Proceedings of the 1998 Winter Simulation Conference
D.J. Medeiros, E.F. Watson, J.S. Carson and M.S. Manivannan, eds.

EXPLOITING MULTIPLE REGENERATION SEQUENCES IN SIMULATION OUTPUT ANALYSIS

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

The regenerative method of simulation output analysis exploits the regenerative structure of a stochastic process to break up a path into independent and identically distributed cycles based on a sequence of regeneration times. If a process is regenerative with respect to more than one sequence of regeneration times, the classical regenerative method does not exploit the additional structure. In a previous paper we introduced an efficiency-improvement technique for regenerative simulation of processes having two sequences of regeneration times based on permuting regeneration cycles associated with the second sequence of regeneration points. In this paper we show how the same basic idea can be extended to exploit more than two regeneration sequences. In particular, for birth-death Markov chains, the regenerations associated with hitting times to each state can all be exploited. We present empirical results that show significant variance reductions in some cases.

1 INTRODUCTION

In Calvin and Nakayama (1998a) we introduced a new class of estimators for regenerative simulations of processes with at least two sequences of regeneration times. The estimators were proven to have the same bias and no greater (and often significantly smaller) variance than the standard estimator. The basic idea of the method is to run a simulation of a fixed number of regenerative cycles corresponding to one sequence of regeneration points, permute the cycles corresponding to a second sequence of regeneration points to obtain a new sample path, and compute an estimate based on the new path. The new estimator is the average of the estimates computed for all possible permutations.

In this paper we show how to extend the basic idea of the previous paper to exploit more than two distinct sequences of regeneration times. The idea now is that permuting cycles for any of the regeneration sequences results in a new sample path that has the same distribution as the original path. The new estimator is obtained by averaging over all such transformed paths. We derive closed-form expressions for the average, so the computational overhead is small; typically the variance reduction far outweighs the additional computational costs.

The rest of the paper is organized as follows. In Section 2 we review the basic idea of permuted regenerative estimators using two regeneration sequences and describe how the technique extends to more than two sequences. In Section 3, we derive the permuted estimator based on only two regeneration sequences for the expectation of a product of two quantities defined over a regenerative cycle. We use this representation to derive the permuted estimators for multiple regeneration sequences in Section 4. Section 5 presents the results of numerical experiments.

2 BASIC IDEA

We first review the basic idea of the approach using two distinct regeneration sequences and then describe how the method extends to more than two sequences; see Calvin and Nakayama (1998a) for details on the technique in the setting of two regenerative sequences. Let $X = \{X(t) : t \geq 0\}$ be a continuous-time stochastic process having sample paths that are right continuous with left limits on a state space $S \subset \Re^d$. We can handle discrete-time processes $\{X_n : n = 0, 1, 2, \ldots\}$ in this framework by letting $X(t) = X_{\lfloor t \rfloor}$ for all $t \geq 0$, where $\lfloor a \rfloor$ is the greatest integer less than or equal to a.

Let $T = \{T(i) : i = 0, 1, 2, \ldots\}$ be an increasing sequence of nonnegative finite stopping times. Consider

the random pair (X,T) and the shift

$$\theta_{T(i)}(X,T) = ((X(T(i)+t)_{t\geq 0},\, (T(k)-T(i))_{k\geq i})\,.$$

We define the pair (X,T) to be a (possibly delayed) *regenerative process* (in the classic sense) if

(i) $\{\theta_{T(i)}(X,T) : i = 0,1,2,\ldots\}$ are identically distributed; and

(ii) for each $i \geq 0$, $\theta_{T(i)}(X,T)$ does not depend on the "prehistory"

$$\big((X(t))_{t<T(i)},\, T(0), T(1), \ldots, T(i)\big)\,.$$

See p. 19 of Kalashnikov (1994) for further details. Simulation methods for regenerative processes have been extensively investigated; e.g., see Shedler (1993) and references therein.

We assume that there are $s \geq 2$ distinct increasing sequences of nonnegative finite stopping times, $T_1 = \{T_1(i) : i = 0,1,2,\ldots\}$ with $T_1(0) = 0$, $T_2 = \{T_2(i) : i = 0,1,2,\ldots\}$, $\ldots$, $T_s = \{T_s(i) : i = 0,1,2,\ldots\}$, such that (X,T_1) and (X,T_j), $2 \leq j \leq s$, are all regenerative processes. For example, if X is an irreducible, positive-recurrent, discrete-time or continuous-time Markov chain on the state space $\{1,2,\ldots,s\}$, then we can define T_j to be the sequence of hitting times to the state j for $1 \leq j \leq s$, where we assume that $X(0) = 1$. We use the terminology that for $1 \leq j \leq s$, the kth *j-cycle* is the segment of the sample path from time $T_j(k-1)$ to $T_j(k)$. Also, we call any T_j stopping time a *j-regeneration*.

