

DTIC FILE COPY

ARO 25839. 30-MA

②

LIKELIHOOD RATIO SENSITIVITY ANALYSIS FOR
MARKOVIAN MODELS OF HIGHLY DEPENDABLE SYSTEMS

by

Marvin K. Nakayama, Ambuj Goyal and Peter W. Glynn

AD-A227 992

TECHNICAL REPORT No. 54

January 1990

Prepared under the Auspices
of
U.S. Army Research Contract
DAAL-03-88-K-0063

DTIC
ELECTE
OCT 24 1990
S E D

Approved for public release: distribution unlimited.

Reproduction in whole or in part is permitted for any
purpose of the United States government.

DEPARTMENT OF OPERATIONS RESEARCH
STANFORD UNIVERSITY
STANFORD, CALIFORNIA

90 10 24 003

(A)

Likelihood Ratio Sensitivity Analysis for Markovian Models of Highly Dependable Systems

Marvin K. Nakayama¹
Department of Operations Research
Stanford University
Stanford, California 94305



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

Peter W. Glynn²
Department of Operations Research
Stanford University
Stanford, California 94305

Accession For	
NTIS GRA&I	☒
DTIC TAB	☒
Unannounced	☐
Justification	
By	
Distribution/	
Availability Codes	
Dist	Avail and/or Special
A-1	

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.

(KR) ←

¹Research supported by the IBM Corporation under SUR-SST Contract 12480042 and by the U.S. Army Research Office under contract number DAAL-03-88-K-0063.

²Research supported by the IBM Corporation under SUR-SST Contract 12480042, U.S. Army Research Office under contract number DAAL-03-88-K-0063 and by a Natural Sciences and Engineering Research Council of Canada grant.

1 Introduction

This paper discusses the application of the likelihood ratio gradient estimator to simulations of highly dependable systems. We believe that 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, [6], [7], [8], [20], [22], and [23], 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 6), 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, [7], [19], and [20]). 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 paper describes one of the few successful implementations of sophisticated variance reduction techniques within a widely distributed simulation software package, namely the System Availability Estimator (SAVE) (see [11] and [12]) developed within IBM. The variance reduction techniques that are described within this paper have been implemented so as to be invisible to the user.

3. 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.
4. This paper shows how "discrete-time conversion" can be applied to the steady-state likelihood ratio gradient estimator (see also [5] and [20]). 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.
5. 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 5 describes various ideas used within SAVE to improve the computational efficiency of the estimator.
6. 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 [20] and used in [5] 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 gradient estimators is guaranteed to give a variance reduction.

This paper is organized as follows. Section 2 describes the basic mathematical model that is simulated by SAVE. In Sections 3 and 4, respectively likelihood

ratio gradient estimation for transient and steady-state performance measures is discussed. Section 4.3 also discusses certain insights that were obtained by analytically analyzing the behavior of the likelihood ratio gradient estimator for a couple of (very) small models. In Section 5, implementation issues are discussed. Section 6 is devoted to a description and discussion of the experimental results obtained through extensive simulations of several large models having more than a million states. Section 7 discusses future research directions. The concluding Appendix contains most of the theoretical material alluded to in Item 6 above.

2 Problem Setting

In this section we briefly discuss the modeling problems being addressed by the SAVE package [10] and describe the basic mathematical model being simulated. We also describe various performance measures associated with the models we consider here.

2.1 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 [4]). 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 [17]). Component interactions often have a substantial effect on system availability and must therefore be considered in addition to the individual component behaviors. 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, the 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 computing 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 [12]). In this paper, we consider the gradient estimation problem for these measures. We use one change of measure to compute the gradient using the likelihood ratio gradient estimation technique, and we use another change of measure (importance sampling) to compute these gradients quickly.

2.2 Markovian Model

Suppose $Y = \{Y_s : s \geq 0\}$ is an irreducible, continuous time Markov chain with state space E and infinitesimal generator $Q(\theta) = \{q(\theta, i, j) : i, j \in E\}$, where θ is in some open 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 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.

Let $X = \{X_n : n \geq 0\}$ 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 $X_n \equiv (X_0, X_1, \dots, X_n)$. Recall X is a discrete time Markov chain (DTMC) with transition matrix $P(\theta)$ defined by $P(\theta, i, j) = q(\theta, i, j)/q(\theta, i)$ for $i \neq j$ and $P(\theta, i, i) = 0$, where $q(\theta, i) = -q(\theta, i, i)$. Furthermore, conditional on 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 Y , i.e. $T_0 = 0$, and $T_n = t_0 + t_1 + \dots + t_{n-1}$ for $n \geq 1$. Then define $N(t) = \sup\{n \geq 0 : T_n \leq t\}$.

Let T denote a stopping time satisfying assumption A5 in the appendix. Also, for any set of states A , 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 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.

2.3 Performance Measures

We will be interested in two types of dependability measures associated with the CTMC Y : transient measures and so-called steady-state measures. Considering the transient measures first, the interval availability, $A(t)$, is defined by

$$A(t) = \frac{1}{t} \int_0^t 1_{\{Y_s \in O\}} ds.$$

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 and 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 $(0, t)$:

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

For steady-state measures we assume that Y is irreducible, in which case $Y_s \Rightarrow Y'$ as $s \rightarrow \infty$, where $\Rightarrow$ denotes convergence in distribution and Y' is a rv having the steady-state distribution $\pi = \{\pi_i, i \in E\}$ (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. Also, we assume that when computing steady-state measures, we can express $Z(\theta)$ as $Z(\theta) = \lim_{t \rightarrow \infty} \frac{1}{t} \int_0^t f(\theta, Y_s) ds$, where $f(\theta, \cdot)$ is a real-valued function on E which satisfies assumptions A7 and A8 in the appendix. By regenerative process theory (see [2]), our steady-state measures take the form of a ratio of two expected values:

$$r = E_\theta[Z(\theta)] = \frac{E_\theta[\int_0^{\alpha_0} f(\theta, Y_s) ds]}{E_\theta[\alpha_0]}.$$

If $f(\theta, i) = 1_{\{i \in O\}}$, then $E_\theta[Z(\theta)]$ is the long run fraction of time the system is operational and is called the steady-state availability; which we denote by $A = \lim_{t \rightarrow \infty} E_\theta[A(t)]$. We will sometimes find it convenient to consider the expected 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 [12]. 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 4 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.

3 Estimating 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 (measurable) set of events.
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 satisfying assumptions A7 and A8 in the appendix.

We define the "likelihood" of a sample path under parameter θ as

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

and the likelihood ratio is given by

$$L(T, \theta, \theta_0) = d\mu(T, \theta) / d\mu(T, \theta_0), \quad (3.1)$$

where θ_0 is some fixed value of θ .

Our performance measure is given by

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

We call this transformation a "change of measure" since we are now computing the expectation based on a different parameter value. The validity of the change of measure is discussed in [1] and [12]. 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,

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

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

$$L'(T, \theta, \theta_n) = \left[\sum_{k=0}^{N(T)} \left\{ \frac{q'(\theta, X_k)}{q(\theta, X_k)} - q'(\theta, X_k) t_k + \frac{P'(\theta, X_k, X_{k+1})}{P(\theta, X_k, X_{k+1})} \right\} - q'(\theta, X_{N(T)+1})(T - T_{N(T)}) \right] L(T, \theta, \theta_n). \quad (3.2)$$

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 $\theta = \theta_n$. In this case, we have that since $L(T, \theta_n, \theta_n) = 1$,

$$r'(\theta_n) = E_{\theta_n} Z'(\theta_n) + E_{\theta_n} Z(\theta_n) L'(T, \theta_n, \theta_n) \quad (3.3)$$

and

$$L'(T, \theta_n, \theta_n) = \sum_{k=0}^{N(T)} \left[\frac{q'(\theta_n, X_k)}{q(\theta_n, X_k)} - q'(\theta_n, X_k) t_k + \frac{P'(\theta_n, X_k, X_{k+1})}{P(\theta_n, X_k, X_{k+1})} \right] - q'(\theta_n, X_{N(T)+1})(T - T_{N(T)}). \quad (3.4)$$

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 last exponential term drops out.

The stopping time, say T_1 , used in the likelihood ratio need not be the same as the stopping time, say T_2 , used for the function $Z(\theta_0)$. However, we always need to take $T_1 \geq T_2$, with strict inequality possible. For example, when computing reliability at a time t , note that $Z(\theta_0)$ is $\mathcal{F}_t$ -measurable, while we can use the likelihood ratio based on $\mathcal{F}_{T_{N(T)+1}}$, where $\mathcal{F}_t$ is the sigma-field generated by the process up to time t .

