relative half width of the 90% confidence intervals. For each replication, this relative value is computed using the point estimate and not the true value. Thus, since the mean percent relative half width of the 90% confidence intervals is defined as the ratio of the absolute half width of the 90% confidence interval over the point estimate, the mean is computed only over the replications with a nonzero point estimate. (By excluding the zero point estimates, the results presented are, in some sense, overly optimistic since we only have computed the mean percent relative half width over all of the "good" observations.) The results are listed in Table IV. Note that, as also seen in Goyal et al. (1992), the estimates using standard simulation are significantly more biased than those using importance sampling, and their confidence intervals are about an order of magnitude wider. The values of the relative bias and relative half widths for the estimates of the sensitivities with respect to $rr2$ are about the same as those for the performance measure estimate, while these values for the estimates of the other sensitivity are generally worse. This agrees with the results given in Section 3.4.2. Furthermore, for the small

Table IV: Coverage results for performance-measure estimators and sensitivity estimators for the balanced system (100 replications).

Events per Replication	Standard Simulation			Bias1/Bias2 Method		
	Average Relative Bias (Std. Dev.) (%)	Average Relative Half-Width (%)	Coverage (%)	Average Relative Bias (Std. Dev.) (%)	Average Relative Half-Width (%)	Coverage (%)
(a) Steady-State Unavailability						
2,000	6.95 (12.82)	144.40	54	0.74 (1.20)	18.88	85
20,000	-3.94 (3.41)	65.47	90	0.39 (0.34)	5.99	82
200,000	1.29 (1.09)	19.60	96	0.05 (0.13)	1.90	90
(b) Sensitivity of Steady-State Unavailability wrt $rr2$						
2,000	14.03 (19.15)	155.96	46	0.37 (1.54)	23.26	84
20,000	-1.86 (4.96)	82.30	84	0.40 (0.43)	7.46	92
200,000	2.23 (1.45)	25.27	94	0.03 (0.16)	2.38	86
(c) Sensitivity of Steady-State Unavailability wrt $c1fr$						
2,000	20.73 (36.64)	432.36	11	3.73 (2.86)	46.45	90
20,000	-2.13 (10.66)	147.02	62	-0.18 (1.04)	14.86	82
200,000	5.25 (3.62)	56.79	86	0.10 (0.27)	4.68	95

run length, the coverage drops significantly below 90% when using standard simulation. Using our variance reduction technique, all the coverages are close to the nominal 90% value.

The good behavior of the regenerative-based steady-state gradient estimates described here can be expected to typically hold for the types of models generated by the SAVE package. Because the failure rates are usually orders of magnitude smaller than the repair rates, regenerative cycles tend to be short, with a typical cycle consisting of one failure transition and one repair transition. Thus, when using importance sampling, regenerative cycles typically consist of only a few failure and repair transition since we turn off failure biasing once a system failure occurs in a cycle.

5.2.2 Transient Measures

In this section we discuss the results of our experiments for estimating reliability and its sensitivity with respect to both $rr2$ and $c1fr$. To estimate these measures, we did not use conditional Monte Carlo in our simulations. Recall that for transient measures we not only want the system to move quickly towards the set of system failed states F , but also to get there before the observation period expires. For Markov chain simulations, these issues are (in some sense) orthogonal, since the holding

times that determine the hitting time are conditionally independent of the embedded DTMC that is biased towards hitting F . We therefore use the same technique as in the steady-state case to bias the embedded Markov chain towards the system failed set, in addition to another independent technique (e.g., forcing as discussed in Section 2.2) to reduce the variance due to holding times in the various states. The likelihood ratios corresponding to these two aspects of simulation are conditionally independent and can be formulated as in Section 2.2 and in Goyal et al. (1992). The goal of the simulation is to study the relative accuracies of the performance measure estimate versus estimates of its sensitivities and to compare the effects of the forcing technique on these quantities. We considered only the balanced system. For each measure, we allowed each method to run for 400,000 events. The results are given in Table V.

For all methods, we notice that the confidence intervals are smaller for some range of intermediate time periods and wider at the ends. Also, the three tables indicate that forcing is most effective for short time intervals. These characteristics are discussed in Goyal et al. (1992).

It is interesting to note that the relative accuracy of the sensitivity estimates with respect to $rr2$ are consistently slightly worse than that of the performance measure estimate, which strays from the result that we obtained for the steady-state measures. This is because we are working with transient measures, and so the likelihood ratio includes terms for the (random) holding times. Thus, when we compute the gradient of the likelihood ratio, we are including additional random variables corresponding to the holding times in the sum, thereby increasing variability (compare (5) and (24)). It is also interesting to note that the relative accuracy of the sensitivity estimates degrades compared to that of the performance measure estimate as the time horizon increases. This is because the length of each observation increases as the time horizon grows, thus increasing the number of random variables included in the sum for the gradient of the likelihood ratio, thereby increasing the variance. This is similar to the results from the n -component parallel system.