Throughout the paper, we make the following assumption regarding the regeneration sequences: any i-regeneration and any $(i+2)$-regeneration are separated by an $(i+1)$-regeneration, and any $(i+2)$-regeneration and any i-regeneration are separated by an $(i+1)$-regeneration. In other words, no "cycles" are allowed in regeneration sequences. For example, the assumption disallows a process in which a 1-regeneration is followed by a 2-regeneration, which is followed by a 3-regeneration, which is followed by a 1-regeneration. Our assumption is satisfied, for example, by birth-death processes where the jth regeneration sequence is defined by the hitting times to state j.

We start by reviewing the technique that uses two regeneration sequences. The simplest way of explaining the technique is as follows. Suppose that we want to estimate some performance measure α, and we have available a "standard" estimator of α, which we will call $\widehat{\alpha}_m(\vec{X})$, based on a sample path $\vec{X}$ of a fixed number m of 1-cycles. Choose a regeneration sequence j, $2 \leq j \leq s$. Construct a new sample path $\vec{X}'$ from $\vec{X}$ by permuting the j-cycles in the path. More specifically, let M_j be the number of occurrences of stopping times from the sequence T_j that occur during the m 1-cycles in the path $\vec{X}$. Note that if $M_j = 0$ or 1, then the path $\vec{X}$ has no j-cycles. If $M_j = 2$, then there is only one j-cycle. Assume now that $M_j \geq 3$. Then for the given path $\vec{X}$, we can now look at the $(M_j - 1)$ j-cycles in the path. We generate a uniform random permutation of the $(M_j - 1)$ j-cycles within the path $\vec{X}$, and this gives us our new sample path $\vec{X}'$, which also has m 1-cycles. Calvin and Nakayama (1998a) show that $\vec{X}'$ and $\vec{X}$ have the same distribution.

Now for the new sample path $\vec{X}'$, we can calculate $\widehat{\alpha}_m(\vec{X}')$, which is just the standard estimator applied to the new sample path $\vec{X}'$. Denote the number of possible paths $\vec{X}'$ we can construct from $\vec{X}$ by permuting cycles as $N(\vec{X})$, which depends on $\vec{X}$ and is therefore random. We label these paths $\vec{X}^{(1)} = \vec{X}, \vec{X}^{(2)}, \ldots, \vec{X}^{(N(\vec{X}))}$, each of which has the same distribution as $\vec{X}$, and for each one we compute $\widehat{\alpha}_m(\vec{X}^{(i)})$. We finally define our new estimator for α to be

$$\widetilde{\alpha}_m(\vec{X}) = \frac{1}{N(\vec{X})} \sum_{i=1}^{N(\vec{X})} \widehat{\alpha}_m(\vec{X}^{(i)}). \tag{1}$$

In Section 3 we present closed-form expressions for a specific permuted estimator $\widetilde{\alpha}_m(\vec{X})$ when only using two regeneration sequences.

Our purpose in this paper is to extend this idea to exploit all regeneration sequences. To do this, we create a new sample path $\vec{X}'$ from the original path $\vec{X}$ by permuting all regeneration sequences simultaneously. We accomplish this by first permuting the 1-cycles, then the 2-cycles, and so on. We do this for all s regeneration sequences to get the new path $\vec{X}'$, which has the same distribution as the original path. We let $N(\vec{X})$ denote the number of possible paths that can be constructed in this manner, and denote the paths by $\vec{X}^{(i)}$, $i = 1,2,\ldots,N(\vec{X})$, as before. We then compute the new estimator as (1). In Section 4 we present closed-form expressions for a specific permuted estimator $\widetilde{\alpha}_m(\vec{X})$ when using all regeneration sequences simultaneously.

