A NEW VARIANCE-REDUCTION TECHNIQUE FOR REGENERATIVE SIMULATIONS OF MARKOV CHAINS

James M. Calvin
Marvin K. Nakayama

Department of Computer and Information Science
New Jersey Institute of Technology
Newark, NJ 07102-1982, U.S.A.

ABSTRACT

We propose a new estimator for some performance measures obtained from a regenerative simulation of a discrete-time Markov chain. Our new estimator is based on the idea of generating (uniform) random permutations of cycles corresponding to a certain state, and it has no larger (and typically smaller) mean squared error than the standard estimator. We show that our method can be used to derive a new estimator for the time-average variance parameter of a regenerative simulation.

1 INTRODUCTION

The regenerative method is a simulation output analysis technique for estimating performance measures of regenerative stochastic systems; see Crane and Iglehart (1975). The basis of the approach is to divide the sample path into independent and identically distributed (i.i.d.) segments (cycles), where the endpoints of the segments are often determined by hitting times to a fixed return state. Many stochastic systems have been shown to be regenerative (see Shedler 1993), and the regenerative method results in asymptotically valid confidence intervals.

In this paper we propose a new estimator for a performance measure α of a discrete-time Markov chain obtained from a regenerative simulation of a fixed number of regenerative cycles for a fixed return state v . Our new estimator is based on the idea of generating (uniform) random permutations of cycles corresponding to a state $w \neq v$, and the new estimator has the same expected value as the standard estimator and no greater (and typically strictly less) variance; thus, it has no larger mean squared error. We develop our method for irreducible discrete-time Markov chains defined on a finite state space, but there are natural extensions to Markov chains with countable state spaces and to more general regenerative

processes. These ideas are explored more fully in Calvin and Nakayama (1997).

While our method has no effect on the standard regenerative ratio estimator for certain steady-state performance measures, the basic technique can still be beneficially applied to a rich class of other performance measures. One example is the standard regenerative variance estimator. Hence, our estimator will result in a variance estimator having no more variability than the standard one. This is important because one measure of the quality of a particular output-analysis methodology is the variability of the half width of the resulting confidence interval (Glynn and Iglehart 1987), which is largely influenced by the variance of the variance estimator.

The rest of the paper is organized as follows. In Section 2, we discuss the standard estimator of α and describe how to construct our new estimator. Section 3 shows how these results can be used to derive a new estimator of the variance parameter arising in a regenerative simulation. We analyze the storage and computational costs of our new estimator in Section 4. We present in Section 5 the results of some numerical experiments comparing our new regenerative-variance estimator to the standard one.

2 OUR NEW ESTIMATORS

Let $X = (X_n : n = 0, 1, 2, \dots)$ be an irreducible aperiodic discrete-time Markov chain taking values on a finite state space S . Suppose our goal is to estimate a performance measure α of the Markov chain by simulating some fixed number of regenerative cycles for a fixed return state. More specifically, for a state $v \in S$, set $X_0 = v$, and define the successive return times

$$\begin{aligned} T_v(0) &= 0, \\ T_v(n+1) &= \inf\{m > T_v(n) : X_m = v\} \end{aligned}$$

for $n \geq 0$. Let $\vec{X} = \{X_n : n = 0, 1, 2, \dots, T_v(m_v)\}$ be a sample path of m_v regenerative v -cycles of our

discrete-time Markov chain, where m_v is fixed. Now we define the standard estimator of α based on the sample path $\vec{X}$ to be

$$\hat{\alpha}(\vec{X}) = h(\vec{X}) \quad (1)$$

where h is some function. This general framework includes many performance measures of interest. We describe two examples below; another performance measure is treated in Calvin and Nakayama (1997).

Example 1 Suppose

$$\alpha = E \left[\left(\sum_{j=T_v(0)}^{T_v(1)-1} g(X_j) \right)^p \right]$$

for some function $g : S \rightarrow \mathfrak{R}$, where $p \geq 1$. Then we can define $h(\vec{X})$ by

$$h(\vec{X}) = \frac{1}{m_v} \sum_{k=1}^{m_v} Y(g; k)^p = \hat{\alpha}(\vec{X}),$$

where

$$Y(g; k) \equiv Y_v(g; k) = \sum_{j=T_v(k-1)}^{T_v(k)-1} g(X_j).$$

Note that $\hat{\alpha}(\vec{X})$ is an unbiased estimator of α in this example. ■