These results are similar to the results derived in In [20], only the specialization to Poisson processes is discussed. Also, the "amiability" assumption discussed in [5] and [20] holds in the current context, and examples of performance measures satisfying this condition are discussed in Section 2.3 and in the Appendix.

3.1 Importance Sampling to Reduce Variance

We can now apply another change of measure to implement the importance sampling to obtain

$$r'(\theta_0) = E_{\theta} \cdot Z'(\theta_0) L(T, \theta_0, \theta^*) + E_{\theta} \cdot Z(\theta_0) L'(T, \theta_0, \theta_0) L(T, \theta_0, \theta^*), \quad (3.5)$$

where θ^* is the parameter value used for the importance sampling change of measure. $L(T, \theta_0, \theta^*)$ is the same as Equation 3.1 except with (θ_0, θ^*) playing the role of (θ, θ_0) (as it should be). Note that we can use two different changes of measure, i.e. two different values of θ^* , to estimate the two expectations on the right hand side of Equation 3.5. Also, the value of θ^* used after the k^{th} transition can depend on 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 5 and also in [10].

We can actually separate the likelihood ratio into two different components, the first including only the transition probabilities of the embedded DTMC and the second incorporating only the random holding times, i.e.

$$L(T, \theta_0, \theta^*, \theta^{**}) = L_1(T, \theta_0, \theta^*) L_2(T, \theta_0, \theta^{**}),$$

where

$$L_1(T, \theta_0, \theta^*) = \prod_{k=0}^{N(T)} \frac{P(\theta_0, X_k, X_{k+1})}{P(\theta^*, X_k, X_{k+1})} \quad (3.6)$$

$$L_2(T, \theta_0, \theta^{**}) = \prod_{k=0}^{N(T)} \frac{q(\theta_0, X_k) \exp\{-q(\theta_0, X_k)t_k\}}{q(\theta^{**}, X_k) \exp\{-q(\theta^{**}, X_k)t_k\}} \cdot \frac{\exp\{-q(\theta_0, X_{N(T)+1})(T - T_{N(T)})\}}{\exp\{-q(\theta^{**}, X_{N(T)+1})(T - T_{N(T)})\}}. \quad (3.7)$$

Thus, we can apply different changes of measure to the two components of the likelihood ratio, which allows us to tailor each change of measure for a specific purpose.

Lewis and Böhm [15] 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 the estimation of 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+1} is forced to be between zero and $t - T_n$ by selecting t_{n+1} from the conditional density given by

$$h(t_{n+1}|X_n, t_n) = \frac{\lambda_0 e^{-\lambda_0 t_{n+1}}}{1 - e^{-\lambda_0(t - T_n)}},$$

where $0 \leq t_{n+1} \leq t - T_n$ 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 $T = \min(\alpha_F, N(t) + 1)$.

Note that $h(t_{n+1}|X_n, t_n)$ is not positive whenever the exponential density is, and so the standard theory of importance sampling says that this is not a legitimate change of measure. However, in this case $h(t_{n+1}|X_n, t_n)$ is positive over that part of the sample space ($\{\omega : \alpha_F < t\}$) that counts, which is sufficient (see [9]).

4 Estimating 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 satisfying assumptions A7 and A8 in the appendix. If we assume that the CTMC Y has finite state space E and the transition matrix P for the embedded DTMC is irreducible, then we have that $Y_t \Rightarrow Y^*$ as $t \rightarrow \infty$. In this case, we define our stopping time T to be the time of the first return to the initial state 0, i.e. $T = \sigma_0$. Let τ_0 be the first return time of the embedded DTMC to state 0. Since Y is a CTMC, T is a regeneration time. Hence, assuming that $E_\theta |Z(\theta)| < \infty$, we can express $r(\theta)$ using the ratio formula

$$r(\theta) = \frac{E_\theta Z_T(\theta)}{E_\theta T},$$

where

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

We also have that 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 $\hat{r}_m = \sum_{j=1}^m \hat{Z}_{T,j}(\theta) / \sum_{j=1}^m \hat{T}_j$, $\hat{Z}_{T,j}(\theta)$ and $\hat{T}_j$, $j = 1, 2, \dots$, denote independent, identically distributed observations of $Z_T(\theta)$ and T , respec-

tively, and $\sigma_r^2 = \text{Var}[Z_{T,j}(\theta) - rT_j]/E[T_j]^2$. See [2] for further details.

If one formally differentiates the expression for $r(\theta)$ by interchanging the derivative and expectation, one obtains

$$r'(\theta_0) = \frac{u'(\theta_0)l(\theta_0) - l'(\theta_0)u(\theta_0)}{l^2(\theta_0)}, \quad (4.1)$$

where

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

and

$$\begin{aligned} Z_T'(\theta_0) &= \sum_{k=0}^{r_0-1} f'(\theta_0, X_k) t_k \\ L'(T, \theta_0, \theta_0) &= \sum_{k=0}^{r_0-1} \left[\frac{q'(\theta_0, X_k)}{q(\theta_0, X_k)} - q'(\theta_0, X_k) t_k + \frac{P'(\theta_0, X_k, X_{k+1})}{P(\theta_0, X_k, X_{k+1})} \right]. \end{aligned}$$

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

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

$$\begin{aligned} & \frac{1}{\beta^2} \sigma_A^2 + \left[\frac{2\gamma\delta - \alpha\beta}{\beta^2} \right]^2 \sigma_B^2 + \frac{\delta^2}{\beta^4} \sigma_C^2 + \frac{\gamma^2}{\beta^4} \sigma_D^2 \\ & + 2 \left[\frac{2\gamma\delta - \alpha\beta}{\beta^2} \sigma_{AB} - \frac{\delta}{\beta^2} \sigma_{AC} - \frac{\gamma}{\beta^2} \sigma_{AD} \right. \\ & \left. - \frac{2\gamma\delta^2 - \alpha\beta\delta}{\beta^3} \sigma_{BC} - \frac{2\gamma^2\delta - \alpha\beta\gamma}{\beta^3} \sigma_{BD} + \frac{\gamma\delta}{\beta^3} \sigma_{CD} \right], \end{aligned} \quad (4.2)$$

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

$$A = Z_T'(\theta_0) + Z_T(\theta_0) l'(T, \theta_0, \theta_0) \quad (4.3)$$

$$B = T \quad (4.4)$$

$$C = Z_T(\theta_0) \quad (4.5)$$

$$D = TL'(T, \theta_0, \theta_0), \quad (4.6)$$

and

$$\alpha = E_{\theta_0} A$$

$$\beta = E_{\theta_0} B$$

$$\gamma = E_{\theta_0} C$$

$$\delta = E_{\theta_0} D.$$

A proof of the validity of the expression for the variance is given in [20]. However, we give a simpler and cleaner proof in the appendix.

4.1 Conditioning to Reduce Variance

Conditional Monte Carlo is a technique which can be used to reduce the variance in simulations of CTMCs (See [3] and [13]). By conditioning on the embedded DTMC 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. Also, as discussed in [3] and [13], 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 $r'(\theta)$.

Using conditional Monte Carlo, we obtain another ratio formula

$$r(\theta) = \frac{E_{\theta} E_{\theta} [Z_T(\theta) | X]}{E_{\theta} E_{\theta} [T | X]} = \frac{E_{\theta} [G(\theta)]}{E_{\theta} [H(\theta)]}, \quad (4.7)$$

where a straightforward calculation shows that

$$G(\theta) = \sum_{k=0}^{\tau_0-1} g(\theta, X_k) \quad (4.8)$$

$$H(\theta) = \sum_{k=0}^{\tau_0-1} h(\theta, X_k) \quad (4.9)$$

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

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

and τ_0 is the first return time of the DTMC to state 0.

