

ON THE SMALL-SAMPLE OPTIMALITY OF MULTIPLE-REGENERATION ESTIMATORS

James M. Calvin

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

Peter W. Glynn

Department of Engineering-
Economic Systems and
Operations Research
Stanford University
Stanford, CA 94305, U.S.A.

Marvin K. Nakayama

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

ABSTRACT

We describe a simulation output analysis methodology suitable for stochastic processes that are regenerative with respect to multiple regeneration sequences. Our method exploits this structure to construct estimators that are more efficient than those that are obtained with the standard regenerative method. We illustrate the method in the setting of discrete-time Markov chains on a countable state space, and we present a result showing that the estimator is the uniform minimum variance unbiased estimator for finite-state-space discrete-time Markov chains. Some empirical results are given.

1 INTRODUCTION

A regenerative stochastic process has a sequence of random times (regenerations) forming a renewal process such that the process from each regeneration time forward is a probabilistic replica of the original process. The regenerative method of simulation-output analysis exploits the fact that the consecutive regeneration cycles are independent and identically distributed (i.i.d.) by bringing to bear the well-developed theory of i.i.d. sequences; see, e.g., Shedler (1993) for details.

Many stochastic processes are regenerative with respect to more than one regeneration sequence. For example, the successive hitting times to any fixed state of a positive-recurrent discrete-time or continuous-time Markov chain on a countable state space form a regeneration sequence. The standard regenerative method does not take advantage of this multiple-regeneration structure, but there has been some previous work on developing estimators that exploit this structure to increase efficiency. Calvin and Nakayama (1998) developed simulation estimators that use two regeneration sequences, and proved that the estimators have the same mean as and no greater (and often significantly

smaller) variance than the standard regenerative estimator. Calvin and Nakayama (2000) proved strong laws of large numbers and central limit theorems for these estimators so confidence intervals can be constructed. Calvin and Nakayama (1999a) derived estimators for the expectation of a product of two additive cycle rewards, which is the performance measure we consider in the current paper, and certain importance-sampling estimators, both for the case when there are two regeneration sequences. Finally, Calvin and Nakayama (1999b) developed estimators for processes that have more than two regeneration sequences under the assumption that the order of regenerations from the various sequences have a “birth-death-type structure” (i.e., any regeneration from sequence i and any regeneration from sequence $i + 2$ must be separated by a regeneration from sequence $i + 1$, and any regeneration from sequence i and any regeneration from sequence $i - 2$ must be separated by a regeneration from sequence $i - 1$).

In this paper we develop estimators for simulations of processes having multiple regeneration sequences as in Calvin and Nakayama (1999b), but we eliminate the restriction that the regeneration sequences possess a birth-death-type structure. We limit our discussion in this paper to constructing estimators for discrete-time Markov chains on a countable state space; see Calvin, Glynn, and Nakayama (1999) for the extension to general processes with multiple regeneration sequences, and a discussion on the relationship between these types of processes and semi-regenerative processes (Chapter 10 of Çinlar 1975).

The rest of this paper is structured as follows. In Section 2 we describe the basic framework in which we derive our results, and give some examples of performance measures that can be estimated using our approach. In Section 3 we derive some identities for processes with multiple regeneration sequences, and these identities form the basis for our estimator, which we present in Section 4. Section 5 contains a result showing that our estimator is the

uniform minimum variance unbiased estimator for discrete-time Markov chains on a finite state space, which is a form of small-sample optimality. In Section 6 we present some empirical results.