3 PERMUTED PRODUCT ESTIMATORS

We consider estimating

$$\alpha = E[U(1)V(1)], \tag{2}$$

where

$$U(k) = \int_{T_1(k-1)}^{T_1(k)} f_U(X(t))dt, \tag{3}$$

$$V(k) = \int_{T_1(k-1)}^{T_1(k)} f_V(X(t))dt, \tag{4}$$

and $f_U, f_V : S \to \Re$ are some "reward" functions. The standard estimator of α is

$$\widehat{\alpha}_m(\vec{X}) = \frac{1}{m}\sum_{k=1}^{m} U(k)V(k), \tag{5}$$

and we want to derive a formula for the permuted estimator of α. In Calvin and Nakayama (1998b) we describe a class of performance measures that can be estimated within this framework, including likelihood-ratio derivative estimators, Tin estimators, the second moment of the cumulative reward over a cycle, and the time-average variance constant.

For $Z = U, V$, let $Z_{ii}^{\uparrow}(k)$ (resp., $Z_{ii}^{\downarrow}(k)$) be the kth instance of the Z functional over an excursion from an i-regeneration back to an i-regeneration with no intervening j-regenerations for $j < i$ (resp., $j > i$), but with regenerations from the other sequences possible in between. For example, if the first six regenerations are from sequences 1, 2, 3, 2, 1 and 2, then

$$U_{11}^{\uparrow}(1) = \int_{T_1(0)}^{T_1(1)} f_U(X(t))\,dt,$$

and

$$U_{22}^{\downarrow}(1) = \int_{T_2(1)}^{T_2(2)} f_U(X(t))\,dt.$$

For $i, j \in S$, let h_{ij} be the number of times that a stopping time from the i sequence is followed next by a stopping time from the j sequence, and number the instances of the Z functional between these times

$$Z_{ij}(1), Z_{ij}(2), \ldots, Z_{ij}(h_{ij}).$$

We also number the $Z_{ii}^{\uparrow}$'s

$$Z_{ii}^{\uparrow}(1), Z_{ii}^{\uparrow}(2), \ldots, Z_{ii}^{\uparrow}(h_{ii} + h_{i,i+1}),$$

and similarly number the $Z_{ii}^{\downarrow}$'s

$$Z_{ii}^{\downarrow}(1), Z_{ii}^{\downarrow}(2), \ldots, Z_{ii}^{\downarrow}(h_{ii} + h_{i,i-1}).$$

We first specialize to the case $s = 2$. The proof of the following theorem follows from a similar result that appears in Calvin and Nakayama (1998b).

Theorem 1 *Suppose we want to estimate α defined in (2), and assume that $E[U(1)^2V(1)^2] < \infty$. The permuted estimator is given by*

$$\begin{aligned}
\widetilde{\alpha}_m(1,\vec{X}) \;=\; & \\
& \frac{1}{h_{1,1}+h_{1,2}} \Bigg[\sum_{k=1}^{h_{1,1}} U_{11}(k)V_{11}(k) \\
& + \sum_{k=1}^{h_{1,2}} (U_{12}(k)V_{12}(k) + U_{21}(k)V_{21}(k)) \\
& + \frac{1}{h_{1,2}} \sum_{k=1}^{h_{1,2}} U_{12}(k) \sum_{j=1}^{h_{1,2}} V_{21}(j) \\
& + \frac{1}{h_{1,2}} \sum_{k=1}^{h_{1,2}} V_{1,2}(k) \sum_{j=1}^{h_{1,2}} U_{21}(j) \\
& + \frac{1}{h_{1,2}} \sum_{k=1}^{h_{1,2}} (U_{12}(k) + U_{21}(k)) \sum_{j=1}^{h_{2,2}} V_{22}^{\uparrow}(j) \\
& + \frac{1}{h_{1,2}} \sum_{k=1}^{h_{1,2}} (V_{12}(k) + V_{21}(k)) \sum_{j=1}^{h_{2,2}} U_{22}^{\uparrow}(j) \\
& + \sum_{l=1}^{h_{2,2}} U_{22}^{\uparrow}(l)V_{22}^{\uparrow}(l) \\
& + \frac{2}{1+h_{12}} \sum_{\substack{i,j \le h_{2,2}, \\ i \ne j}} U_{22}^{\uparrow}(i)V_{22}^{\uparrow}(j) \Bigg]
\end{aligned}$$

if $M_2 \ge 3$, and $\widetilde{\alpha}_m(\vec{X}) = \widehat{\alpha}_m(\vec{X})$ otherwise, where $\widehat{\alpha}_m(\vec{X})$ is the standard estimator of α as defined in (5). The estimator satisfies $E[\widetilde{\alpha}_m(\vec{X})] = \alpha$ and

$$\mathit{Var}(\widetilde{\alpha}_m(\vec{X})) \le \mathit{Var}(\widehat{\alpha}_m(\vec{X})).$$

The estimator $\widetilde{\alpha}_m(1,\vec{X})$ of Theorem 1 is the expectation of the standard estimator $\widehat{\alpha}_m(\vec{X}^{(i)})$ with respect to a random permutation of the 2-cycles and 1-cycles. The similar result from Calvin and Nakayama (1998b) alluded to above is the expectation with respect to 2-cycles only, and is slightly more complicated and less efficient.