In [6], it is shown that under certain conditions (viz., assumption A4 given in the appendix) that

$$r(\theta) = \frac{E_{\theta_0} G(\theta) \dot{L}(\tau_0, \theta, \theta_0)}{E_{\theta_0} H(\theta) \dot{L}(\tau_0, \theta, \theta_0)},$$

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

$$\dot{L}(\tau_0, \theta, \theta_0) = \prod_{k=0}^{\tau_0-1} \frac{P(\theta, X_k, X_{k+1})}{P(\theta_0, X_k, X_{k+1})}. \quad (4.10)$$

A simple calculation shows that $\dot{L}(\tau_0, \theta, \theta_0) = E_{\theta}[L(T_{\tau_0}, \theta, \theta_0)|\mathbf{X}]$. Note that $\dot{L}$ is the same as L_1 defined in Equation 3.6. If one formally differentiates this expression by interchanging the derivative and expectation, one obtains

$$r'(\theta_0) = \frac{\dot{u}'(\theta_0)\dot{l}(\theta_0) - \dot{l}'(\theta_0)\dot{u}(\theta_0)}{\dot{l}^2(\theta_0)}, \quad (4.11)$$

where

$$\dot{u}(\theta_0) = E_{\theta_0} G(\theta_0)$$

$$\dot{u}'(\theta_0) = E_{\theta_0} G'(\theta_0) + E_{\theta_0} G(\theta_0) \dot{L}'(\tau_0, \theta_0, \theta_0)$$

$$\dot{l}(\theta_0) = E_{\theta_0} H(\theta_0)$$

$$\dot{l}'(\theta_0) = E_{\theta_0} H'(\theta_0) + E_{\theta_0} H(\theta_0) \dot{L}'(\tau_0, \theta_0, \theta_0)$$

and

$$G'(\theta_0) = \sum_{k=0}^{\tau_0-1} \frac{f'(\theta_0, X_k)q(\theta_0, X_k) - q'(\theta_0, X_k)f(\theta_0, X_k)}{q^2(\theta_0, X_k)} \quad (4.12)$$

$$H'(\theta_0) = \sum_{k=0}^{\tau_0-1} -\frac{q'(\theta_0, X_k)}{q^2(\theta_0, X_k)} \quad (4.13)$$

$$\dot{L}'(\tau_0, \theta_0, \theta_0) = \sum_{k=0}^{\tau_0-1} \frac{P'(\theta_0, X_k, X_{k+1})}{P(\theta_0, X_k, X_{k+1})}. \quad (4.14)$$

The proof of the validity of the interchange of the operators in this case is given in [6].

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 Equation 4.2, is the asymptotic variance of the estimator of Equation 4.1, and σ_2^2 is the asymptotic variance of the estimator of Equation 4.11. 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 4 in the appendix).

4.2 Importance Sampling to Reduce Variance

As in Section 3.1, we can use importance sampling by applying another change of measure. However, in this case since 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 Equation 3.6, or equivalently, $\hat{L}$, given in Equation 4.10.

4.3 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 was also analyzed in [12], and the second is a two-dimensional five state birth and death process.

Because of their simple structure, we are able to do an extensive amount of analysis on these models. Recall that we defined the sensitivity of a performance measure with respect to a certain parameter to be the product of parameter itself multiplied by the gradient of the performance measure with respect to the parameter. Therefore, a sensitivity measures the effect of relative changes to parameter values on a performance measure, and so relative changes in the parameters corresponding to the largest sensitivities will cause the largest change in the performance measure. We will show that when one sensitivity is much larger in magnitude than another, its relative accuracy is much greater than that of the smaller sensitivity. In addition, we can estimate the sensitivities with the largest orders of magnitude with about the same relative accuracy as the regular 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. We measure the relative accuracy by the squared coefficient of variation. Much of the analysis was done using the symbolic manipulator Scratchpad [25].

We define the vector of parameters θ which we compute sensitivities with respect to as the vector of all continuous-valued parameters of the model. Note that the above claims 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, the claims seem to hold.

4.3.1 A Three State Example

The three state example can be viewed as a reliability system in which there is one type of component with redundancy of two and the components can fail and be 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.

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 can further assume that $\mu \approx 1$, since this only fixes the time scale.

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 Equation 4.7. In this example, we set $f(0) = f(1) = 0$ and $f(2) = 1$. We assume that state 0 is the regenerative state. The numerator in our ratio formula can be explicitly written as $G = n_F h_2$, and the denominator can be expressed as $H = h_0 + h_1 + n_F(h_1 + h_2)$, 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 [12], we have that $r = \Theta(\lambda^2)$ and

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

where $CV^2(r, m)$ denotes the asymptotic squared coefficient of variation of our estimate of r after m regenerative cycles (we modify Knuth's Θ -notation [14] to mean $f(x) = \Theta(g(x))$ 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 $r_\lambda = \Theta(\lambda^2)$ and $r_\mu = \Theta(\lambda^2)$, where we use the notation $r_\theta = \theta \cdot \partial r / \partial \theta$. Using the asymptotic variance from the Central Limit Theorem for gradient estimators from Section 4, we can arrive at the asymptotic squared coefficients of variation of our sensitivity estimates, which are given by $CV^2(r_\lambda, m) = \Theta(1/\lambda)/m$ and $CV^2(r_\mu, m) = \Theta(1/\lambda)/m$. 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 terms involving the variances of the down time in a cycle G and its gradients.

Thus, when $\lambda \ll \mu$, we have that both of the sensitivities are of the same order as the regular estimate, and the relative accuracies of the regular estimate and the sensitivities are about the same.

4.3.2 A Five State Example

The five state example models a system with two types of components, each of which has 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 assume that the regenerative state is the

state in which all components are operational. We consider the system to be operational in states $(0,0)$, $(1,0)$, and $(0,1)$, but failed in states $(2,0)$ and $(0,2)$. We 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:

$$P((0,0), (1,0)) = \lambda_1 / (\lambda_1 + \lambda_2)$$

$$P((1,0), (2,0)) = \lambda_1 / (\lambda_1 + \mu)$$

$$P((1,0), (0,0)) = \mu / (\lambda_1 + \mu)$$

$$P((0,0), (0,1)) = \lambda_2 / (\lambda_1 + \lambda_2)$$

$$P((0,1), (0,2)) = \lambda_2 / (\lambda_2 + \mu)$$

$$P((0,1), (0,0)) = \mu / (\lambda_2 + \mu)$$

$$P((2,0), (1,0)) = P((0,2), (0,1)) = 1.$$

We 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 the numerator can be explicitly written as

$$G = 1_{\{X_1=(1,0)\}} n_1 h_{(2,0)} + 1_{\{X_1=(0,1)\}} n_2 h_{(0,2)},$$

and the denominator can be expressed as

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

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 $r = \Theta(\lambda_1^2)$ and

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

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

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

and

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

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.

From this example, we see that when $\lambda_2 \ll \lambda_1 \ll \mu = 1$, the sensitivity with respect to λ_1 is much larger in magnitude than the sensitivity with respect to λ_2 , and the relative accuracy of the former is much better than that of the latter.

Thus, we see that in these two examples the relative accuracy of the sensitivity with the larger magnitude is of the same order of the relative accuracy of the regular estimate. Though these results were derived for simple examples, we can see that this is also true for the models used in experimentation in Section 6

5 Implementation Issues

In this section we consider the implementation of the different variance reduction techniques described in the previous sections. These techniques have been implemented in the SAVE package so that large models can be simulated. One salient feature of our implementation is that we use one simulation run for estimating all the measures and sensitivities. Regenerative simulation is used with the "all components operational" state as the regenerative state. The event generator simulates only the embedded Markov Chain. For the steady-state measures we accumulate functions of the mean holding times in the various states and also functions of the gradients of mean holding times and transition probabilities. For the transient measures we accumulate functions of the sample holding times (from exponential distributions) in the various states and also gradients of the transition probabilities and densities of the holding times. The likelihood ratio for transient measures is different for each of the different transient estimators allowing us to tailor it for the specific application.

We have employed various techniques in the implementation of the SAVE package which allow the CTMC to be generated quickly. We do not generate the entire state space explicitly since the models which we solve using SAVE can have millions of states. Instead, we use arrays to keep track of the number of operational and spare components of each type and the order of components in the repair queue. This information is sufficient to determine the state that the

system is in. In addition, arrays are used to store the total failure rate and total repair rate of each type of component. Also, we keep track of the sum of all of the failure rates of all operational and spare components and the sum of all the repair rates of components currently being repaired. Since each transition typically affects only one component (the program has been implemented to allow for the possibility of the statuses of more than one component changing on a single transition, e.g. one component failure causing others to fail also), we are able to update the arrays and the two sums quickly. By having these quantities readily available, we can easily determine the transition probabilities of the embedded DTMC and also the total rate out of a state.