2 REGENERATIVE MIXED-MOMENT ESTIMATORS

Let $X = (X_n : n = 0, 1, 2, \dots)$ be an irreducible, positive-recurrent discrete-time Markov chain (DTMC) with finite or countable state space $S \subset \{1, 2, 3, \dots\}$. Denote by $R = (R_{i,j} : i, j \in S)$ the transition probability matrix of X , and let $\pi = (\pi_i : i \in S)$ denote the stationary distribution. Let $S_0 = \{x_1, x_2, \dots, x_s\} \subset S$, $s \geq 1$, be a finite subset of the states, and let $T(0) = \inf\{k \geq 0 : X_k \in S_0\}$, and $T(n+1) = \inf\{k > T(n) : X_k \in S_0\}$, $n \geq 0$, denote the successive hitting times to the set S_0 . For $1 \leq i \leq s$, denote the successive hitting times to $x_i \in S_0$ by $T_i(0) = \inf\{k \geq 0 : X_k = x_i\}$, and $T_i(n+1) = \inf\{k > T_i(n) : X_k = x_i\}$, $n \geq 0$. We call the sample-path segment $(X_k : T_i(j-1) \leq k < T_i(j))$, $j \geq 1$, the j th T_i -cycle.

Define another Markov chain $W = (W_n : n \geq 0)$ by $W_n = X_{T(n)}$; i.e., W_n is the state of the X chain on the $(n+1)$ th visit to the set S_0 . We let $Q = (Q_{i,j} : i, j \in S)$ denote the transition probability matrix of W , and $\nu = (\nu_i : i \in S_0)$ its stationary distribution, which is given by $\nu_i = 1_{(i \in S_0)} \pi_i / \pi(S_0)$, where $\pi(S_0) = \sum_{j \in S_0} \pi_j$.

We consider estimating the mixed moment

$$\alpha = E[U(1) V(1)], \tag{1}$$

where

$$U(k) = \sum_{n=T_1(k-1)}^{T_1(k)-1} f_U(X_n, X_{n+1}), \tag{2}$$

$$V(k) = \sum_{n=T_1(k-1)}^{T_1(k)-1} f_V(X_n, X_{n+1}), \tag{3}$$

for $k \geq 1$, and $f_U, f_V : S \times S \rightarrow \Re$ are “reward” functions. Based on a sample path \vec{X} of exactly m T_1 -cycles, i.e., $\vec{X} = (X_n : 0 \leq n < T_1(m))$ with $X_0 = x_1$, we construct the “standard estimator” of α as

$$\hat{\alpha}(\vec{X}) = \frac{1}{m} \sum_{k=1}^m U(k) V(k).$$

(Throughout this paper we will assume that all of the estimators are based on sample paths \vec{X} consisting of exactly m T_1 -cycles with m fixed, but the estimators can also be constructed from more general sample paths; e.g., the path does not have to start and end in state x_1 nor does it have to consist of a fixed number of T_1 -cycles. However, to

simplify the presentation we consider only the case when \vec{X} consists of exactly m T_1 -cycles.)

Calvin and Nakayama (1999a,1999b) developed estimators for (1) in more restrictive settings than we consider here and showed that the following performance measures can be expressed as functions of quantities having the form in (1).

Example 1 *Time-Average Variance Constant.* Consider estimating

$$\sigma^2 = \frac{E[Y(f; 1)^2] - 2rE[Y(f; 1)\tau(1)] + r^2E[\tau(1)^2]}{E[\tau(1)]}, \tag{4}$$

where $f : S \rightarrow \Re$ is a “reward” function,

$$r = \frac{E[Y(f; 1)]}{E[\tau(1)]} \tag{5}$$

is a steady-state average reward, and for $k \geq 1$,

$$\begin{aligned} \tau(k) &= T_1(k) - T_1(k-1), \\ Y(f; k) &= \sum_{n=T_1(k-1)}^{T_1(k)-1} f(X_n). \end{aligned}$$

The constant σ^2 is known as the time-average variance constant and arises in the central limit theorem for the time-average reward:

$$n^{1/2}(r_n - r) \xrightarrow{D} N(0, \sigma^2),$$

as $n \rightarrow \infty$, where $r_n = \frac{1}{n} \sum_{k=0}^n f(X_k)$ and \xrightarrow{D} denotes convergence in distribution; see Theorem 2.3 of Shedler (1993) for details. Observe that in (4), the terms $E[Y(f; 1)^2]$, $E[Y(f; 1)\tau(1)]$, and $E[\tau(1)^2]$ each have the form given in (1).