Example 2 Suppose that

$$\alpha = \sigma^2 = \frac{E[Z(1)^2]}{E[\tau(1)]},$$

where

$$\begin{aligned} \tau(k) &\equiv \tau_v(k) = T_v(k) - T_v(k-1), \\ Z(f; k) &= Y(f; k) - r\tau(k), \end{aligned}$$

$f : S \rightarrow \mathfrak{R}$ is some “cost” function, $r = \pi f \equiv \sum_{x \in S} \pi_x f(x)$, and $\pi = (\pi_x : x \in S)$ is the stationary distribution of X . Observe that α in this case is the time-average variance parameter of the chain. (More details are given in Section 3.) Then we can define $h(\vec{X})$ by

$$h(\vec{X}) = \frac{\sum_{k=1}^{m_v} (Y(f; k) - \hat{r}\tau(k))^2}{\sum_{k=1}^{m_v} \tau(k)} = \hat{\alpha}(\vec{X}), \quad (2)$$

where

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

Note that $\hat{\alpha}(\vec{X})$ is the standard regenerative estimator of σ^2 . We will return to this particular example in Section 3. ■

Our goal now is to create a new estimator for α , which we do as follows. Given the original sample path $\vec{X}$, we begin by constructing a new sample path $\vec{X}'$ from $\vec{X}$ such that $\vec{X}' \stackrel{\mathcal{D}}{=} \vec{X}$, where “ $\stackrel{\mathcal{D}}{=}$ ” denotes equality in distribution. This is done by first taking the original sample path $\vec{X}$ and fixing a new state $w \neq v$. The state w is hit a random number of times M_w within the original path $\vec{X}$, and we can now look at the $(M_w - 1)$ w -cycles in the path. We generate a uniform random permutation of the $(M_w - 1)$ w -cycles within the path $\vec{X}$, and this gives us our new sample path $\vec{X}'$. More specifically, define

$$M_w = |\{0 \leq n \leq t : X_n = w\}|,$$

where for notational simplicity we define $t = T_v(m_v)$. If $M_w \leq 2$, then let $\vec{X}' = \vec{X}$. If $M_w \geq 3$, then define

$$T_w(1) = \min\{n > 0 : X_n(\vec{X}) = w\}$$

and

$$T_w(k) = \min\{n > T_w(k-1) : X_n(\vec{X}) = w\}$$

for $k = 2, 3, \dots, M_w$. Hence, we can break up the path $\vec{X}$ into

$$\vec{X} = (\vec{X}_1, C(1), C(2), \dots, C(M_w - 1), \vec{X}_2),$$

where

$$\begin{aligned} \vec{X}_1 &= \{X_n : 0 \leq n < T_w(1)\}, \\ \vec{X}_2 &= \{X_n : T_w(M_w) \leq n \leq t\}, \end{aligned}$$

and

$$C(k) = \{X_n : T_w(k) \leq n < T_w(k+1)\},$$

$k = 1, 2, \dots, M_w - 1$, which is the k th w -cycle of the original path $\vec{X}$. Let $\zeta(1), \zeta(2), \dots, \zeta(M_w - 1)$ be a uniform random permutation of $1, 2, \dots, M_w - 1$. Then we define our new path $\vec{X}' = \{X'_n : 0 \leq n \leq t\}$ to be

$$\vec{X}' = (\vec{X}_1, C(\zeta(1)), C(\zeta(2)), \dots, C(\zeta(M_w - 1)), \vec{X}_2),$$

which is the original path $\vec{X}$ with the w -cycles permuted. Hence, $\vec{X}'$ starts in state v and has exactly m_v v -cycles. It can be shown that $\vec{X} \stackrel{\mathcal{D}}{=} \vec{X}'$; see Calvin and Nakayama (1997) for a proof.

Now for the new sample path $\vec{X}'$, we can calculate

$$\hat{\alpha}(\vec{X}') = h(\vec{X}'),$$

which is just the estimator obtained from the new sample path $\vec{X}'$. We finally define our new estimator for α to be

$$\tilde{\alpha}(\vec{X}) = E_*[\hat{\alpha}(\vec{X}')], \quad (3)$$

where E_* is the conditional expectation operator with respect to a random (uniform) permutation of w -cycles (as was done when constructing the path $\vec{X}'$ from $\vec{X}$) given the original sample path $\vec{X}$. The following result, proved in Calvin and Nakayama (1997), then holds:

Theorem 1 *Let $v, w \in S$ be two distinct states and construct the estimator $\tilde{\alpha}(\vec{X})$ defined by (3). Then $E[\tilde{\alpha}(\vec{X})] = E[\hat{\alpha}(\vec{X})]$ and*

$$\text{Var}[\tilde{\alpha}(\vec{X})] \leq \text{Var}[\hat{\alpha}(\vec{X})], \quad (4)$$

and so the mean squared error of our new estimator $\tilde{\alpha}(\vec{X})$ is no greater than that of the original estimator $\hat{\alpha}(\vec{X})$. Strict inequality is obtained in (4) unless $P\{h(\vec{X}') = h(\vec{X})\} = 1$, where P is the probability measure corresponding to generating the original sample path $\vec{X}$ and generating the permutation for constructing $\vec{X}'$.

An outline of the proof of Theorem 1 is as follows. Using the Markov property, it can be shown that $\vec{X}' \stackrel{D}{=} \vec{X}$. It then follows that

$$E[\tilde{\alpha}(\vec{X})] = E[E_*[\hat{\alpha}(\vec{X}')]] = E[\hat{\alpha}(\vec{X}')] = E[\hat{\alpha}(\vec{X})].$$

Also, decomposing the variance by conditioning on $\vec{X}$ establishes the variance reduction.

In Theorem 1 we see that there is no variance reduction when for every possible original sample path $\vec{X}$, the value of the function h in (1) is unaffected by permutations of the w -cycles. For example, this is the case in Example 1 with $p = 1$ since

$$h(\vec{X}) = \frac{1}{T_v(m_v)} \sum_{j=0}^{T_v(m_v)-1} g(X_j) = h(\vec{X}'),$$

and so $\tilde{\alpha}(\vec{X}) = \hat{\alpha}(\vec{X})$. Similarly, by choosing $g(x) \equiv 1$, we see that permuting w -cycles does not alter the estimator for $E[\tau(1)]$. Thus, our method has no effect on the standard ratio estimator for steady-state performance measures α that can be expressed as $\alpha = E[Y(g; 1)]/E[\tau(1)]$.

However, for $p \neq 1$, we have that $h(\vec{X}) \neq h(\vec{X}')$ in general, and so typically $\tilde{\alpha}(\vec{X}) \neq \hat{\alpha}(\vec{X})$. Also, we usually have that the standard time-average variance estimator in Example 2 for a regenerative simulation will differ from the new estimator defined by (3).

Our new estimator $\tilde{\alpha}(\vec{X})$ is only of use in practice when we can explicitly compute the conditional expectation in (3). One situation where this can be done is in Example 1 for $p = 2$; i.e., when

$$\alpha = E[Y(g; 1)^2] \quad (5)$$

and our standard estimator of α is

$$\hat{\alpha}(\vec{X}) = \frac{1}{m_v} \sum_{k=1}^{m_v} Y(k)^2, \quad (6)$$

where we have dropped the dependence of Y on g to simplify the notation. Our new estimator of α is then

$$\tilde{\alpha}(\vec{X}) = E_* \left[\frac{1}{m_v} \sum_{k=1}^{m_v} Y'(k)^2 \right], \quad (7)$$

where $Y'(k)$ is the same as $Y(k)$ except that it is for the sample path $\vec{X}'$ rather than $\vec{X}$.

Now to explicitly calculate (7) in this particular setting, we need some new notation. For $v, w \in S$, let $H(v; w) \subset \{1, 2, \dots, m_v\}$ denote the set of indices of the v -cycles that hit state w , and define the complementary set $J(v; w) = \{1, 2, \dots, m_v\} - H(v; w)$. Let $h_{vw} = |H(v; w)|$. For $k \in H(v; w)$, define $T'_w(k) = \inf\{n > T_v(k-1) : X_n = w\}$, which is the first time that state w is hit after the $(k-1)$ st time state v is hit. Similarly define $\tilde{T}_w(k) = \sup\{n < T_v(k) : X_n = w\}$, which is the last time state w is hit before the k th regeneration for state v . Then, for $k \in H(v; w)$, we let

$$Y_{vw}(k) = \sum_{j=T_v(k-1)}^{T'_w(k)-1} g(X_j),$$

which is the contribution to $Y(k)$ until state w is hit, and let

$$Y_{wv}(k) = \sum_{j=\tilde{T}_w(k)}^{T_v(k)-1} g(X_j),$$

which is the contribution to $Y(k)$ from the last time state w is hit in the k th v -cycle until the end of the cycle. Also, for $l \in J(w; v)$, let