The importance sampling for the embedded Markov chain is based on the following heuristics. As suggested in [12], 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. This has been called "failure biasing" in [15]. 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 6 and in [12]. We call this the *Bias1/Balancing* method. The second is to give higher combined probability, *bias2*, to those failure transitions which correspond to componen-

t 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 6 and in [12]. We call this the *Bias1/Bias2* method.

For the steady-state unavailability each regenerative cycle corresponds to a sample. We use either direct simulation or the DIS method given in Section 4.2 to estimate steady-state unavailability and its sensitivities. For the mean time to failure, a sample ends when either the regeneration occurs or the system enters one of the system failed states from the set F . In the latter case, we continue to simulate the embedded Markov chain until the regeneration occurs before starting a new sample. This wastes only a few events as typically a regenerative cycle is short (the number of events per cycle is approximately twice the average redundancy, which is typically two or three). Once again, we use either direct simulation or the DIS method to estimate the mean time to failure and its sensitivities. For the transient measures multiple regenerative cycles may be contained in a single sample. Moreover, a sample typically ends either when a failure occurs or when the time interval expires, which is usually in the middle of some regenerative cycle. As in the case for the mean time to failure, we continue to simulate the embedded Markov chain until the next regeneration occurs before starting a new sample. Separate accumulators for the appropriate likelihood ratios and their derivatives are maintained for each transient estimator, time horizon of interest, and parameter a sensitivity is computed with respect to. Thus, all measures can be estimated simultaneously from a single simulation run.

In the SAVE package, the user is able to compute sensitivities of all of the performance measures with respect to any continuous valued parameter of

the model, where a sensitivity of a performance measure with respect to some parameter is defined as the product of the parameter itself multiplied by the gradient of the performance measure with respect to the parameter. In its most general form, SAVE employs a symbolic differentiator to compute the derivatives needed during the simulation. This allows the use of complicated expressions to describe the parameters of the system. For example, one could specify that a failure rate of some component to be $5\lambda_1 + 4\lambda_2\lambda_3$, and then compute the sensitivity of some performance measure with respect to the parameter λ_2 . However, the computation of sensitivities is somewhat slow using this technique, with the extra CPU time needed to compute each sensitivity being about the same as the time needed to compute the regular (non-gradient) performance measure. Therefore, we have employed special techniques in the implementation of the SAVE package to allow the user to compute sensitivities with respect to certain parameters with little extra computational effort. If the user desires to compute sensitivities with respect to only component failure rates and repair rates and these rates are not themselves functions of other parameters, then the additional CPU time needed is small, as shown in Section 6.2.3. Note that when in expressions for the gradient of a performance measure given Equations 3.3 and 3.4, the only derivative terms which we have to compute are $Z'(\theta)$ and $q'(\theta, x, y)$ (since both $q(\theta, x)$ and $P(\theta, x, z)$ can be expressed solely in terms of $q(\theta, x, y)$). For the transient performance measures we consider in the SAVE package, we have $Z'(\theta) = 0$. In the case of steady-state performance measures computed in SAVE, we have that the role that $Z(\theta)$ played in the transient measure case has been replaced by $G(\theta)$ and $H(\theta)$ since we use conditional Monte Carlo (see Section 4.1), which depend on the parameter θ only through $q(\theta, x, y)$. Now note that $q(\theta, x, y)$ is an integer multiple, say k , of either a component failure

rate or repair rate (where k is either the redundancy of the component or the number of busy repairmen). So by only computing sensitivities with respect to simple failure and repair rates, the derivative of $q(\theta, x, y)$ is either 0 or k , thus allowing us to bypass the symbolic differentiator.

6 Experimental Results

In this section, we will discuss the results of simulations of two different models in order to analyze the behavior of gradient estimates via the likelihood ratios method and to demonstrate the effectiveness of different variance reduction techniques. We compare the sensitivities to the regular estimates in many cases in order to benchmark our results, where we define the sensitivity of a measure with respect to a parameter to be the product of the gradient of the measure with respect to the parameter multiplied by the value of the parameter itself. All numerical (non-simulation) and simulation results were obtained using the SAVE package ([11]).

6.1 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 we varied the number of components n from 2 to 12. In order 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 and formed estimates of the steady-state unavailability and the sensitivities of it with respect to λ . Table

1 contains the values of the numerical results of steady-state unavailability and its sensitivity with respect to the failure rate λ , and also their respective point estimates and percentage relative half-widths of the 90% confidence intervals obtained using direct simulation for each of the experiments. The percentage relative half-width of a confidence interval 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 sensitivities 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. 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 in order that the value of the steady-state unavailability remains constant. Since the gradient of the likelihood ratio turns out to be a sum of random variables, as the regenerative cycles become longer, we are summing up more random variables, which in turn leads to more variability.

6.2 Balanced and Unbalanced Systems

The next model we experimented with is a large computing system, whose block diagram is shown in Figure 6.2. This model is also discussed in [12] and [16]. We use two different parameter sets to create a "balanced" and an "unbalanced" system. In order 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 in order 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




Figure 1: A block diagram of the computing system modeled.

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

We make minor changes to the above parameters' settings in order 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.

6.2.1 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 regular estimates, and the sensitivities of smaller magnitude are not estimated as precisely, as shown by the example in Section 4.3.2. Numerical (non-simulation) results for these measures and their sensitivities were obtained using the SAVE package [11]. 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 [16]. These bounds are very tight and typically do not differ from the exact results significantly. We simulate both the balanced and the unbalanced systems. The goal of the simulation experiments is to demonstrate that we can obtain estimates of certain sensitivities that have approximately the same relative error as the regular estimate. Also, we see that the various variance reduction techniques have the same effect on the sensitivity estimates as they do on the regular estimates. Significant variance reductions can be obtained using the *Bias1/Bias2* method for the balanced systems and *Bias1/Balancing* method for the unbalanced systems, as is shown in [12]. These results hold for both the regular estimates and the sensitivities.

Tables 2 and 3 show the results obtained for the balanced and the unbalanced systems, respectively. We ran the simulation 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%. The percentage relative half-width of a confidence interval is defined to be 100% times the confidence interval half-width divided by the point estimate. This corresponds to approximately 100,000 events for each entry in Table 2 and 1,000,000 events for each entry in Table 3, respectively. Based on empirical results obtained in [10], the values for $bias1 = 0.5$ and $bias2 = 0.5$ were selected for DIS.

There are a few important points to note in the tables. For the balanced system, we used the *Bias1/Bias2* method, and *Bias1/Balancing* is used for the unbalanced system. As is shown in [10], these methods are most effective for their respective models when estimating the regular (non-gradient) 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 regular estimate.

The relative precision of the regular estimates and 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 4.3.1 Also, as claimed in Section 4.3.2, we do not obtain as accurate estimates for the sensitivity with respect to $c1fr$ since it is of smaller magnitude. It is also interesting to note that the amount of improvement from importance sampling over direct simulation in the sensitivities is about the same as the improvement in the regular estimates. This is because the same likelihood ratio needed for importance sampling is used in both the regular estimate and the sensitivities, 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 direct simulation given in Table 3 are very poor. This is because the value of $c1fr$ is much smaller than the value of parameter $procfr$, 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 and the variance of the sensitivity with respect to $c1fr$. However, when using *Bias1/Balancing*, we are able to obtain much better estimates of these quantities.

We next performed coverage experiments (see e.g., [18]) to determine the validity of the confidence intervals that are formed based on the asymptotic central limit theorems described in Section 4. Such studies are important since certain variance reduction techniques sometimes do not produce valid confidence intervals, except for very long run-lengths (see e.g., [18]).

We performed experiments on estimates of the steady-state unavailability, U , and its sensitivities with respect to both $rr2$ and $c1fr$, denoted by $U_{,,2}$ and

$U_{cl/r}$, 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 regular estimate and $\hat{U}_{\theta,1}, \dots, \hat{U}_{\theta,R}$, for $\theta = rr2$ and cl/r , 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 $U_{cl/r}$, respectively. If the confidence interval is valid, then by definition, the 90% coverage should be equal to 90%.