Example 2 *Derivative Estimation.* Suppose we want to estimate the derivative of r defined in (5) with respect to some system parameter λ , and assume that $f = f_\lambda$ is continuously differentiable in λ . For example, r may be the steady-state availability of a reliability system, and we want to compute its derivative with respect to the failure rate λ of some component. Assume that the random times $T_1(0), T_1(1), \dots$ do not depend on λ (but their distribution may depend on λ). Differentiating the ratio formula in (5) yields

$$\frac{\partial r}{\partial \lambda} = \frac{(\partial E[Y(f; 1)]) E[\tau(1)] - E[Y(f; 1)] (\partial E[\tau(1)])}{E^2[\tau(1)]}, \tag{6}$$

where we use the notation ∂A to denote the derivative of A with respect to λ . Under appropriate regularity conditions, applying the likelihood-ratio method of differentiation gives

$$\partial E[Y(f; 1)] = E[Y(f; 1) \partial L(1)] + E[\partial Y(f; 1)] \quad (7)$$

and

$$\partial E[\tau(1)] = E[\tau(1) \partial L(1)], \quad (8)$$

where

$$\partial Y(f; k) = \sum_{n=T_1(k-1)}^{T_1(k)-1} \partial f(X_n)$$

for $k \geq 1$, and $\partial L(1)$ is the derivative of the likelihood ratio; see Glynn (1990) for more details. In our context with X a discrete-time Markov chain on a countable state space S , the derivative of the likelihood ratio over the k th T_1 -cycle is

$$\partial L(k) = \sum_{j=T_1(k-1)}^{T_1(k)-1} \frac{\partial R_{X_j, X_{j+1}}}{R_{X_j, X_{j+1}}}.$$

Note that $E[Y(f; 1) \partial L(1)]$ in (7) and $E[\tau(1) \partial L(1)]$ in (8) each have the form given in (1).

Example 3 *Low-Bias Ratio Estimation.* Suppose our goal is to estimate some steady-state performance measure r having the form in (5). The standard ratio estimator based on \bar{X} is then

$$\hat{r}(m) = \frac{\sum_{k=1}^m Y(f; k)}{\sum_{k=1}^m \tau(k)}.$$

One can show that

$$E[\hat{r}(m)] - r = -\frac{1}{m} \frac{E[Y(f; 1) \tau(1)] - r E[\tau(1)^2]}{E^2[\tau(1)]} + O(m^{-2}); \quad (9)$$

see p. 104 of Shedler (1993). Hence, the standard estimator $\hat{r}(m)$ is typically biased, and various techniques have been introduced to try to reduce the bias. One approach proposed by Tin (1965) is to modify the standard estimator by adding to it terms that estimate the first-order bias term in (9). Observe that $E[Y(f; 1) \tau(1)]$ and $E[\tau(1)^2]$ in (9) both have the form given in (1).

3 IDENTITIES FOR PROCESSES WITH MULTIPLE REGENERATION SEQUENCES

We now derive a representation of the performance measure α that explicitly exploits the multi-regeneration structure of X . To simplify the development, we specialize the derivation to the setting when the reward functions f_U and f_V in (2) and (3), respectively, depend only on the first parameter; i.e., $f_U, f_V : S \rightarrow \mathfrak{R}$.

For $1 \leq i, j \leq s$ and $k \geq 0$, let

$$\vec{U}_{i,j}(k) = \sum_{l=T_i(k)}^{\inf\{n>T_i(k):X_n=x_j\}-1} f_U(X_l),$$

$$\vec{V}_{i,j}(k) = \sum_{l=T_i(k)}^{\inf\{n>T_i(k):X_n=x_j\}-1} f_V(X_l)$$

denote the cumulative rewards over the k th x_i to x_j meander (or excursion if $x_i = x_j$). By an x_i to x_j meander, we mean a path segment starting in state x_i and going to the next visit to x_j . With this notation, our goal can be expressed as estimating $\alpha = E[\vec{U}_{11}(1) \vec{V}_{11}(1)]$.