$$Y_{ww}(l) = \sum_{j=T_w(l-1)}^{T_w(l)-1} g(X_j),$$

which is the sum of the $g(X_j)$ over the l th w -cycle that does not hit state v . Also, define

$$\bar{Y}_{vw} = \frac{1}{h_{vw}} \sum_{k \in H(v; w)} Y_{vw}(k)$$

and

$$\bar{Y}_{wv} = \frac{1}{h_{vw}} \sum_{k \in H(v; w)} Y_{wv}(k).$$

Finally, we define β_l to be the l th smallest element of the set $H(v; w)$ for $l = 1, 2, \dots, h_{vw}$, and define $\beta_0 = \beta_{h_{vw}}$. For $k = \beta_l \in H(v; w)$ for some $l =$

$1, 2, \dots, h_{vw}$, define $\psi(k) = \beta_{l-1}$; i.e., $\psi(k)$ is the index in $H(v; w)$ that occurs just before k if k is not the first index and is the last element in $H(v; w)$ if k is the first element. Then the following is proved in Calvin and Nakayama (1997):

Theorem 2 Suppose we want to estimate α defined in (5). Then, our new estimator is given by $\tilde{\alpha}(\vec{X}) = \hat{\alpha}(\vec{X})$ if $M_w < 3$, and otherwise by

$$\begin{aligned} & \frac{1}{m_v} \left(\sum_{k \in J(v; w)} Y(k)^2 + \sum_{k \in H(v; w)} [Y_{vw}(k)^2 + Y_{wv}(k)^2] \right. \\ & + \sum_{k \in H(v; w)} \frac{2Y_{vw}(k)}{h_{vw} - 1} \left(\sum_{j \in H(v; w)} Y_{wv}(j) - Y_{wv}(\psi(k)) \right) \\ & + \sum_{k \in J(w; v)} Y_{ww}(k)^2 + 2(\bar{Y}_{vw} + \bar{Y}_{wv}) \sum_{k \in J(w; v)} Y_{ww}(k) \\ & \left. + \frac{2}{1 + h_{vw}} \sum_{\substack{l, m \in J(w; v) \\ l \neq m}} Y_{ww}(l) Y_{ww}(m) \right) \end{aligned} \quad (8)$$

The estimator $\tilde{\alpha}(\vec{X})$ satisfies $E[\tilde{\alpha}(\vec{X})] = \alpha$, and also $\text{Var}(\tilde{\alpha}(\vec{X})) \leq \text{Var}(\hat{\alpha}(\vec{X}))$ when $\hat{\alpha}(\vec{X})$ is the standard estimator of α as defined in (6).

3 NEW ESTIMATOR FOR THE REGENERATIVE VARIANCE

We can use Theorem 2 to construct a new estimator for the variance parameter in a regenerative simulation of a discrete-time Markov chain X . We start by first giving a more complete explanation of Example 2 in Section 2. Recall that X possesses a unique stationary distribution π . Given a cost function $f : S \rightarrow \mathfrak{R}$, the goal of the regenerative method of steady-state simulation is to estimate the parameter $r \equiv \pi f$. To estimate r , we simulate X for n transitions, and form the estimator

$$r_n = \frac{1}{n+1} \sum_{i=0}^n f(X_i).$$

There exists a finite positive constant σ such that

$$n^{1/2} \frac{(r_n - r)}{\sigma} \xrightarrow{D} N(0, 1) \quad (9)$$

as $n \rightarrow \infty$, where $\xrightarrow{D}$ denotes convergence in distribution. The constant σ^2 is called the *time-average variance* of X . Given the central limit theorem described by (9), construction of confidence intervals for r therefore effectively reduces to developing a consistent estimator for σ^2 . The quality of the resulting confidence

interval is largely dependent upon the quality of the associated time-average variance estimator.

The standard consistent estimator of σ^2 is $\hat{\sigma}^2(\vec{X}) = \hat{\alpha}(\vec{X})$ defined in (2). Note that $\hat{\sigma}^2(\vec{X})$ can be expressed as

$$\hat{\sigma}^2(\vec{X}) = \frac{\sum_{k=1}^{m_v} Y(f - \hat{r}; k)^2}{\sum_{k=1}^{m_v} \tau(k)}.$$