We also computed 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. The mean is computed over all replications with a nonzero point estimate. The results are listed in Table 4. Note that, as also seen in [12], the estimates using direct simulation are significantly more biased than those using importance sampling, and that its confidence intervals are about an order of magnitude wider. Also note that the values of the relative bias and relative half widths for the sensitivities with respect to *rr2* are about the same as those for the regular estimate, while these values for the other sensitivity are generally worse. This agrees with the results given in Section 4.3.2. Furthermore, for the small run length, the coverage drops significantly below

90% when using direct 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. Even 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. As a consequence we found it unnecessary to implement alternative variance reduction techniques to be used for steady-state regenerative gradient estimation in the SAVE package.

6.2.2 Transient Measures

In this section we discuss the results of our experiments for estimating reliability and its sensitivity with respect to both $rr2$ and c/fr . Recall that for transient measures we not only want the system to move quickly towards the set of system failed states F , but also 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 3.1) 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 3.1 and in [12]. The goal of the simulation is to study the relative accuracy of the regular estimate versus 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 5.

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 [12].

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 regular estimate, which strays from the result that we obtained for the steady-state measures. This is because we are working with transient measures. The likelihood ratio therefore 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. It is also interesting to note that the relative accuracy of the sensitivity estimates degrades compared to that of the regular estimate as the time horizon increases. This is because the length of each observation increases as the time horizon increases, 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.

6.2.3 Timing Experiments

Finally, Table 6 shows the results from some timing experiments which we performed in order to determine how much extra CPU time is required to compute

sensitivities. The experiments consisted of different simulation runs in which we varied the number of sensitivities computed and recorded the amount of CPU time taken in each run. All of the experiments were carried out on an IBM 3090 computer using the SAVE package, simulating the balanced system with the *bina1/bina2* (0.5/0.5) technique for 100,000 events. As one can see, there is a fairly large fixed cost in CPU time for computing any gradients, but the marginal cost in CPU time for computing each additional gradient is small. It is interesting to note that the additional time required to compute eight sensitivities is about the same as the amount of time needed to run SAVE when computing no gradients.

7 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 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 4.3).
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.

Acknowledgments. We would like to credit Adrian Conway for initial work on combining likelihood ratio gradient estimation with importance sampling for steady-state measures when using discrete-time conversion.

References

- [1] P. Billingsley (1961), *Statistical Inference for Markov Processes*, University of Chicago Press, Chicago.
- [2] M. A. Crane, and D. L. Iglehart (1974), "Simulating Stable Stochastic Systems, II: Markov Chains," *J. Assoc. Comput. Mach.* **21**, pp. 114-123.
- [3] B. L. Fox and P. W. Glynn (1986), "Discrete time conversion for simulating semi-Markov processes," *Operations Research Letters* **5**, pp. 191-196.
- [4] R. M. Geist and K. S. Trivedi (1983), "Ultra-High Reliability Prediction for Fault-Tolerant Computer Systems," *IEEE Transactions on Computers* **C-32**, pp. 1118-1127.
- [5] P. Glasserman (1989), "Derivative Estimates From Simulation of Continuous Time Markov Chains," Manuscript.
- [6] P. W. Glynn (1986), "Stochastic Approximation For Monte Carlo Optimization," *Proceedings of the 1986 Winter Simulation Conference*, J. Wilson, J. Henriksen, and S. Roberts (eds.), IEEE Press, pp. 356-364.
- [7] P. W. Glynn (1987), "Likelihood Ratio Gradient Estimation: An Overview," *Proceedings of the 1987 Winter Simulation Conference*, A. Thesen, H. Grant, and W. D. Kelton (eds.), IEEE Press, pp. 366-375.
- [8] P. W. Glynn (1989), "Likelihood Ratio Derivative Estimators for Stochastic Systems," *Proceedings of the 1989 Winter Simulation Conference*, E. A. MacNair, K. J. Musselman, and P. Heidelberger (eds.), IEEE Press, pp. 374-380.

- [9] P. W. Glynn and D. L. Iglehart (1986), "Importance Sampling for Stochastic Simulations," to appear in *Management Science*.
- [10] A. Goyal, P. Heidelberger, P. Shahabuddin (1987), "Measure Specific Dynamic Important Sampling for Availability Simulations," *Proceedings of the 1987 Winter Simulation Conference*. A. Thesen, H. Grant, and W. D. Kelton (eds.), IEEE Press, pp. 351-357.
- [11] A. Goyal and S. S. Lavenberg (1987), "Modeling and Analysis of Computer System Availability," *IBM Journal of Research and Development* **31** 6, pp. 651-664.
- [12] A. Goyal, P. Shahabuddin, P. Heidelberger, V. F. Nicola, and P. W. Glynn (1989), "A Unified Framework for Simulating Markovian Models of Highly Dependable Systems," IBM Research Report RC 14772, Yorktown Heights, New York.
- [13] A. Hordijk, D. L. Iglehart, and R. Schassberger (1976), "Discrete-time methods of simulating continuous-time Markov chains," *Adv. Appl. Prob.* **8**, pp. 772-788.
- [14] D. E. Knuth (1976), "Big Omicron, Big Omega and Big Theta," *SIGACT News (ACM)* **8**, 2, pp. 18-24.
- [15] E. E. Lewis and F. Böhm (1984), "Monte Carlo Simulation of Markov Unreliability Models," *Nuclear Engineering and Design* **77**, pp. 49-62.
- [16] R. R. Muntz, 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.

- [17] J. C. Laprie (1984), "Dependability Evaluation of Software Systems in Operation," *IEEE Transactions Software Engineering* **SE-10**, pp. 701-713.
- [18] S. S. Lavenberg, T. L. Moeller, and C. H. Sauer (1979), "Concomitant Control Variables Applied to the Regenerative Simulation of Queuing Systems," *Operations Research* **27**, 1, pp. 134-160.
- [19] P. L'Ecuyer, N. Giroux, and P. W. Glynn (1989), "Stochastic Optimization by Simulation: Some Experiments with a Simple Steady-State Queue," Technical Report, Department of Operations Research, Stanford University, Stanford, California.
- [20] M. I. Reiman, and A. Weiss (1989). "Sensitivity Analysis for Simulations Via Likelihood Ratios," *Operations Research* **37**, 5, pp. 830-844.
- [21] S. M. Ross (1983), *Stochastic Processes*, John Wiley, New York.
- [22] R. Y. Rubenstein (1986), "The Score Function Approach for Sensitivity Analysis of Computer Simulation Models," *Math. and Computers in Simulation* **28**, pp. 351-379.
- [23] R. Y. Rubenstein (1989), "Sensitivity Analysis and Performance Extrapolation for Computer Simulation Models," *Operations Research* **37**, 1, pp. 72-81.
- [24] R. J. Serfling (1980), *Approximation Theorems for Mathematical Statistics*, John Wiley, New York.
- [25] R. S. Sutor (1989). "A Guide to Programming in the Scratchpad II Interpreter," IBM T. J. Watson Research Center, Yorktown Heights, New York.

8 Appendix

Now we will justify the interchange of derivative and expectation. In order to do so, we will make the following assumptions:

A1. State space E is finite.

A2. $Q(\cdot)$ is continuously differentiable for $\theta \in \Theta$.

A3. $P(\theta)$ is irreducible for $\theta \in \Theta$.

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

A5. T is a stopping time satisfying the following two conditions:

1. $P_\theta\{T > 0\} = 1$ for all $\theta \in \Theta$.
2. 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 (measurable) set of even;
2. $Z(\theta) = \int_0^T f(s, Y_s) ds$, where T is some stopping time satisfying assumption A5 and f is a real-valued function defined on (Θ, E) satisfying assumptions A7 and A8.

A7. $\|f\| = \sup\{|f(\theta, x)| : \theta \in \Theta, x \in E\} < \infty$.

A8. $\|f'\| = \sup\{|f'(\theta, x)| : \theta \in \Theta, x \in E\} < \infty$.

Note that under assumptions A1 and A3, our CTMC Y is a regenerative process. Also note that when we express a steady-state measure $r(\theta)$ using the ratio formula, we have that $r(\theta)$ is a ratio of two expectations of random variables, each having the second form of $Z(\theta)$ given in assumption A6.

Also, note that by assumption A1, we can let $\Gamma = \Gamma(\theta)$, i.e. $\Gamma(\theta)$ is independent of θ .

The first condition of assumption A5 is necessary to insure the validity of the ratio formula. Also, there are numerous examples of stopping times which satisfy assumption A5.

Proposition 1 Define $T = T_{N(t)+1}$, for some deterministic time t , i.e. T is the time of the first transition after time t . Suppose assumptions A1 and A2 hold. Then, T satisfies assumption A5.