For $1 \leq i, j \leq s$, set

$$\begin{aligned} \beta_{i,j}^U &= E[\vec{U}_{i,j}(1)], \\ \beta_{i,j}^V &= E[\vec{V}_{i,j}(1)], \\ \beta_{i,j}^{UV} &= E[\vec{U}_{i,j}(1) \vec{V}_{i,j}(1)], \end{aligned}$$

and set

$$\begin{aligned} \gamma_{i,j}^U &= E_i \left[\sum_{l=0}^{T(1)-1} f_U(X_l) \mid X_{T(1)} = x_j \right], \\ \gamma_{i,j}^V &= E_i \left[\sum_{l=0}^{T(1)-1} f_V(X_l) \mid X_{T(1)} = x_j \right], \\ \gamma_{i,j}^{UV} &= E_i \left[\sum_{l=0}^{T(1)-1} f_U(X_l) \sum_{l=0}^{T(1)-1} f_V(X_l) \mid X_{T(1)} = x_j \right], \end{aligned}$$

where E_i is the expectation for the chain starting in state $x_i \in S_0$. Then

$$\begin{aligned} \beta_{i,j}^{UV} &= E_i \left[\vec{U}_{i,j}(1) \vec{V}_{i,j}(1) \right] \\ &= \sum_{k=1}^s Q_{i,k} E_i \left[\vec{U}_{i,j}(1) \vec{V}_{i,j}(1) \mid X_{T(1)} = x_k \right] \\ &= \sum_{k=1}^s Q_{i,k} E_i \left[\left(\sum_{l=0}^{T(1)-1} f_U(X_l) + \vec{U}_{k,j}(1) 1_{(k \neq j)} \right) \right. \\ &\quad \left. \times \left(\sum_{l=0}^{T(1)-1} f_V(X_l) + \vec{V}_{k,j}(1) 1_{(k \neq j)} \right) \mid X_{T(1)} = x_k \right] \\ &= \sum_{k=1}^s Q_{i,k} \gamma_{i,k}^{UV} \\ &\quad + \sum_{k \neq j} Q_{i,k} \left(\gamma_{i,k}^U \beta_{k,j}^V + \gamma_{i,k}^V \beta_{k,j}^U + \beta_{k,j}^{UV} \right). \end{aligned}$$

For $1 \leq i, j \leq s$, set

$$A_{i,j} = \sum_{k=1}^s Q_{i,k} \gamma_{i,k}^{UV} + \sum_{k \neq j} Q_{i,k} \left(\gamma_{i,k}^U \beta_{k,j}^V + \gamma_{i,k}^V \beta_{k,j}^U \right),$$

and let $A = (A_{i,j} : 1 \leq i, j \leq s)$. Then letting ${}_1Q$ denote the matrix Q but with the elements of the first column changed to zero, we have that

$$(I - {}_1Q)\beta_{\cdot,1}^{UV} = A_{\cdot,1},$$

where we use the notation that $M_{\cdot j}$ denotes the j th column of a matrix $M = (M_{i,j})$. Similar calculations yield

$$(I - {}_1Q)\beta_{\cdot,1}^U = B^U \tag{10}$$

and

$$(I - {}_1Q)\beta_{\cdot,1}^V = B^V, \tag{11}$$

where $B^U = (B_i^U : i \in S_0)$, $B^V = (B_i^V : i \in S_0)$, and $B^{UV} = (B_i^{UV} : i \in S_0)$ are vectors with

$$\begin{aligned} B_i^U &= \sum_{k=1}^s Q_{i,k} \gamma_{i,k}^U, \\ B_i^V &= \sum_{k=1}^s Q_{i,k} \gamma_{i,k}^V, \\ B_i^{UV} &= \sum_{k=1}^s Q_{i,k} \gamma_{i,k}^{UV}. \end{aligned}$$