5.2.3 Timing Experiments

Table VI shows the results from some timing experiments which we performed to determine how much extra CPU time is required to estimate sensitivities. The experiments consisted of different simulation runs in which we varied the number of sensitivities estimated and recorded the amount of CPU time taken in each run. We performed all of the experiments on an IBM 3090 computer using the SAVE package, simulating the balanced system with the *bias1/bias2* (0.5/0.5) technique for 100,000 events. As one can see, there is a fairly large fixed cost in CPU time for estimating the first gradient, but the marginal cost in CPU time for estimating each additional gradient is small. It is interesting to note that the additional time required to estimate eight sensitivities is about the same as the amount of time needed to run SAVE when estimating no gradients.

6 Summary and Directions for Future Work

In this paper we have shown that the likelihood ratio gradient estimation technique can be an effective practical tool for computing parameter sensitivities in large Markovian models of highly

Table V: Estimates of the unreliability and its sensitivities for different time horizons t for the balanced system.

Time (t)	Numerical Result	Standard Simulation		Bias1/Bias2 Method	
		No Forcing	Forcing	No Forcing	Forcing
(a) Unreliability					
4	0.153×10^{-4}	0.103×10^{-4} $\pm 116.3\%$	0.148×10^{-4} $\pm 23.9\%$	0.158×10^{-4} $\pm 7.0\%$	0.152×10^{-4} $\pm 1.4\%$
16	0.873×10^{-4}	1.072×10^{-4} $\pm 37.7\%$	0.943×10^{-4} $\pm 22.8\%$	0.869×10^{-4} $\pm 3.3\%$	0.870×10^{-4} $\pm 1.3\%$
64	0.380×10^{-3}	0.355×10^{-3} $\pm 24.0\%$	0.342×10^{-3} $\pm 24.9\%$	0.380×10^{-3} $\pm 1.8\%$	0.384×10^{-3} $\pm 1.1\%$
256	0.155×10^{-2}	0.146×10^{-2} $\pm 16.8\%$	0.158×10^{-2} $\pm 19.9\%$	0.157×10^{-2} $\pm 1.5\%$	0.158×10^{-2} $\pm 1.0\%$
1,024	0.623×10^{-2}	0.560×10^{-2} $\pm 14.9\%$	0.625×10^{-2} $\pm 16.0\%$	0.628×10^{-2} $\pm 4.9\%$	0.623×10^{-2} $\pm 4.3\%$
(b) Sensitivity of Unreliability wrt $rr2$					
4	-0.435×10^{-5}	-0.178×10^{-5} $\pm 117.0\%$	-0.513×10^{-5} $\pm 39.8\%$	-0.411×10^{-5} $\pm 13.4\%$	-0.428×10^{-5} $\pm 2.9\%$
16	-0.489×10^{-4}	-0.639×10^{-4} $\pm 59.8\%$	-0.480×10^{-4} $\pm 48.0\%$	-0.485×10^{-4} $\pm 6.1\%$	-0.485×10^{-4} $\pm 2.5\%$
64	-0.246×10^{-3}	-0.208×10^{-3} $\pm 46.7\%$	-0.168×10^{-3} $\pm 44.6\%$	-0.242×10^{-3} $\pm 3.6\%$	-0.247×10^{-3} $\pm 2.3\%$
256	-0.103×10^{-2}	-0.093×10^{-2} $\pm 35.0\%$	-0.088×10^{-2} $\pm 42.4\%$	-0.104×10^{-2} $\pm 4.0\%$	-0.106×10^{-2} $\pm 2.7\%$
1,024	-0.416×10^{-2}	-0.520×10^{-2} $\pm 33.4\%$	-0.450×10^{-2} $\pm 40.8\%$	-0.422×10^{-2} $\pm 15.0\%$	-0.397×10^{-2} $\pm 12.6\%$
(c) Sensitivity of Unreliability wrt $c1fr$					
4	0.379×10^{-5}	-0.349×10^{-5} $\pm 116.7\%$	0.459×10^{-5} $\pm 59.2\%$	0.404×10^{-5} $\pm 19.5\%$	0.372×10^{-5} $\pm 4.3\%$
16	0.217×10^{-4}	-0.011×10^{-4} $\pm 40.6\%$	0.253×10^{-4} $\pm 61.1\%$	0.212×10^{-4} $\pm 9.3\%$	0.214×10^{-4} $\pm 3.8\%$
64	0.945×10^{-4}	0.805×10^{-4} $\pm 74.2\%$	1.097×10^{-4} $\pm 63.9\%$	0.930×10^{-4} $\pm 5.5\%$	0.937×10^{-4} $\pm 3.5\%$
256	0.385×10^{-3}	0.306×10^{-3} $\pm 57.4\%$	0.670×10^{-3} $\pm 44.6\%$	0.397×10^{-3} $\pm 6.2\%$	0.396×10^{-3} $\pm 4.2\%$
1,024	0.154×10^{-2}	0.142×10^{-2} $\pm 60.6\%$	0.214×10^{-2} $\pm 48.7\%$	0.125×10^{-2} $\pm 27.1\%$	0.138×10^{-2} $\pm 22.8\%$

Table VI: CPU times to estimate different numbers of sensitivities using simulation (100,000 events).