4 MULTIPLE REGENERATIVE SEQUENCES

In the previous section we presented a closed-form expression for the estimator obtained by permuting two regeneration sequences. We now present the formula for the estimator using $s > 2$ distinct regeneration sequences; for the proof and details see Calvin and Nakayama (1998c). Our new estimator will be the expectation of the standard estimator with respect to a random transformation of the original path.

To present the new estimator, we need to introduce some notation. First, fix a "return state" v, $1 \le v \le s$. We consider estimating α defined by (2)–(4), but with the T_1 sequence replaced by T_v.

For $1 \le i \le s$, define

$$\xi_i = \sum_{k=1}^{h_{i,i}} U_{ii}(k)V_{ii}(k) \tag{6}$$

and

$$\begin{aligned}\beta_{i,i+1} \;=\; & \sum_{k=1}^{h_{i,i+1}} [U_{i,i+1}(k)V_{i,i+1}(k) \\ & + U_{i+1,i}(k)V_{i+1,i}(k)] \\ & + \frac{1}{h_{i,i+1}} \left(\sum_{k=1}^{h_{i,i+1}} U_{i,i+1}(k) \sum_{j=1}^{h_{i,i+1}} V_{i+1,i}(j) \right. \\ & \left. + \sum_{k=1}^{h_{i,i+1}} V_{i,i+1}(k) \sum_{j=1}^{h_{i,i+1}} U_{i+1,i}(j) \right).\end{aligned}$$

Also define

$$\eta_{i,i+1}^U = \sum_{k=1}^{h_{i,i+1}} \left(U_{i,i+1}(k) + U_{i+1,i}(k) \right),$$

$$\eta_{i,i+1}^V = \sum_{k=1}^{h_{i,i+1}} \left(V_{i,i+1}(k) + V_{i+1,i}(k) \right),$$

(by assumption, $h_{i,i+1} = h_{i+1,i}$) and

$$\gamma_i^U = \sum_{k=1}^{h_{i,i}} U_{i,i}(k), \quad \gamma_i^V = \sum_{k=1}^{h_{i,i}} V_{i,i}(k).$$

For $1 \le i \le s$, define

$$S_i^{\uparrow U} = \sum_{k=1}^{h_{i,i}+h_{i,i+1}} U_{ii}^{\uparrow}(k) = \sum_{k=i}^{s} \eta_{k,k+1}^U + \gamma_k^U,$$

$$S_i^{\uparrow V} = \sum_{k=1}^{h_{i,i}+h_{i,i+1}} V_{ii}^{\uparrow}(k) = \sum_{k=i}^{s} \eta_{k,k+1}^V + \gamma_k^V,$$

and

$$S_i^{\downarrow U} = \sum_{k=1}^{h_{i,i}+h_{i,i+1}} U_{ii}^{\downarrow}(k) = \sum_{k=1}^{i} \eta_{k,k+1}^U + \gamma_k^U,$$

$$S_i^{\downarrow V} = \sum_{k=1}^{h_{i,i}+h_{i,i+1}} V_{ii}^{\downarrow}(k) = \sum_{k=i}^{s} \eta_{k,k+1}^V + \gamma_k^V.$$

Set

$$\theta^s = \frac{1}{h_{s,s}} \sum_{k=1}^{h_{s,s}} U_{ss}(k)V_{ss}(k) = \frac{\xi_s}{h_{s,s}},$$

and

$$\theta_1 = \frac{1}{h_{1,1}} \sum_{k=1}^{h_{1,1}} U_{11}(k)V_{11}(k) = \frac{\xi_1}{h_{1,1}}.$$