Now we define our new estimator $\tilde{\sigma}^2(\vec{X})$ to be the conditional expectation of $\hat{\sigma}^2(\vec{X}')$ with respect to a random permutation of w -cycles, given the original sample $\vec{X}$. Hence, letting $\hat{r}'$, $Y'(f - \hat{r}'; k)$, and $\tau'(k)$ be the corresponding values of $\hat{r}$, $Y(f - \hat{r}; k)$, and $\tau(k)$ for the sample path $\vec{X}'$, we get that

$$\begin{aligned} \tilde{\sigma}^2(\vec{X}) &= E_* \left[\frac{\sum_{k=1}^{m_v} Y'(f - \hat{r}'; k)^2}{\sum_{k=1}^{m_v} \tau'(k)} \right] \\ &= \frac{E_* [\sum_{k=1}^{m_v} Y'(f - \hat{r}'; k)^2]}{\sum_{k=1}^{m_v} \tau(k)}, \end{aligned}$$

since $\sum_{k=1}^{m_v} \tau'(k) = T_v(m_v) = \sum_{k=1}^{m_v} \tau(k)$ is independent of the permutation of w -cycles. Also, observe that

$$\hat{r}' = \frac{1}{T_v(m_v)} \sum_{j=0}^{T_v(m_v)-1} f(X_j) = \hat{r}$$

is independent of the permutation of w -cycles, so

$$\tilde{\sigma}^2(\vec{X}) = \frac{E_* [\sum_{k=1}^{m_v} Y'(f - \hat{r}; k)^2]}{\sum_{k=1}^{m_v} \tau(k)}; \quad (10)$$

i.e., we can replace $\hat{r}'$ with $\hat{r}$. Therefore, we can compute the numerator of the right-hand side of (10) using (8) with the function $g = f - \hat{r}$. It follows from Theorem 1 that $E[\tilde{\sigma}^2(\vec{X})] = E[\hat{\sigma}^2(\vec{X})]$ and also $\text{Var}[\tilde{\sigma}^2(\vec{X})] \leq \text{Var}[\hat{\sigma}^2(\vec{X})]$.

4 STORAGE AND COMPUTATION COST

We now discuss the implementation issues associated with constructing our new estimator $\tilde{\alpha}(\vec{X})$ given in (8) for the case when α is defined in (5). First observe that the second line of (8), without the factor $2/(h_{vw} - 1)$, satisfies

$$\begin{aligned} & \sum_{k \in H(v; w)} Y_{vw}(k) \left(\sum_{j \in H(v; w)} Y_{wv}(j) - Y_{wv}(\psi(k)) \right) \\ &= \sum_{k \in H(v; w)} Y_{vw}(k) \sum_{j \in H(v; w)} Y_{wv}(j) \\ & \quad - \sum_{k \in H(v; w)} Y_{vw}(k) Y_{wv}(\psi(k)). \end{aligned}$$

Also, the last line of (8), excluding the $2/(1 + h_{vw})$ term, satisfies

$$\begin{aligned} & \sum_{\substack{l, m \in J(w; v) \\ l \neq m}} Y_{ww}(l) Y_{ww}(m) \\ &= \left(\sum_{k \in J(w; v)} Y_{ww}(k) \right)^2 - \left(\sum_{k \in J(w; v)} Y_{ww}(k)^2 \right). \end{aligned}$$

Hence, to construct our estimator $\tilde{\alpha}$, we need to calculate the following quantities:

- the sum of the $Y(k)^2$ over the v -cycles $k \in J(v; w)$;
- the sums of the $Y_{vw}(k)$, $Y_{wv}(k)$, $Y_{vw}(k)^2$, and $Y_{wv}(k)^2$ over the v -cycles $k \in H(v; w)$;
- the sum of the $Y_{vw}(k)Y_{wv}(\psi(k))$ over the v -cycles $k \in H(v; w)$;
- the sums of the $Y_{ww}(k)$ and $Y_{ww}(k)^2$ over the w -cycles $k \in J(w; v)$.

To compute these quantities in a simulation, we do not have to store the entire sample path, but rather we only need to keep track of the various cumulative sums as the simulation progresses. Therefore, compared to the standard estimator, the new estimator can be constructed with little additional computational effort and storage.

At first glance, it may appear that to calculate the expression on the right-hand side of (10) for the new estimator of the time-average variance parameter of a regenerative simulation, we must store the entire sample path. This is because as it currently stands, two passes must be made through the sample path to construct our estimator $\tilde{\sigma}^2(\vec{X})$, where $\hat{r}$ is calculated in the first pass, and then on the second pass, we can compute the cycle quantities $Y(f - \hat{r}; k)$. However, one can derive an equivalent expression for $\tilde{\sigma}^2(\vec{X})$ (which we do not give here) that can be calculated in a single pass through the sample path using only accumulators.