Number of Sensitivites Estimated	CPU Seconds Taken
0	11.38
1	17.18
2	17.92
4	19.00
8	22.60
16	28.74

dependable systems. In fact, both our analysis and our computational experience suggests that the gradient estimates considered here are not significantly noisier than the estimates of the performance measures themselves. In addition to discussing implementation issues that arise in calculating and computing such gradient estimators, we also show that the derivative and expectation interchange implicit in obtaining the validity of the estimators does in fact hold for a wide class of performance measures associated with finite-state continuous-time Markov chains.

A number of interesting research directions present themselves for future work:

- 1 Development of additional variance reduction techniques for the likelihood ratio gradient estimator;
- 2 An analytic proof, for the general Markovian model of a highly dependable system, that the variability of the gradient estimator is roughly of the same order as that of the performance measure itself (thereby extending the results of this paper beyond our current three and five state examples given in Section 3.4);
- 3 Extending the methods of this paper to non-Markovian models, in which the failure and repair times are no longer necessarily exponential. This will necessitate the development of efficient non-regenerative techniques for estimating steady-state gradients in a rare-event setting.
- 4 Fully examining the applicability of infinitesimal perturbation analysis techniques to our setting.

7 Appendix

Now we justify the interchange of derivative and expectation. To do so, we make the following assumptions:

- A1.** State space E is finite.
- A2.** $Q(\cdot)$ is continuously differentiable on Θ .

A3. $P(\theta_0)$ is irreducible.

A4. $\Gamma(\theta) = \{(x, y) : P(\theta, x, y) > 0\}$ is independent of θ , for $\theta \in \Theta$.

A5. T is a stopping time such that there exists some $z_0 > 0$ for which the moment generating function $M_\theta^{N(T)}(z)$ of $N(T)$ converges for all $z \in (-z_0, z_0)$ and all $\theta \in \Theta$.

A6. $Z(\theta)$ has one of the following forms:

- 1 $Z(\theta) = 1_S$, where S is some event determined by the process \mathbf{Y} up to time T , where T is some stopping time satisfying Assumption A5.
- 2 $Z(\theta) = \int_0^T f(\theta, Y_s) ds$, where T is some stopping time satisfying Assumption A5 and f is a real-valued function defined on (Θ, E) for which $f(\cdot, x)$ is continuously differentiable on Θ for each $x \in E$.

We now argue that Assumptions A1–A4 are typically satisfied by models of highly reliable systems. In addition we show that the performance measures considered in Section 1.3 satisfy Assumptions A5 and A6.

Recall that a system we are considering consists of a finite number, C , of components, where each component can be in a finite number of states (e.g., fully operational, degraded status, failed, or under repair). In the case when components can only be in one of two states (i.e., operating or failed) and the repair discipline is processor sharing (i.e., the repairperson simultaneously works on all failed components, where the amount of effort devoted to each component is proportional to its repair rate), the state space is given by

$$E = \{(x_1, \dots, x_C) : 0 \leq x_i \leq n_i \text{ for } 1 \leq i \leq C\},$$

where n_i is the total number of components of type i in the system and x_i represents the number of failed components of type i in state x . Hence, the resulting model has a finite state space, and so Assumption A1 is satisfied.

The parameter θ usually is a vector consisting of the components' failure and repair rates; i.e., $\theta = (\lambda_1, \dots, \lambda_C, \mu_1, \dots, \mu_C)$, where λ_i and μ_i are the failure and repair rates, respectively, of components of type i . When there is no failure propagation, the elements of the generator matrix $Q(\theta)$ typically have the form

$$q(\theta, x, y) = \begin{cases} n_i(x)\lambda_i & \text{if } y = x + e_i \text{ for some } i \\ c(x, y)\mu_j & \text{if } y = x - e_j \text{ for some } j \\ 0 & \text{otherwise} \end{cases}, \quad (25)$$

for $x \neq y$, where $n_i(x)$ is the number of components of type i that are operational in state x , $c(x, y) \geq 0$ is some constant, and e_i is the i -th unit vector in \mathfrak{R}^C . When we allow for failure propagation, the generator matrix has a similar form except that there are transitions (x, y) for which $q(\theta, x, y) > 0$ and $y \geq x$, $y \neq x$, and $y \neq x + e_i$. Hence, in either case, Assumption A2 usually is satisfied.