In particular,

$$\alpha = \beta_{11}^{UV} = (I - {}_1Q)_1^{-1} A_{\cdot,1}.$$

Now

$$\begin{aligned} (I - {}_1Q)_1^{-1} &= \sum_{n=0}^{\infty} {}_1Q_{1,j}^n \\ &= E_1 \left[\sum_{n=0}^{\infty} 1_{\{W_n = x_j, T_1(1) > n\}} \right] = v_j/v_1, \end{aligned}$$

where $1_{\{\cdot\}}$ is the indicator function of the set $\{\cdot\}$, and so

$$\begin{aligned} \beta_{11}^{UV} &= \frac{1}{v_1} \sum_{k=1}^s v_k A_{k,1} \\ &= \frac{1}{v_1} \sum_{k=1}^s v_k \left(\sum_{j=1}^s Q_{k,j} \gamma_{k,j}^{UV} \right. \\ &\quad \left. + \sum_{j \neq 1} Q_{k,j} \left(\gamma_{k,j}^U \beta_{j,1}^V + \gamma_{k,j}^V \beta_{j,1}^U \right) \right). \end{aligned}$$

Using expressions for $\beta_{\cdot,1}^U$ and $\beta_{\cdot,1}^V$ derived from (10) and (11), respectively, we can rewrite the last expression as

$$\begin{aligned} \alpha &= \frac{1}{v_1} \left(v B^{UV} + v({}_1Q \circ \gamma^U)(I - {}_1Q)^{-1} B^V \right. \\ &\quad \left. + v({}_1Q \circ \gamma^V)(I - {}_1Q)^{-1} B^U \right), \tag{12} \end{aligned}$$

where $(A \circ B)_{i,j} = A_{i,j} B_{i,j}$.

4 MIXED-MOMENT ESTIMATORS

We can use (12) to construct an estimator $\hat{\alpha}_*$ for α as follows. First generate a sample path $\vec{X} = (X_k : 0 \leq k < T_1(m))$, with $X_0 = x_1$. We then form the estimator $\hat{\alpha}_*(\vec{X})$ by inserting the sample quantities into the formula (12) for α . More precisely, for $1 \leq i \leq s$, let $\zeta_i(0) = \inf\{k \geq 0 : W_k = x_i\}$ and $\zeta_i(n+1) = \inf\{k > \zeta_i(n) : W_k = x_i\}$, $n \geq 0$, be the sequence of hitting times to state x_i for the W chain. For $1 \leq i, j \leq s$, define $N_{i,j} = \sum_{k=0}^{\zeta_1(m)-1} 1_{\{W_k = x_i, W_{k+1} = x_j\}}$, which is the number of (x_i, x_j) transitions that the W chain makes up to the $(m+1)$ th hit to state x_1 . Define \hat{Q} as the sample transition probability matrix of W , with $\hat{Q}_{i,j} = N_{i,j}/N_i$, where $N_i = \sum_j N_{i,j}$ is the total number of transitions out of state x_i . Form the natural sample averages

$\hat{\gamma}^U$, $\hat{\gamma}^V$, and $\hat{\gamma}^{UV}$, and let $\hat{v}_k = N_k/N$, where $N = \sum_i N_i$ is the total number of transitions of the W chain. Define

$$\hat{B}_i^U = \sum_{k=1}^s \hat{Q}_{i,k} \hat{\gamma}_{i,k}^U,$$

$$\hat{B}_i^V = \sum_{k=1}^s \hat{Q}_{i,k} \hat{\gamma}_{i,k}^V,$$

and

$$\hat{B}_i^{UV} = \sum_{k=1}^s \hat{Q}_{i,k} \hat{\gamma}_{i,k}^{UV}.$$

Then calculate $\hat{\beta}^U$ and $\hat{\beta}^V$ as the solution to the equations

$$(I - {}_1\hat{Q})\hat{\beta}_1^U = \hat{B}^U$$

and

$$(I - {}_1\hat{Q})\hat{\beta}_1^V = \hat{B}^V.$$

Next, for $1 \leq i, j \leq s$, let

$$\hat{A}_{i,j} = \sum_{k=1}^s \hat{Q}_{i,k} \hat{\gamma}_{i,k}^{UV} + \sum_{k \neq j} \hat{Q}_{i,k} \left(\hat{\gamma}_{i,k}^U \hat{\beta}_{k,j}^V + \hat{\gamma}_{i,k}^V \hat{\beta}_{k,j}^U \right).$$