For $i = s-1, s-2, \ldots, v+1$, define θ^i recursively by

$$\begin{aligned}\theta^i \;=\; & \frac{1}{h_{i,i} + h_{i,i+1} + h_{i,i-1}} (\xi_i + \beta_{i,i+1} \\ & + \left(\frac{1}{h_{i,i+1}} \right) \left(\eta_{i,i+1}^U S_{i+1}^{\uparrow V} + \eta_{i,i+1}^V S_{i+1}^{\uparrow U} \right) \\ & + \left(\frac{2}{1 + h_{i,i+1}} \right) S_{i+1}^{\uparrow U} S_{i+1}^{\uparrow V} \\ & + \left(\frac{(h_{i+1,i+1} + h_{i+1,i+2})(h_{i,i+1} - 1)}{1 + h_{i,i+1}} \right) \theta^{i+1}\end{aligned}$$

and for $i = 2, 3, \ldots, v-1$, define θ_i recursively by

$$\begin{aligned}\theta_i \;=\; & \frac{1}{h_{i,i} + h_{i,i+1} + h_{i,i-1}} (\xi_i + \beta_{i-1,i} \\ & + \left(\eta_{i-1,i}^U S_{i-1}^{\downarrow V} + \eta_{i-1,i}^V S_{i-1}^{\downarrow U} \right) \\ & + \left(\frac{2}{1 + h_{i,i-1}} \right) S_{i-1}^{\downarrow U} S_{i-1}^{\downarrow V} \\ & + \left(\frac{(h_{i-1,i-1} + h_{i-1,i-2})(h_{i,i-1} - 1)}{1 + h_{i,i-1}} \right) \theta_{i-1} \Big).\end{aligned}$$

Finally, the permuted estimator is given by

$$\begin{aligned}\widetilde{\alpha}_m(v) \;=\; & \frac{1}{h_{v,v} + h_{v,v+1} + h_{v,v-1}} (\xi_v + \beta_{v,v+1} \\ & + \left(\frac{1}{h_{v,v+1}} \right) \left(\eta_{v,v+1}^U S_{v+1}^{\uparrow V} + \eta_{v,v+1}^V S_{v+1}^{\uparrow U} \right) \\ & + \left(\frac{2}{1 + h_{v,v+1}} \right) S_{v+1}^{\uparrow U} S_{v+1}^{\uparrow V} \\ & + \left(\frac{(h_{v+1,v+1} + h_{v+1,v+2})(h_{v,v+1} - 1)}{1 + h_{v,v+1}} \right) \theta^{v+1} \\ & + \beta_{v-1,v} + \left(\eta_{v-1,v}^U S_{v-1}^{\downarrow V} + \eta_{v-1,v}^V S_{v-1}^{\downarrow U} \right) \\ & + \left(\frac{2}{1 + h_{v,v-1}} \right) S_{v-1}^{\downarrow U} S_{v-1}^{\downarrow V} \\ & + \left(\frac{(h_{v-1,v-1} + h_{v-1,v-2})(h_{v,v-1} - 1)}{1 + h_{v,v-1}} \right) \theta_{v-1} \Big).\end{aligned}$$

The following three steps summarize the calculation of the estimator $\widetilde{\alpha}_m(v)$.

Step 1. Generate a sample path of m v-cycles. After each transition is generated, increment the appropriate counters and accumulators for h_{ij}, $\sum U_{ij}$, $\sum V_{ij}$, and $\sum U_{ij}V_{ij}$ for $1 \le i, j \le s$.

Step 2. From the data collected in Step 1, construct the ξ_i's, β_{ij}'s, η_{ij}^U's, η_{ij}^V's, γ_{ij}^U's, γ_{ij}^V's, $S_i^{\uparrow U}$'s, $S_i^{\uparrow V}$'s, $S_i^{\downarrow U}$'s, and $S_i^{\downarrow V}$'s, for $1 \le i, j \le s$.

Step 3. Use the quantities calculated in Step 2 in the recursive formula given above to calculate the θ^i, $v < i \leq s$, and the θ_i, $1 \leq i < v$, and from them $\widetilde{\alpha}_m(v)$.

The total storage required is proportional to s. The work of the calculation in Step 1 is proportional to the sample path length. The calculation in Step 2 requires $O(s)$ work, and the calculations in Step 3 also require $O(s)$ work, both independent of the simulation run length.