The transition matrix is irreducible for $\theta \in \Theta$, which can be seen as follows. Since all operational components continue to fail regardless of the state of the system, the state (n_1, \dots, n_C) (i.e., the state with all components failed) is accessible from any other state of the system. Also, since the repairperson is always busy if some component is failed (i.e., for all $x \neq 0$, there exists some i such that $y = x - e_i$ and $q(\theta, x, y) > 0$), all states (including state 0) are accessible from state (n_1, \dots, n_C) . Thus, the transition matrix is irreducible, and so Assumption A3 is satisfied.

By the structure of the generator matrix $Q(\theta)$ given in (25), we can see that $\Gamma(\theta) = \{(x, y) : P(\theta, x, y) > 0\}$ typically is independent of θ , for $\theta \in \Theta$. Thus, we can let $\Gamma = \Gamma(\theta)$, and so Assumption A4 is satisfied.

Many interesting examples of stopping times satisfy Assumption A5. In particular, T satisfies Assumption A5 if Assumptions A1–A4 are in force and any of the following holds:

- 1 $T = T_{N(t)+1}$, where t is deterministic; i.e., T is the time of the first transition after time t .
- 2 $T = t$, where $t > 0$ is deterministic.
- 3 $T = \alpha_A = \inf\{t > 0 : Y_{t-} \notin A, Y_t \in A\}$ for some set of states A ; i.e., T is the hitting time to the set of states A .

We now show that all of the performance measures described in Section 1.3 satisfy Assumption A6. First consider the transient measures (except for MTTF). All of these measures have functions $Z(\theta)$ that depend on a deterministic stopping time T ; i.e., $T = t$, where t is fixed. By item 2 above, stopping times of this form satisfy Assumption A5. Thus, the performance measures reliability and the distribution of interval availability have functions $Z(\theta)$ satisfying form 1 of Assumption A6. Also, the function associated with the expected interval availability satisfies form 2 of Assumption A6. Now consider the steady-state performance measures and the MTTF. Under Assumptions A1 and A3, our CTMC \mathbf{Y} is a regenerative process. Hence, these performance measures can be expressed as ratios of expectations of functions $Z(\theta)$. For these measures, the $Z(\theta)$ depend on stopping times T that are hitting times, which satisfy Assumption A5 by item 3 above. The numerator and denominator terms in the steady-state availability and the numerator in the MTTF ratio expressions have functions $Z(\theta)$ of the form 2 in Assumption A6. Also, the denominator in the ratio formula for the MTTF can easily be seen to be of form 1 of Assumption A6. Hence, all of the performance measures discussed in Section 1.3 satisfy Assumption A6.

Now we justify the interchange of the derivative and expectation.

Theorem 1 *If Assumptions A1 – A6 hold, then*

$$\frac{\partial}{\partial \theta} [E_{\theta_0}[Z(\theta)L(T, \theta, \theta_0)]]_{\theta=\theta_0} = E_{\theta_0}[Z'(\theta_0)] + E_{\theta_0}[Z(\theta_0)L'(T, \theta_0, \theta_0)],$$

where θ_0 lies in the interior of Θ .

Before we prove the theorem, we need to prove some preliminary results. The first lemma shows that Assumption A5 implies that the stopping time T has a moment generating function which converges in some neighborhood of 0.

Lemma 1 *Suppose T is a stopping time satisfying Assumption A5. Then there exists $z_0 > 0$ such that the moment generating function $\bar{M}_\theta(z)$ of T converges for all $z \in (-z_0, z_0)$ and all $\theta \in \Theta$.*

Proof. Note that

$$\bar{M}_\theta(z) = E_\theta \left[e^{zT} \right].$$

For $z \leq 0$, $\bar{M}_\theta(z) < \infty$ since $e^{zT} \leq 1$. So now assume that $z > 0$. Note that over the set $\{N(T) = n\}$, we have $T \leq \sum_{k=0}^n t_k$, and so

$$\begin{aligned} \bar{M}_\theta(z) &= \sum_{n=0}^{\infty} E_\theta \left[e^{zT} 1_{\{N(T)=n\}} \right] \\ &\leq \sum_{n=0}^{\infty} E_\theta \left[\exp \left\{ z \sum_{k=0}^n z t_k \right\} 1_{\{N(T)=n\}} \right] \\ &\leq \sum_{n=0}^{\infty} E_\theta^{1/2} \left[\exp \left\{ 2z \sum_{k=0}^n t_k \right\} \right] P_\theta^{1/2} \{N(T) = n\} \end{aligned}$$

by the Cauchy-Schwarz inequality.

We now show that

$$\hat{M}_\theta(n, z) \equiv E_\theta \left[\exp \left\{ 2z \sum_{k=0}^n t_k \right\} \right]$$

converges for all $z > 0$ sufficiently small. Define $q_* = \inf\{q(\theta, x) : \theta \in \Theta, x \in E\} > 0$. Then we have

$$\begin{aligned} \hat{M}_\theta(n, z) &= E_\theta \left[E_\theta \left[\exp \left\{ 2z \sum_{k=0}^n t_k \right\} \middle| \mathbf{X} \right] \right] = E_\theta \left[\prod_{k=0}^n E_\theta \left[e^{2z t_k} \middle| \mathbf{X} \right] \right] \\ &= E_\theta \left[\prod_{k=0}^n \left(\frac{q(\theta, X_k)}{q(\theta, X_k) - 2z} \right) \right] \leq \left(\frac{q_*}{q_* - 2z} \right)^{n+1} < \infty \end{aligned}$$

for all $z < q_*/2$.