Finally, set $\hat{\alpha}_*(\vec{X}) = \hat{\beta}_{11}^{UV}$, where $\hat{\beta}_1^{UV}$ is the solution to

$$(I - {}_1\hat{Q})\hat{\beta}_1^{UV} = \hat{A}_1.$$

The new estimator is unbiased; that is, $E[\hat{\alpha}_*(\vec{X})] = \alpha$. To see why this is true, it is useful to take another viewpoint in deriving the estimator $\hat{\alpha}_*(\vec{X})$. The cycles between successive hits to a fixed state are i.i.d.; therefore, permuting the cycles between visits to, say, state x_i does not change the distribution of the path (see Calvin and Nakayama 1998). Therefore, if we form a ‘‘standard’’ regenerative estimator $\hat{\alpha}(\vec{X})$ based on a sample path \vec{X} of m standard regenerative T_1 -cycles, then

$$E[\hat{\alpha}(\vec{X})] = E[\hat{\alpha}(\vec{X}')],$$

where \vec{X}' is obtained from \vec{X} by permuting cycles for any collection of states in S_0 . It turns out that $\hat{\alpha}_*(\vec{X})$ is precisely the average of $\hat{\alpha}$ over all such cycle permutations. Therefore, it has the same expectation as $\hat{\alpha}(\vec{X})$, which is unbiased. (The above argument for unbiasedness relies on the sample path \vec{X} consisting of exactly m T_1 -cycles, m fixed.)

5 SMALL-SAMPLE OPTIMALITY

Now assume that $|S| < \infty$ and $S_0 = S = \{1, 2, 3, \dots, s\}$; i.e., we will construct our estimator using all states in the state space S . Thus, we take the state $x_i \in S_0$ to be $x_i = i$ for $i = 1, 2, \dots, s = |S|$, and so the T_i sequence consists of the successive hitting times to state i , $i \in S$. We will consider our estimator based on a sample path $\vec{X} = (X_k : 0 \leq k < T_1(m))$, with $X_0 = 1$ and $m \geq 2$ fixed. In this setting the estimator $\hat{\alpha}_*(\vec{X})$ is the uniform minimum variance unbiased estimator (UMVUE) of α .

To state this result precisely, we need some additional definitions. Define $\mathcal{F}_{T_1(m)} = \sigma(X_0, X_1, \dots, X_{T_1(m)})$, which is the sigma-field generated by the process X up to time $T_1(m)$. Define \mathcal{R}_+ to be the class of transition probability matrices R with $R_{i,j} > 0$ for all $i, j \in S$. Finally define the family \mathcal{P}_m of probability measures P on $\mathcal{F}_{T_1(m)}$ induced by a transition probability matrix $R \in \mathcal{R}_+$.

Our goal is to estimate $\alpha_P = E_P[U(1)V(1)]$, where we use a subscript P to emphasize the measure used to generate the DTMC. Then the following result holds, which is a form of small-sample optimality; see Calvin, Glynn, and Nakayama (1999) for the proof.

Theorem 1 *The permuted estimator $\hat{\alpha}_*(\vec{X})$ is the UMVUE of α_P over $P \in \mathcal{P}_m$, $m \geq 2$.*

Remarks:

1. Theorem 1 establishes that $\hat{\alpha}_*(\vec{X})$ has the smallest variance of any unbiased estimator of α_P over $P \in \mathcal{P}_m$. However, the result can be extended to showing that $\hat{\alpha}_*(\vec{X})$ is the unique (w.p. 1) unbiased estimator that uniformly minimizes the risk for any convex loss function; see p. 88 of Lehmann and Casella (1998).
2. Observe that (4), (6), and the first-order bias term in (9) are all nonlinear functions of means. We obtain the standard estimators of these quantities by replacing the means with their respective sample means. Thus, the standard estimators of these quantities are nonlinear functions of sample means and so are typically biased. Thus, even though our estimator $\hat{\alpha}_*(\vec{X})$ is unbiased for α_P , our overall new estimators for (4), (6), and the first-order bias term in (9) are typically biased.
3. For simplicity we assumed in Theorem 1 that the probability transition matrix R used to generate the sample path was strictly positive; i.e., $R \in \mathcal{R}_+$. We can relax this assumption to allow for non-negative R as follows. Let $\Theta \subset S \times S$ such that for any pair of states $i, j \in S$, there exists a sequence of states $i_1 = i, i_2, i_3, \dots, i_n = j$, $n \geq 2$, such that $(i_k, i_{k+1}) \in \Theta$, $k = 1, 2, \dots, n - 1$. Now