Proof. First, define $q^* = \sup\{q(\theta, i) : \theta \in \Theta, i \in E\}$ and $N^*(t) = \sup\{n \geq 0 : t_0^* + \dots + t_n^* \leq t\}$, where t_k^* is an exponentially distributed random variable with mean $1/q^*$ for all k . By assumptions A1 and A2, $q^* < \infty$. Therefore, $\{N^*(t) : t \geq 0\}$ is a Poisson process with rate q^* . Let $M^*(z)$ be the moment generating function of $N^*(t)$. Then

$$M^*(z) = e^{(e^z - 1)q^*t} < \infty$$

for all finite t and z . Hence, $N^*(t)$ has a convergent moment generating function for all finite t and z . Now, $P_\theta\{N(t) > k\} \leq P_\bullet\{N^*(t) > k\}$ for all $\theta \in \Theta$, where $P_\bullet\{\cdot\}$ is the probability measure corresponding to $\{N^*(t) : t \geq 0\}$ (see [21]). Thus, we have that $N(t)$ also has a convergent moment generating function for all finite t and z . ■

Proposition 2 Define $T = t$, where $t > 0$ is some deterministic time. Suppose assumptions A1 and A2 hold. Then, T satisfies assumption A5.

Proof. Since $t < T_{N(t)+1}$, the result follows from Proposition 1. ■

Proposition 3 Define $T = \alpha_A = \inf\{t > 0 : Y_{t-} \notin A, Y_t \in A\}$ for some set of states A , i.e. T is the hitting time to some set of states A . Suppose assumptions A1 - A4 hold. Then, T satisfies assumption A5.

Proof. Since $T = \inf\{t : Y_{t-} \notin A, Y_t \in A\}$ for some set A , we have that $N(T) = \inf\{n \geq 0 : X_n \in A\}$. Let $x \in A$. By assumptions A3, and A4, we have that for each $y \in E$, there exists an integer $m(y, x)$ such that $P^{m(y, x)}(\theta, y, x) > 0$, where $P^k(\theta)$ is the k -step transition matrix of the embedded DTMC X . Let $m = \max\{m(y, x) : y \in E, x \in A\}$, which is finite since $|E| < \infty$. Now we have that for all $x \in A$,

$$\begin{aligned} P_\theta\{N(T) > m | X_0 = y\} &\leq P_\theta\{X_{m(y, x)} \neq x | X_0 = y\} \\ &= 1 - P^{m(y, x)}(\theta, y, x) \\ &\leq \rho, \end{aligned}$$

where $\rho = \sup\{1 - P^{m(y, x)}(\theta, y, x) : \theta \in \Theta, y \in E, x \in A\}$. By assumptions A1, A2, and A4, $\inf\{P(\theta, x, y) : \theta \in \Theta, (x, y) \in \Gamma\} > 0$, and so $\rho < 1$.

Now we have

$$\begin{aligned} P_\theta\{N(T) > mn\} &= \sum_{y \in A} P_\theta\{N(T) > m(n-1), X_{m(n-1)} = y\} P_\theta\{N(T) > m | X_0 = y\} \\ &\leq \rho \sum_{y \in A} P_\theta\{N(T) > m(n-1), X_{m(n-1)} = y\} \\ &= \rho P_\theta\{N(T) > m(n-1)\}. \end{aligned}$$

By induction, we obtain

$$P_\theta\{N(T) > mn\} \leq \rho^n.$$

So we have the moment generating function

$$\begin{aligned} M_\theta^{N(T)}(z) &= \sum_{n=0}^{\infty} e^{zn} P_\theta\{N(T) = n\} \\ &= P_\theta\{N(T) = 0\} + \sum_{n=0}^{\infty} \sum_{l=1}^m e^{z(nm+l)} P_\theta\{N(T) = nm+l\} \\ &\leq P_\theta\{N(T) = 0\} + \sum_{n=0}^{\infty} P_\theta\{N(T) > nm\} \sum_{l=1}^m e^{z(nm+l)}. \end{aligned}$$

For $z \leq 0$, we have that

$$\begin{aligned} M_{\theta}^{N(T)}(z) &\leq P_{\theta}\{N(T) = 0\} + \sum_{n=0}^{\infty} m P_{\theta}\{N(T) > nm\} e^{znm} \\ &\leq P_{\theta}\{N(T) = 0\} + m \sum_{n=0}^{\infty} [\rho e^{zm}]^n. \end{aligned}$$

For $z > 0$, we have that

$$\begin{aligned} M_{\theta}^{N(T)}(z) &\leq P_{\theta}\{N(T) = 0\} + \sum_{n=0}^{\infty} m P_{\theta}\{N(T) > nm\} e^{z(n+1)m} \\ &\leq P_{\theta}\{N(T) = 0\} + m e^{zm} \sum_{n=0}^{\infty} [\rho e^{zm}]^n. \end{aligned}$$

Hence, there exists some $z_0 > 0$ such that $M_{\theta}^{N(T)}(z) < \infty$ for all $z \in (-z_0, z_0)$ and all $\theta \in \Theta$. ■

We now state a lemma.

Lemma 1 *If assumptions A1-A8 hold, then*

$$E_{\theta} Z(\theta)^k < \infty$$

$$E_{\theta} Z'(\theta)^k < \infty$$

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

Proof. When $Z(\theta) = I_S$, 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 $|Z'(\theta)| \leq \|f'\|T$. Now assumption A5 implies that $E_{\theta} T^k < \infty$ for all k and all $\theta \in \Theta$, which, along with assumptions A7 and A8, gives us our result. ■

Hence, we have that the performance measures discussed in Section 2.3 satisfy assumptions A5-A8. Now we will justify the interchange of the derivative and expectation.

Theorem 1 If assumptions A1 - A8 hold, then

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

Proof. 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

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

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

From Equation 3.2, we have that $L'(T, \eta, \theta_0)$ is equal to

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

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

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

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

$$\begin{aligned} \phi_1(h) &= (N(T) + 1)(\|q'/q\| + \|P'/P\|) + \|q'\| \sum_{j=0}^{N(T)+1} t_j \\ \phi_2(h) &= \|q/q(\theta_0)\|^{N(T)+1} \|P/P(\theta_0)\|^{N(T)+1} \exp \left\{ \|q - q(\theta_0)\| \sum_{k=0}^{N(T)+1} t_k \right\}. \end{aligned} \quad (8.1)$$

Note that we can bound $|Z(\theta_0)L'(T, \eta, \theta_0)|$ by $\phi(h) \equiv |Z(\theta_0)|\phi_1(h)\phi_2(h)$. So we now want to show that $\phi(h)$ is integrable for h sufficiently small. To do this, we will show that there exists some $z_0 > 0$ such that $M_\theta(z)$, which we define as the moment generating function of $\sum_{k=0}^{N(T)} t_k$, converges for all $z \in (-z_0, z_0)$ and all $\theta \in \Theta$.

First define $q^* = \inf\{q(\theta, x) : \theta \in \Theta, x \in E\}$. Then

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

for z sufficiently small since $N(T)$ has a convergent moment generating function in a neighborhood of 0 by assumption A5. Since $M_\theta^{N(T)}(z)$ and $M_\theta(z)$ both converge in a neighborhood of 0 for all $\theta \in \Theta$, we have that $N(T)$ and $\sum_{k=0}^{N(T)} t_k$ have finite moments of all orders for all $\theta \in \Theta$.

Now note that by assumption A2, $\|q/q(\theta_0)\| \rightarrow 1$, $\|P/P(\theta_0)\| \rightarrow 1$, and $\|q - q(\theta_0)\| \rightarrow 0$ as $h \downarrow 0$. Hence, by repeated applications of the Schwarz inequality and using Lemma 1, 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. ■

Now we will give a proof of the asymptotic variance of our estimator of $r'(\theta_0)$ given in Equation 4.1. In order to do this, we need the following result (see [21], p. 118, for the proof).

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

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

as $n \rightarrow \infty$, where

$$\bar{X}_n = \frac{1}{n} \sum_{i=1}^n X_i$$

$$\sigma^2 = \nabla g(\mu)^T \Sigma \nabla g(\mu)$$

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 - A8 hold, then the asymptotic variance of the estimate of $r'(\theta_0)$ using Equation 4.1 is given by Equation 4.2.*

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

$$A = Z_T'(\theta) + Z(\theta)L'(T, \theta_0, \theta_0)$$

$$B = T$$

$$C = Z_T(\theta)$$

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

as in Equations 4.3-4.6. In order to apply Theorem 2, we first need to show that $E_{\theta_0} \|V\|^2 < \infty$. Assumption A5 and Lemma 1 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 Equation 8.1 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 $E_{\theta_0} \|V\|^2 < \infty$.