By Assumption A5, there exists $\eta > 0$ such that $M_\theta(\eta) < \infty$ for all $\theta \in \Theta$, where $M_\theta(\cdot)$ is the moment generating function of $N(T)$. Thus,

$$\begin{aligned} P_\theta \{N(T) = n\} &\leq P_\theta \{N(T) \geq n\} = P_\theta \left\{ e^{\eta N(T)} \geq e^{\eta n} \right\} \\ &\leq e^{-\eta n} M_\theta(\eta) \end{aligned}$$

by Markov's inequality, and so

$$\begin{aligned} \bar{M}_\theta(z) &\leq \left(\frac{M_\theta(\eta) q_*}{q_* - 2z} \right)^{1/2} \sum_{n=0}^{\infty} \left(\frac{e^{-\eta} q_*}{q_* - 2z} \right)^{n/2} \\ &< \infty \end{aligned}$$

for all $z < (1 - e^{-\eta})q_*/2$, which completes the proof. ■

We now state another lemma.

Lemma 2 *If Assumptions A1–A6 hold, then*

$$E_\theta[Z(\theta)^k] < \infty$$

$$E_\theta[Z'(\theta)^k] < \infty$$

for all k and all $\theta \in \Theta$.

Proof. When $Z(\theta) = 1_S$, the result obviously holds. So now assume $Z(\theta) = \int_0^T f(\theta, Y_s) ds$. Then we have that $|Z(\theta)| \leq \|f\|T$ and $|Z'(\theta)| \leq \|f'\|T$, where

$$\begin{aligned} \|f\| &= \sup\{|f(\theta, x)| : \theta \in \Theta, x \in E\} \\ \|f'\| &= \sup\{|f'(\theta, x)| : \theta \in \Theta, x \in E\}. \end{aligned}$$

Since $f(\cdot, x)$ is continuously differentiable on Θ for each $x \in E$, we have $\|f\| < \infty$ and $\|f'\| < \infty$ by Assumption A1. Now Lemma 1 implies that $E_\theta[T^k] < \infty$ for all k and all $\theta \in \Theta$, from which the result easily follows. \blacksquare

Now we prove the validity of the interchange of the derivative and expectation operators.

Proof of Theorem 1. To justify the interchange, we will show that the difference quotients

$$h^{-1}[Z(\theta_0 + h)L(T, \theta_0 + h, \theta_0) - Z(\theta_0)]$$

are dominated by an integrable random variable. By the mean value theorem, we have that the difference quotient is equal to

$$Z'(\eta)L(T, \eta, \theta_0) + Z(\eta)L'(T, \eta, \theta_0),$$

for some $\eta \in (\theta_0, \theta_0 + h)$.

Define

$$\begin{aligned} \|q'\|_h &= \sup\{|q'(\theta, x)| : |\theta - \theta_0| \leq h, x \in E\} \\ \|q'/q\|_h &= \sup\{|q'(\theta, x)/q(\theta, x)| : |\theta - \theta_0| \leq h, x \in E\} \\ \|q/q(\theta_0)\|_h &= \sup\{|q(\theta, x)/q(\theta_0, x)| : |\theta - \theta_0| \leq h, x \in E\} \\ \|q - q(\theta_0)\|_h &= \sup\{|q(\theta, x) - q(\theta_0, x)| : |\theta - \theta_0| \leq h, x \in E\} \\ \|P'/P\|_h &= \sup\{|P'(\theta, x, y)/P(\theta, x, y)| : |\theta - \theta_0| \leq h, (x, y) \in \Gamma\} \\ \|P/P(\theta_0)\|_h &= \sup\{|P(\theta, x, y)/P(\theta_0, x, y)| : |\theta - \theta_0| \leq h, (x, y) \in \Gamma\}, \end{aligned}$$

where $\Gamma = \Gamma(\theta)$. By Assumptions A1, A2, and A4, all of these terms are finite.