consider the class \mathcal{R} of probability transition matrices R such that $R_{i,j} > 0$ if and only if $(i, j) \in \Theta$. Any $R \in \mathcal{R}$ is irreducible by our definition of Θ . Finally define the family \mathcal{P}_m of probability measures P on $\mathcal{F}_{T_1(m)}$ induced by a transition probability matrix $R \in \mathcal{R}$. Then with these changes, Theorem 1 holds.

6 EMPIRICAL RESULTS

We now present results from applying our proposed technique in simulations of two different models. We first consider a discrete-time Markov chain $X = (X_n : n = 0, 1, 2, \dots)$ on a finite state space $S = \{0, 1, \dots, s\}$ with transition probability matrix defined by $R_{i,i+1} = \lambda/(i+\lambda) = 1 - R_{i,i-1}$ for $0 < i < s$, and $R_{0,1} = R_{s,s-1} = 1$. This chain is the discrete-time version of the Erlang loss system. Our goal is to estimate the time-average variance constant σ^2 of X using the proposed technique with $S_0 = S$. We ran numerical experiments with $s = 15$ and $\lambda = s/2$.

Table 1 reports the results of simulations of 1,000 independent replications for various choices for the state $x_1 \in S_0$ used to determine the T_1 -cycles. In each replication we fixed the number of T_1 -cycles simulated. Since the expected length of the T_1 -cycles depends on the choice of the state x_1 , we adjusted the number of T_1 -cycles in a replication for each choice of x_1 so that the expected total number of transitions of the Markov chain is about 1,000,000. Thus, the results across the various rows are comparable. The columns labeled with “ $\hat{\alpha}_*$ ” and “ $\hat{\alpha}$ ” are the sample variances of the proposed and standard estimators, respectively, over the 1,000 replications. The last column contains the ratio of the estimated efficiency of the proposed estimator over that of the standard estimator, where the efficiency is defined as the inverse of the product of the work and variance, and the work is defined as the expected CPU time required to run the simulation and construct the estimator.

First observe that our proposed method can significantly reduce the variance and increase the efficiency; e.g., see the row for $x_1 = 0$ in Table 1. Also, note that the variance of the standard estimator varies quite a bit over the various choices for x_1 , showing that the standard estimator of σ^2 is quite sensitive to the choice of the regeneration sequence used to control the simulation. On the other hand, the variance of our proposed estimator does not seem to depend on the choice of x_1 .

We also ran experiments on a large queueing system. The system consisted of 8 stations, each with a single server and a queue that can hold up to 14 waiting customers. The service discipline at each station is first-come-first-served, and the service distribution is exponential with rate 1. If a customer arrives to a station and the queue is full, the customer immediately leaves the system. The interarrival distribution of customers to the system is exponential with

Table 1: Variances of Estimators of σ^2

x_1	$\hat{\alpha}_*$	$\hat{\alpha}$	Efficiency
0	0.10	130.55	1318.90
1	0.10	12.45	120.48
2	0.10	2.57	26.22
3	0.10	0.81	8.50
4	0.09	0.39	4.16
5	0.09	0.22	2.31
6	0.09	0.14	1.54
7	0.09	0.12	1.32

rate 5, and an arriving customer is equally likely to go to any of the stations. We simulated the process $X = (X_n : n = 0, 1, 2, \dots)$, which is the discrete-time version of this queueing system, where the state space of X is $S = \{(n_1, n_2, \dots, n_8) : 0 \leq n_i \leq 15, i = 1, 2, \dots, 8\}$. Note that $|S| = 16^8 \approx 4 \times 10^9$. The goal of our simulation experiment is to estimate the time-average variance constant σ^2 of the total number of customers in the system.