In order to apply Theorem 2, we define $g : \mathbb{R}^4 \rightarrow \mathbb{R}$ as

$$g(a, b, c, d) = \frac{ab - cd}{b^2}. \quad (8.2)$$

The first condition of assumption A5 assures that g is differentiable at the point $(\alpha, \beta, \gamma, \delta)$. Thus, by computing the gradient of g and plugging in the appropriate values into the expression for the variance given in Theorem 2, the proof is complete. ■

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. We define the vector $\mathbf{V} = (A, B, C, D)$ as in the proof of Theorem 3, and we let $\mathbf{W} = E_{\theta_n}[\mathbf{V}|\mathbf{X}] = (\hat{A}, \hat{B}, \hat{C}, \hat{D})$, where

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

where $G(\theta)$, $H(\theta)$, $G'(\theta)$, $H'(\theta)$, and $\dot{J}'(\tau_0, \theta_0, \theta_n)$ are defined in Equations 4.8, 4.9, 4.12, 4.13, and 4.14, respectively. Let $\mu = E_{\theta_n}\mathbf{V} = E_{\theta_n}\mathbf{W}$. 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 Equation 8.2, and

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

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 4 $\sigma_2^2 \leq \sigma_1^2$.

Proof. By noting that

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

we have the result by the principle of conditional Monte Carlo (see [3]). ■

Number of Comps (n)	Failure Rate (λ)	Avg Number of Events per Regen Cycles	Steady-State Unavailability		Sensitivity w.r.t. λ	
			Numerical Result	Direct Simulation	Numerical Result	Direct Simulation
2	0.0229	2.05	0.1000×10^{-2}	$0.0991 \times 10^{-2} \pm 1.6\%$	0.1954×10^{-2}	$0.1934 \times 10^{-2} \pm 1.6\%$
4	0.0884	2.63	0.1000×10^{-2}	$0.1014 \times 10^{-2} \pm 4.4\%$	0.3593×10^{-2}	$0.3559 \times 10^{-2} \pm 4.7\%$
6	0.1233	4.12	0.1000×10^{-2}	$0.0999 \times 10^{-2} \pm 6.5\%$	0.4901×10^{-2}	$0.4872 \times 10^{-2} \pm 7.3\%$
8	0.1375	7.53	0.1000×10^{-2}	$0.1007 \times 10^{-2} \pm 6.5\%$	0.5859×10^{-2}	$0.5677 \times 10^{-2} \pm 9.0\%$
10	0.1426	16.39	0.1000×10^{-2}	$0.1010 \times 10^{-2} \pm 8.1\%$	0.6457×10^{-2}	$0.6230 \times 10^{-2} \pm 10.4\%$
12	0.1442	43.89	0.1000×10^{-2}	$0.1020 \times 10^{-2} \pm 8.3\%$	0.6755×10^{-2}	$0.6784 \times 10^{-2} \pm 13.1\%$

Table 1: Estimates of steady-state unavailability and sensitivities with relative 90% confidence intervals for an n-component system using direct simulation (1,024,000 events)

Performance Measure	Regular Estimate		
	Numerical Result	Direct Simulation	Bias1/Bias2 (0.5/0.5)
Unavailability	0.9309×10^{-5}	1.0171×10^{-5} $\pm 27.1\%$	0.9395×10^{-5} $\pm 2.7\%$
MTTF	$0.1637 \times 10^{+6}$	$0.1524 \times 10^{+6}$ $\pm 25.7\%$	$0.1626 \times 10^{+6}$ $\pm 2.5\%$

(a) Results for regular estimates

Performance Measure	Sensitivity w.r.t. $rr2$		
	Numerical Result	Direct Simulation	Bias1/Bias2 (0.5/0.5)
Unavailability	$-.1252 \times 10^{-4}$	$-.1256 \times 10^{-4}$ $\pm 33.0\%$	$-.1265 \times 10^{-4}$ $\pm 3.3\%$
MTTF	$0.1099 \times 10^{+6}$	$0.0879 \times 10^{+6}$ $\pm 33.4\%$	$0.1109 \times 10^{+6}$ $\pm 2.6\%$

(b) Results for sensitivities w.r.t. $rr2$

Performance Measure	Sensitivity w.r.t. $clfr$		
	Numerical Result	Direct Simulation	Bias1/Bias2 (0.5/0.5)
Unavailability	0.2315×10^{-5}	0.3720×10^{-5} $\pm 64.1\%$	0.2590×10^{-5} $\pm 6.3\%$
MTTF	$-.4065 \times 10^{+5}$	$-.5893 \times 10^{+5}$ $\pm 58.6\%$	$-.4418 \times 10^{+5}$ $\pm 6.0\%$

(c) Results for sensitivities w.r.t. $clfr$

Table 2: Estimates of steady-state unavailability, MTTF, and sensitivities with relative 90% confidence intervals for the balanced system (100,000 events)

Performance Measure	Regular Estimate		
	Numerical Result	Direct Simulation	Bias/Balancing (0.5)
Unavailability	0.6967×10^{-7}	0.4165×10^{-7} $\pm 164.5\%$	0.6976×10^{-7} $\pm 2.4\%$
MTTF	$0.2188 \times 10^{+8}$	$0.4703 \times 10^{+8}$ $\pm 164.5\%$	$0.2183 \times 10^{+8}$ $\pm 2.3\%$

(a) Results for regular estimates

Performance Measure	Sensitivity w.r.t. $rr2$		
	Numerical Result	Direct Simulation	Bias/Balancing (0.5)
Unavailability	-9.436×10^{-7}	-7.939×10^{-7} $\pm 164.5\%$	-9.384×10^{-7} $\pm 3.1\%$
MTTF	$0.1481 \times 10^{+8}$	$4.555 \times 10^{+8}$ $\pm 164.5\%$	$0.1470 \times 10^{+8}$ $\pm 2.4\%$

(b) Results for sensitivities w.r.t. $rr2$

Performance Measure	Sensitivity w.r.t. $c1fr$		
	Numerical Result	Direct Simulation	Bias/Balancing (0.5)
Unavailability	0.2324×10^{-7}	-3.425×10^{-10} $\pm 166.8\%$	0.2361×10^{-7} $\pm 6.0\%$
MTTF	$-7.298 \times 10^{+7}$	$0.1104 \times 10^{+8}$ $\pm 191.0\%$	$-7.318 \times 10^{+7}$ $\pm 5.2\%$

(c) Results for sensitivities w.r.t. $c1fr$

Table 3: Estimates of steady-state unavailability, MTTF, and sensitivities with relative 90% confidence intervals for the unbalanced system (1,000,000 events)

Events per Rep	Steady-State Unavailability					
	Direct Simulation			Bias1/Bias2 (0.5/0.5)		
	Rel Bias (Std Dev)	Rel HW	Coverage	Rel Bias (Std Dev)	Rel HW	Coverage
2000	6.95% (12.82%)	144.40%	54%	0.74% (1.20%)	18.88%	85%
20000	-3.94% (3.41%)	65.47%	90%	0.39% (0.34%)	5.99%	92%
200000	1.29% (1.09%)	19.60%	96%	0.05% (0.13%)	1.90%	90%

(a) Results for steady-state unavailability

Events per Rep	Sensitivity of unavailability w.r.t. $rr2$					
	Direct Simulation			Bias1/Bias2 (0.5/0.5)		
	Rel Bias (Std Dev)	Rel HW	Coverage	Rel Bias (Std Dev)	Rel HW	Coverage
2000	14.03% (19.15%)	155.96%	46%	0.37% (1.54%)	23.26%	84%
20000	-1.86% (4.96%)	82.30%	84%	0.40% (0.43%)	7.46%	92%
200000	2.23% (1.45%)	25.27%	94%	0.03% (0.16%)	2.38%	86%