From (3), we have that $L'(T, \eta, \theta_0)$ is equal to

$$\begin{aligned} &\left[\sum_{j=0}^{N(T)-1} \left\{ \frac{q'(\eta, X_j)}{q(\theta_0, X_j)} - q'(\eta, X_j)t_j + \frac{P'(\eta, X_j, X_{j+1})}{P(\theta_0, X_j, X_{j+1})} \right\} - q'(\eta, X_{N(T)})(T - T_{N(T)}) \right] \\ &\cdot \left[\prod_{k=0}^{N(T)} \frac{q(\eta, X_k)}{q(\theta_0, X_k)} \exp\{-(q(\eta, X_k) - q(\theta_0, X_k))t_k\} \frac{P(\eta, X_k, X_{k+1})}{P(\theta_0, X_k, X_{k+1})} \right] \\ &\cdot \exp\{-(q(\eta, X_{N(T)}) - q(\theta_0, X_{N(T)}))(T - T_{N(T)})\}. \end{aligned}$$

Now we can bound $|L'(T, \eta, \theta_0)|$ by

$$\begin{aligned} & \left[(N(T) + 1)(\|q'/q\|_h + \|P'/P\|_h) + \|q'\|_h \sum_{j=0}^{N(T)-1} t_j + \|q'\|_h (T - T_{N(T)}) \right] \\ & \cdot \|q/q(\theta_0)\|_h^{N(T)+1} \|P/P(\theta_0)\|_h^{N(T)} \exp \left\{ \|q - q(\theta_0)\|_h \sum_{k=0}^{N(T)-1} t_k \right\} \\ & \cdot \exp \{ \|q - q(\theta_0)\|_h (T - T_{N(T)}) \}, \end{aligned}$$

which in turn can be bounded by $\phi_1(h)\phi_2(h)$, where

$$\begin{aligned} \phi_1(h) &= N(T)(\|q'/q\|_h + \|P'/P\|_h) + \|q'\|_h T \\ \phi_2(h) &= \|q/q(\theta_0)\|_h^{N(T)} \|P/P(\theta_0)\|_h^{N(T)} \exp \{ \|q - q(\theta_0)\|_h T \}. \end{aligned} \tag{26}$$

Note that we can bound $|Z(\theta_0)L'(T, \eta, \theta_0)|$ by $\phi(h) \equiv |Z(\theta_0)|\phi_1(h)\phi_2(h)$. Lemma 1 and Assumption A5 imply that T and $N(T)$ have finite moments of all orders. By Assumption A2, $\|q/q(\theta_0)\|_h \rightarrow 1$, $\|P/P(\theta_0)\|_h \rightarrow 1$, and $\|q - q(\theta_0)\|_h \rightarrow 0$ as $h \downarrow 0$. Hence, by repeated applications of the Schwarz inequality and using Lemma 2, we have that $\phi(h_0)$ is integrable for some $h_0 > 0$ which is sufficiently small. Now noting that for $0 < h_1 < h_0$, we have $\phi(h_1) < \phi(h_0)$, and so we can use $\phi(h_0)$ as our dominating random variable for $Z(\theta_0)L'(T, \eta, \theta_0)$. Thus, we have shown that $Z(\theta_0)L'(T, \eta, \theta_0)$ is integrable. Similarly, we can show that $Z'(\eta)L(T, \eta, \theta_0)$ can also be dominated by an integrable random variable. Hence, by noting that $L(T, \eta, \theta_0) \rightarrow 1$ as $h \downarrow 0$, the proof is complete. \blacksquare

Now we will give a proof of the asymptotic variance of our estimator of $r'(\theta_0)$ given in (10). To do this, we need the following result (see Serfling 1980, p. 118, for the proof).

Theorem 2 (Central Limit Theorem) *Let $\mathbf{X}_i, i = 1, 2, \dots$, be independent and identically distributed d -dimensional random vectors with mean vector μ and covariance matrix Σ , and suppose $g: \mathbb{R}^d \rightarrow \mathbb{R}$ is differentiable at μ . If $E\|\mathbf{X}_1\|^2 < \infty$, then*

$$\sqrt{n} [g(\bar{\mathbf{X}}_n) - g(\mu)] \Rightarrow N(0, \sigma^2)$$

as $n \rightarrow \infty$, where

$$\begin{aligned} \bar{\mathbf{X}}_n &= \frac{1}{n} \sum_{i=1}^n \mathbf{X}_i \\ \sigma^2 &= \nabla g(\mu)^T \Sigma \nabla g(\mu) \end{aligned}$$

and $\nabla g(\cdot)$ is the gradient of g .

Now we give a proof of the expression for the asymptotic variance of the estimate of the gradient when using the ratio formula.

Theorem 3 *If Assumptions A1 – A6 hold, then the asymptotic variance of the estimate of $r'(\theta_0)$ in (10) is given by (11).*

Proof. We define the vector $\mathbf{V} = (A, B, C, D)$, where

$$\begin{aligned} A &= Z_T'(\theta) + Z(\theta)L'(T, \theta_0, \theta_0) \\ B &= T \\ C &= Z_T(\theta) \\ D &= TL'(T, \theta_0, \theta_0) \end{aligned}$$

as in (12)–(15). To apply Theorem 2, we first need to show that $E_{\theta_0}[\|\mathbf{V}\|^2] < \infty$. Lemmas 1 and 2 show that T , $Z_T(\theta)$, and $Z_T'(\theta)$ all have finite moments of all orders. Now note that $|L'(T, \theta_0, \theta_0)|$ is bounded by $\phi_1(h)$ for all $h \geq 0$, where $\phi_1(h)$ is defined in (26) in the proof of Theorem 1. Since $\phi_1(h)$ has a moment generating function which converges for all sufficiently small h , we have that $L'(T, \theta_0, \theta_0)$ also has finite moments of all orders. Hence, by repeated applications of the Schwarz inequality, we have that A , B , C , and D all have finite second moments, which implies $E_{\theta_0}[\|\mathbf{V}\|^2] < \infty$.