In this experiment we let S_0 be a strict subset of the state space S , where $S_0 = \{0, 3\}^8$ is the set of states (n_1, n_2, \dots, n_8) in which each station i has either $n_i = 0$ or $n_i = 3$; thus, $|S_0| = 256$. We let $x_1 = (0, 0, 0, 0, 0, 0, 0, 0)$, i.e., no customers in the system. We ran 1000 independent replications in which we constructed both the standard estimator and our proposed estimator. Each replication consisted of approximately 100,000 transitions. We obtained a 6-fold reduction in variance by using our proposed technique. The time required for the post-processing needed to construct our proposed estimator was negligible compared with the time required to generate the sample path. Thus, the increase in efficiency is also about 6-fold.

ACKNOWLEDGMENTS

The work of the first author was supported in part by the National Science Foundation under grants DMI-9500173 and DMI-9900117. The second author’s research was partially supported by Army Research Office Grant DAAG55-97-1-0377-P0001 and National Science Foundation Grant DMS-9704732-001. The work of the third author was supported in part by the National Science Foundation under grants DMI-9624469 and DMI-9900117.

REFERENCES

Calvin, J. M. and M. K. Nakayama. 1998. Using permutations in regenerative simulations to reduce variance. *ACM Transactions on Modeling and Computer Simulation* 8: 153–193.

- Calvin, J. M. and M. K. Nakayama. 1999a. Permuted product and importance-sampling estimators for regenerative simulations. Submitted.
- Calvin, J. M. and M. K. Nakayama. 1999b. Simulation of processes with multiple regeneration sequences. Submitted.
- Calvin, J. M. and M. K. Nakayama. 2000. Central limit theorems for permuted regenerative estimators. *Operations Research*, to appear.
- Calvin, J. M., P. W. Glynn, and M. K. Nakayama. 1999. The semi-regenerative method of simulation-output analysis. In preparation.
- Çinlar, E. 1975. *Introduction to Stochastic Processes*. Prentice-Hall, New Jersey.
- Glynn, P. W. 1990. Likelihood ratio derivative estimators for stochastic systems. *Comm. ACM* 33: 75–84.
- Lehmann, E. L., and G. Casella. 1998. *Theory of Point Estimation, Second Edition*. New York: Springer-Verlag.
- Shedler, G. S. 1993. *Regenerative Stochastic Simulation*. San Diego: Academic Press.
- Tin, M. 1965. Comparison of some ratio estimates. *Journal of the American Statistical Association* 60: 294–307.
- ests include applied probability, statistics, simulation-output analysis, variance-reduction techniques, gradient estimation, rare-event simulation, and random-variate generation.

AUTHOR BIOGRAPHIES

JAMES M. CALVIN is an assistant professor in the Department of Computer and Information Science at the New Jersey Institute of Technology. He received a Ph.D. in Operations Research from Stanford University. His research interests include global optimization and simulation output analysis.

PETER W. GLYNN received his Ph.D. from Stanford University, after which he joined the faculty of the Department of Industrial Engineering at the University of Wisconsin-Madison. In 1987, he returned to Stanford, where he is the Thomas Ford Professor of Engineering in the Department of Engineering-Economic Systems and Operations Research. He was a co-winner of the 1993 Outstanding Simulation Publication Award sponsored by the TIMS College on Simulation. His research interests include discrete-event simulation, computational probability, queueing, and general theory for stochastic systems.

MARVIN K. NAKAYAMA is an associate professor in the Department of Computer and Information Science at the New Jersey Institute of Technology. He received a Ph.D. in Operations Research from Stanford University. He won second prize in the 1992 George E. Nicholson Student Paper Competition sponsored by INFORMS and is a recipient of a CAREER Award from the National Science Foundation. He is an associate editor for *ACM Transactions on Modeling and Computer Simulation*. His research inter-