5 NUMERICAL EXPERIMENTS

Our numerical example is based on the Ehrenfest urn model with cost function f the identity function. The transition probabilities are given by $P_{01} = P_{s, s-1} = 1$, and

$$P_{i, i+1} = \frac{s-i}{s} = 1 - P_{i, i-1}, \quad 0 < i < s.$$

In our numerical experiments we take $s = 8$. We performed 3 experiments, corresponding to return states

$v = 1, 2$ and 4 . In all three experiments we performed 1,000 independent replications, each comprising 1,000 v -cycles. In Tables 1–3 we present results from estimating the time-average variance of this process, giving the sample average and sample variance of our new estimator over the 1,000 replications.

Table 1 shows, for all choices of alternate state w , the estimated variance and the sample variance of the estimates for return state $v = 1$. (The entry for $w = v$ corresponds to the standard regenerative estimator.) The transition probabilities are symmetric around state 4 (the mean of the binomial stationary distribution), so our first choice of return state is fairly far from the mean. Notice that the variability of the variance estimator is smaller with w near the mean state 4, and that the variance reduction is greater for $w > v$. The reason for this is that the excursions from v that go below 1 have little variability; because of the strong restoring force of the Ehrenfest model, such excursions tend to be very brief. On the other hand, excursions that get as far as the mean are likely to be quite long (and thus the contribution to the variance estimator tends to have large variability). In the second table we ran the same experiment with $v = 2$ and obtained similar results.

In the third table, we examine the same model, but now with our return state v chosen to be the stationary mean, 4. (The average cycle lengths change with different choices of v , so the results in the tables correspond to simulations of differing lengths.) The first thing to notice is that, compared with the other choices of the return state, the variance reduction is relatively small. State 4 is the best return state in the sense of minimizing the variance of the regenerative-variance estimator. Therefore, for this example, it appears that our estimator is a significant improvement over the standard regenerative estimator if the standard regenerative estimator is based on a relatively bad return state. However, if one is able to choose a near optimal return state to begin with, our estimator yields a modest improvement. Our method thus offers some safety against the choice of a bad return state.

ACKNOWLEDGMENTS

The work of the first author was supported in part by the National Science Foundation under grant DMI-9500173. The work of the second author was supported in part by the National Science Foundation under grant DMI-9624469 and by NJIT under SBR grant 421180.

Table 1: Ehrenfest Urn Model, $v = 1$

w	Avg. of $\tilde{\sigma}^2$	Sample Var.
0	13.97	1.35
1	13.97	1.35
2	13.98	0.53
3	13.96	0.26
4	13.97	0.19
5	13.98	0.23
6	13.96	0.39
7	13.97	0.69
8	13.97	1.14

Table 2: Ehrenfest Urn Model, $v = 2$

w	Avg. of $\tilde{\sigma}^2$	Sample Var.
0	13.98	1.72
1	13.98	1.72
2	13.98	1.72
3	13.98	0.86
4	13.99	0.63
5	13.99	0.65
6	14.00	0.92
7	13.99	1.35
8	13.99	1.66

Table 3: Ehrenfest Urn Model, $v = 4$

w	Avg. of $\tilde{\sigma}^2$	Sample Var.
0	13.90	1.75
1	13.90	1.68
2	13.92	1.62
3	13.91	1.56
4	13.90	1.75
5	13.88	1.53
6	13.88	1.56
7	13.90	1.63
8	13.89	1.74

REFERENCES

- Calvin, J. M. and M. K. Nakayama. 1997. A new variance-reduction technique for regenerative simulations. Research Report CIS-97-16. Department of Computer and Information Science, New Jersey Institute of Technology, Newark, NJ.
- Crane, M. and D. L. Iglehart. 1975. Simulating stable stochastic systems, III: Regenerative processes and discrete-event simulations. *Operations Research* 23:33-45.
- Glynn, P. W. and D. L. Iglehart. 1987. A joint central limit theorem for the sample mean and regenerative variance estimator. *Annals of Operations Research* 8:41-55.
- Shedler, G. S. 1993. *Regenerative Stochastic Simulation*. San Diego: Academic Press.

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 simulation output analysis and global optimization.

MARVIN K. NAKAYAMA 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. 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. His research interests include simulation output analysis, gradient estimation, rare-event simulation, random-variate generation, applied probability, and statistics.