To apply Theorem 2, we define $g : \mathfrak{R}^4 \rightarrow \mathfrak{R}$ as

$$g(a, b, c, d) = \frac{ab - cd}{b^2}. \quad (27)$$

Since T is a hitting time, we have $\beta > 0$, and so g is differentiable at the point $(\alpha, \beta, \gamma, \delta)$. Thus, by computing the gradient of g and substituting the appropriate values into the expression for the variance given in Theorem 2, the proof is complete. \blacksquare

Now we show that we obtain better estimates of the gradients when we use conditional Monte Carlo. Before we prove the result, we make some definitions. Define the vector $\mathbf{V} = (A, B, C, D)$ as in the proof of Theorem 3, and we let $\mathbf{W} = E_{\theta_0}[\mathbf{V}|\mathbf{X}] = (\hat{A}, \hat{B}, \hat{C}, \hat{D})$, where

$$\begin{aligned} \hat{A} &= E_{\theta_0}[A|\mathbf{X}] = G'(\theta_0) + G(\theta_0)\hat{L}'(\tau_0, \theta_0, \theta_0) \\ \hat{B} &= E_{\theta_0}[B|\mathbf{X}] = H(\theta_0) \\ \hat{C} &= E_{\theta_0}[C|\mathbf{X}] = G(\theta_0) \\ \hat{D} &= E_{\theta_0}[D|\mathbf{X}] = H'(\theta_0) + H(\theta_0)\hat{L}'(\tau_0, \theta_0, \theta_0), \end{aligned}$$

where $G(\theta)$, $H(\theta)$, $G'(\theta)$, $H'(\theta)$, and $\hat{L}'(\tau_0, \theta_0, \theta_0)$ are defined in (17), (18), (22), (23), and (24), respectively. Let $\mu = E_{\theta_0}[\mathbf{V}] = E_{\theta_0}[\mathbf{W}]$. Then, by Theorem 2, we have that

$$\sqrt{n} [g(\bar{\mathbf{V}}_n) - g(\mu)] \Rightarrow N(0, \sigma_1^2)$$

as $n \rightarrow \infty$, and

$$\sqrt{n} [g(\bar{\mathbf{W}}_n) - g(\mu)] \Rightarrow N(0, \sigma_2^2)$$

as $n \rightarrow \infty$, where g is defined in (27), and

$$\bar{\mathbf{V}}_n = \frac{1}{n} \sum_{i=1}^n \mathbf{V}_i$$

$$\begin{aligned}\bar{\mathbf{W}}_n &= \frac{1}{n} \sum_{i=1}^n \mathbf{W}_i \\ \sigma_1^2 &= \text{Var}[\nabla g(\mu)^T (\mathbf{V} - \mu)] \\ \sigma_2^2 &= \text{Var}[\nabla g(\mu)^T (\mathbf{W} - \mu)].\end{aligned}$$

So we have that σ_1^2 and σ_2^2 are the variances of the gradient estimators when using the ratio formula obtained without and with conditional Monte Carlo, respectively. Then we have the following result.

Proposition 1 $\sigma_2^2 \leq \sigma_1^2$.

Proof. By noting that

$$\nabla g(\mu)^T (\mathbf{W} - \mu) = E_{\theta_0}[\nabla g(\mu)^T (\mathbf{V} - \mu) | \mathbf{X}],$$

we have the result by the principle of conditional Monte Carlo (see Fox and Glynn 1986). ■

8 Acknowledgments

We would like to credit Adrian Conway for his initial work on combining likelihood ratio gradient estimation with importance sampling for steady-state measures when using discrete-time conversion. Also, we would like to thank the associate editor and the referees for their helpful comments.

Marvin Nakayama performed part of this work while he was visiting IBM Research. He performed the rest of the work while he was a graduate student in the Operations Research Department at Stanford University. While at Stanford, he was supported by IBM contract 12480042 and NSF grant ECS 86-17905. Peter Glynn performed part of this work while he was a visiting scientist at IBM research. He also was supported by IBM contract 12480042 and U.S. Army Research Contract ARO DAAL03-88-K-0063.