(b) Results for sensitivity of unavailability w.r.t. $rr2$

Events per Rep	Sensitivity of unavailability w.r.t. $c1fr$					
	Direct Simulation			Bias1/Bias2 (0.5/0.5)		
	Rel Bias (Std Dev)	Rel HW	Coverage	Rel Bias (Std Dev)	Rel HW	Coverage
2000	20.73% (36.64%)	432.36%	11%	3.73% (2.86%)	46.45%	90%
20000	-2.13% (10.66%)	147.02%	62%	-0.18% (1.04%)	14.86%	82%
200000	5.25% (3.62%)	56.79%	86%	0.10% (0.27%)	4.68%	95%

(c) Results for sensitivity of unavailability w.r.t. $c1fr$

Table 4: Coverage experiments for estimates of steady-state unavailability and sensitivities on the balanced system (100 replications)

Time (t)	Numerical Result	Unreliability			
		Direct Simulation		Bias1/Bias2 (0.5/0.5)	
		Standard	Forcing	Standard	Forcing
4	0.1528×10^{-4}	0.1034×10^{-4} $\pm 116.3\%$	0.1481×10^{-4} $\pm 23.9\%$	0.1583×10^{-4} $\pm 7.0\%$	0.1522×10^{-4} $\pm 1.4\%$
16	0.8734×10^{-4}	1.0721×10^{-4} $\pm 37.7\%$	0.9428×10^{-4} $\pm 22.8\%$	0.8693×10^{-4} $\pm 3.3\%$	0.8699×10^{-4} $\pm 1.3\%$
64	0.3804×10^{-3}	0.3552×10^{-3} $\pm 24.0\%$	0.3421×10^{-3} $\pm 21.9\%$	0.3801×10^{-3} $\pm 1.8\%$	0.3841×10^{-3} $\pm 1.1\%$
256	0.1552×10^{-2}	0.1463×10^{-2} $\pm 16.8\%$	0.1578×10^{-2} $\pm 19.9\%$	0.1565×10^{-2} $\pm 1.5\%$	0.1578×10^{-2} $\pm 1.0\%$
1024	0.6225×10^{-2}	0.5598×10^{-2} $\pm 14.9\%$	0.6254×10^{-2} $\pm 16.0\%$	0.6275×10^{-2} $\pm 4.9\%$	0.6233×10^{-2} $\pm 4.3\%$

(a) Results for unreliability

Time (t)	Numerical Result	Sensitivity of Unreliability w.r.t. $rr2$			
		Direct Simulation		Bias1/Bias2 (0.5/0.5)	
		Standard	Forcing	Standard	Forcing
4	-0.4353×10^{-5}	-0.1778×10^{-5} $\pm 117.0\%$	-0.5126×10^{-5} $\pm 39.8\%$	-0.4113×10^{-5} $\pm 13.4\%$	-0.4280×10^{-5} $\pm 2.9\%$
16	-0.4886×10^{-4}	-0.6391×10^{-4} $\pm 59.8\%$	-0.4799×10^{-4} $\pm 48.0\%$	-0.4845×10^{-4} $\pm 6.1\%$	-0.4853×10^{-4} $\pm 2.5\%$
64	-0.2455×10^{-3}	-0.2084×10^{-3} $\pm 46.7\%$	-0.1679×10^{-3} $\pm 41.6\%$	-0.2420×10^{-3} $\pm 3.6\%$	-0.2467×10^{-3} $\pm 2.3\%$
256	-0.1031×10^{-2}	-0.0928×10^{-2} $\pm 35.0\%$	-0.0877×10^{-2} $\pm 42.4\%$	-0.1040×10^{-2} $\pm 4.0\%$	-0.1055×10^{-2} $\pm 2.7\%$
1024	-0.4156×10^{-2}	-0.5200×10^{-2} $\pm 33.4\%$	-0.4504×10^{-2} $\pm 40.8\%$	-0.4220×10^{-2} $\pm 15.0\%$	-0.3965×10^{-2} $\pm 12.6\%$

(b) Results for sensitivities of unreliability w.r.t. $rr2$

Sensitivity of Unreliability w.r.t. $c1fr$					
Time (t)	Numerical Result	Direct Simulation		Bias1/Bias2 (0.5/0.5)	
		Standard	Forcing	Standard	Forcing
4	0.3788×10^{-5}	-0.3493×10^{-7} $\pm 116.7\%$	0.4685×10^{-5} $\pm 59.2\%$	0.4038×10^{-5} $\pm 19.5\%$	0.3719×10^{-5} $\pm 4.3\%$
16	0.2168×10^{-4}	-0.0108×10^{-4} $\pm 40.6\%$	0.2532×10^{-4} $\pm 61.1\%$	0.2118×10^{-4} $\pm 9.3\%$	0.2141×10^{-4} $\pm 3.8\%$
61	0.9448×10^{-4}	0.8050×10^{-4} $\pm 74.2\%$	1.0965×10^{-4} $\pm 63.9\%$	0.9301×10^{-4} $\pm 5.5\%$	0.9367×10^{-4} $\pm 3.5\%$
256	0.3853×10^{-3}	0.3057×10^{-3} $\pm 57.4\%$	0.6703×10^{-3} $\pm 44.6\%$	0.3971×10^{-3} $\pm 6.2\%$	0.3955×10^{-3} $\pm 4.2\%$
1024	0.1542×10^{-2}	0.1415×10^{-2} $\pm 60.6\%$	0.2143×10^{-2} $\pm 48.7\%$	0.1247×10^{-2} $\pm 27.1\%$	0.1377×10^{-2} $\pm 22.8\%$

(c) Results for sensitivities of unreliability w.r.t. $c1fr$

Table 5: Estimates of unreliability and sensitivities with relative 90% confidence intervals for the balanced system (400,000 events)

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

Table 6: CPU times taken for computing sensitivities in balanced system using *bias1/bias2* (0.5/0.5) for 100,000 simulated events

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified			1b. RESTRICTIVE MARKINGS	
2a. SECURITY CLASSIFICATION AUTHORITY			3. DISTRIBUTION/AVAILABILITY OF REPORT Approved for public release; distribution unlimited.	
2b. DECLASSIFICATION/DOWNGRADING SCHEDULE				
4. PERFORMING ORGANIZATION REPORT NUMBER(S) Technical Report No. 54			5. MONITORING ORGANIZATION REPORT NUMBER(S) ARO 25839.30-MA	
6a. NAME OF PERFORMING ORGANIZATION Dept. of Operations Research		6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION U. S. Army Research Office	
6c. ADDRESS (City, State, and ZIP Code) Stanford, CA 94305-4022			7b. ADDRESS (City, State, and ZIP Code) P. O. Box 12211 Research Triangle Park, NC 27709-2211	
8a. NAME OF FUNDING/SPONSORING ORGANIZATION U. S. Army Research Office		8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER DAA603-88-K-0063	
8c. ADDRESS (City, State, and ZIP Code) P. O. Box 12211 Research Triangle Park, NC 27709-2211			10. SOURCE OF FUNDING NUMBERS	
			PROGRAM ELEMENT NO. PROJECT NO. TASK NO. WORK UNIT ACCESSION NO.	
11. TITLE (Include Security Classification) Likelihood ratio sensitivity analysis for Markovian models of highly dependable systems.				
12. PERSONAL AUTHOR(S) Marvin K. Nakayama, Ambuj Goyal, Peter W. Glynn				
13a. TYPE OF REPORT Technical		13b. TIME COVERED FROM TO		14. DATE OF REPORT (Year, Month, Day) January 1990
15. PAGE COUNT 57				
16. SUPPLEMENTARY NOTATION The view, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy, or decision, unless so designated by other documentation.				
17. COSATI CODES			18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)	
FIELD	GROUP	SUB-GROUP	highly dependable systems, likelihood ratios, importance sampling, gradient estimators.	
19. ABSTRACT (Continue on reverse if necessary and identify by block number) (please see next page)				
20. DISTRIBUTION/AVAILABILITY OF ABSTRACT ☐ UNCLASSIFIED/UNLIMITED ☐ SAME AS RPT. ☐ DTIC USERS			21. ABSTRACT SECURITY CLASSIFICATION Unclassified	
22a. NAME OF RESPONSIBLE INDIVIDUAL			22b. TELEPHONE (Include Area Code)	
			22c. OFFICE SYMBOL	

Likelihood Ratio Sensitivity Analysis for Markovian Models of Highly Dependable Systems

Marvin K. Nakayama¹
Department of Operations Research
Stanford University
Stanford, California 94305

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

Peter W. Glynn²
Department of Operations Research
Stanford University
Stanford, California 94305

December 20, 1989

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.

90 10 24 003