9 References

- Billingsley, P. 1961. *Statistical Inference for Markov Processes*. University of Chicago Press, Illinois.
- Crane, M. A., and D. L. Iglehart. 1974. Simulating Stable Stochastic Systems, II: Markov Chains. *J. Assoc. Comput. Mach.* **21**, 114–123.
- Fox, B. L., and P. W. Glynn. 1986. Discrete Time Conversion for Simulating Semi-Markov Processes. *Oper. Res. Lett.* **5**, 191–196.
- Geist, R. M., and K. S. Trivedi. 1983. Ultra-High Reliability Prediction for Fault-Tolerant Computer Systems. *IEEE Trans. Comput.* **C-32**, 1118–1127.
- Glasserman, P. 1992. Derivative Estimates From Simulation of Continuous Time Markov Chains. *Oper. Res.* **40**, 292–308.

- Glynn, P. W. 1986. Stochastic Approximation For Monte Carlo Optimization. *Proceedings of the 1986 Winter Simulation Conference*, Wilson, J., Henriksen, J. and Roberts, S. (eds.), IEEE Press, 356–364.
- Glynn, P. W. 1987. Likelihood Ratio Gradient Estimation: An Overview. *Proceedings of the 1987 Winter Simulation Conference*, Thesen, A., Grant, H. and Kelton, W. D. (eds.), IEEE Press, 366–375.
- Glynn, P. W. 1990. Likelihood Ratio Derivative Estimators for Stochastic Systems. *Comm. ACM* **33**, 75–84.
- Glynn, P. W., and D. L. Iglehart. 1989. Importance Sampling for Stochastic Simulations. *Management Sci.* **35**, 1367–1393.
- Goyal, A., P. Heidelberger, and P. Shahabuddin. 1987. Measure Specific Dynamic Importance Sampling for Availability Simulations. *Proceedings of the 1987 Winter Simulation Conference*, Thesen, A., Grant, H. and Kelton, W. D. (eds.), IEEE Press, 351–357.
- Goyal, A., and S. S. Lavenberg. 1987. Modeling and Analysis of Computer System Availability. *IBM J. Res. Develop.* **31**, 651–664.
- Goyal, A., P. Shahabuddin, P., Heidelberger, V. F. Nicola, P. W. and Glynn. 1992. A Unified Framework for Simulating Markovian Models of Highly Dependable Systems. *IEEE Trans. Comput.* **C-41**, 36–51.
- Ho, Y. C. 1987. Performance Evaluation and Perturbation Analysis of Discrete Event Dynamical Systems, Perspectives and Open Problems. *IEEE Trans. Automat. Control* **AC-32**, 563–572.
- Hordijk, A., D. L. Iglehart, R. Schassberger. 1976. Discrete Time Methods For Simulating Continuous Time Markov Chains. *Adv. in Appl. Probab.* **8**, 772–788.
- Knuth, D. E. 1976. Big Omicron, Big Omega and Big Theta. *SIGACT News (ACM)* **8**, 18–24.
- Laprie, J. C. 1984. Dependability Evaluation of Software Systems in Operation. *IEEE Trans. Software Engrg.* **SE-10**, 701–713.
- Lavenberg, S. S., T. L. Moeller, and C. H. Sauer. 1979. Concomitant Control Variables Applied to the Regenerative Simulation of Queueing Systems. *Oper. Res.* **27**, 134–160.
- L'Ecuyer, P. 1990. A Unified View of the IPA, SF and LR Gradient Estimation Techniques. *Management Sci.* **36**, 1364–1383.
- L'Ecuyer, P., N. Giroux, and P. W. Glynn. 1989. Stochastic Optimization by Simulation: Some Experiments with a Simple Steady-State Queue. *Management Sci.*, to appear.
- Lewis, E. E., and F. Böhm. 1984. Monte Carlo Simulation of Markov Unreliability Models. *Nuclear Engineering and Design* **77**, 49–62.

- Muntz, R. R., E. de Souza e Silva, and A. Goyal. 1989. Bounding Availability of Repairable Computer Systems. *Proceedings of the ACM Sigmetrics and Performance 89*, Berkeley, California, ACM Press, 29–38.
- Reiman, M. I., and A. Weiss. 1989. Sensitivity Analysis for Simulations Via Likelihood Ratios. *Oper. Res.* **37**, 830–844.
- Rubinstein, R. Y. 1986. The Score Function Approach for Sensitivity Analysis of Computer Simulation Models. *Math. Comput. Simulation* **28**, 351–379.
- Rubinstein, R. Y. 1989. Sensitivity Analysis and Performance Extrapolation for Computer Simulation Models. *Oper. Res.* **37**, 72–81.
- Serfling, R. J. 1980. *Approximation Theorems for Mathematical Statistics*. John Wiley, New York.
- Suri, R. 1989. Perturbation Analysis: The State of the Art and Research Issues Explained via the GI/G/1 Queue. *Proceedings of the IEEE* **77**, 114–137.
- Sutor, R. S. 1989. A Guide to Programming in the Scratchpad II Interpreter. IBM T. J. Watson Research Center, Yorktown Heights, New York.