5 NUMERICAL EXPERIMENTS

We now present some numerical results from simulating a discrete-time Markov chain using our permuted estimators. We let each stopping-time sequence correspond to successive hitting times to a fixed state, and we compute the estimators based on only two regeneration sequences and the estimator using all of the sequences. We run different experiments with different choices of return state v. In the case when only two regeneration sequences are used, the T_1 (resp., T_2) sequence corresponds to hits to state v (resp., some state $w \neq v$). We examine the variance reduction resulting from different choices of w and compare these to the reduction from using all sequences simultaneously.

The model is a discrete-time Markov chain $X = \{X_n : n = 0, 1, 2, \ldots\}$ on state space $S = \{0, 1, \ldots, s\}$ with transition probability matrix

$$R(i, i+1) = \frac{\lambda}{i+\lambda} = 1 - R(i, i-1)$$

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. As $n \to \infty$, $X_n \Rightarrow X_\infty$, where $P\{X_\infty = i\} = \pi_i$, $0 \leq i \leq s$, and $\pi \equiv (\pi_0, \pi_1, \ldots, \pi_s)$ is the stationary distribution of X. (Note that π depends on the parameter λ, and so we sometimes write $\pi \equiv \pi(\lambda) \equiv (\pi_0(\lambda), \pi_1(\lambda), \ldots, \pi_s(\lambda))$ to emphasize the dependence.)

Let $\alpha(\lambda) \equiv \lambda E_{\pi(\lambda)}[X_\infty]$ be the steady-state average cost, where $E_{\pi(\lambda)}$ is the expectation operator under the steady-state distribution π parameterized by λ, and our goal is to estimate $\partial\alpha(\lambda)$, the derivative of the average cost with respect to λ.

We ran numerical experiments with $s = 20$ and $\lambda = 5$. The theoretical value of $\partial\alpha(5)$ is approximately 10.5. Tables 1-3 report the results of simulations of 1,000 independent replications for the return states $v = 0$, 5 and 10, respectively. The column labeled with "Avg. of $\widetilde{\alpha}_m$" (resp., "Samp. Var.") is the average (resp., sample variance) of the permuted estimators over the 1,000 replications. In each table the row for $w = v$ corresponds to the standard (unpermuted) regenerative estimator, and the row labeled "all" gives the results for the combined estimator (using all regeneration sequences simultaneously).

Table 1: Derivative estimates, $v = 0$.

w	Avg. of $\widetilde{\alpha}_m$	Samp. Var.
0	10.35	2.43
5	10.30	0.36
10	10.31	0.66
15	10.33	2.33
20	10.35	2.43
all	10.26	0.30

Table 2: Derivative estimates, $v = 5$.

w	Avg. of $\widetilde{\alpha}_m$	Samp. Var.
0	10.47	0.23
5	10.47	0.23
10	10.48	0.20
15	10.47	0.22
20	10.47	0.23
all	10.49	0.16

The expected cycle lengths for different choices of the return state vary considerably, with $v = 5$ having the shortest expected cycle length (about 5.7) over the entire state space. When $v = 0$, the expected cycle length is quite long (about 300), and the expected cycle length for $v = 10$ is about 37. Therefore, for each return state the number of cycles simulated is adjusted so that the total expected number of transitions in each replication is about 20,000. For $v = 0$ as the return state, 67 cycles were simulated; for $v = 5$ as the return state, 3,501 cycles were simulated; and for $v = 10$ as the return state, 544 cycles were required. Hence, we can directly compare the results across the three tables. (The difference in the number of cycles may account for the difference in the apparent bias among the tables corresponding to the different choices of return state v.)

Table 3: Derivative estimates, $v = 10$.

w	Avg. of $\widetilde{\alpha}_m$	Samp. Var.
0	10.44	0.51
5	10.45	0.19
10	10.43	0.58
15	10.43	0.58
20	10.43	0.58
all	10.45	0.17

With the three choices of return state $v = 0, 5, 10$, the sample variance of the standard regenerative estimators are 2.43, 0.23, and 0.58, respectively. For the combined estimator, the corresponding sample variances are 0.30, 0.16, and 0.17, respectively. In this example we see that the variability of the standard regenerative estimator is quite sensitive to the "return state", while the variability of the combined permuted estimator is relatively insensitive to the return state. Even when the "best" return state 5 is used, the combined permuted estimator gives a significant reduction in sample variance. In the case of $v = 0$, the reduction is by a factor of about 8.

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.

REFERENCES

Calvin, J. M. and M. K. Nakayama. 1998a. Using permutations in regenerative simulations to reduce variance. *ACM Transactions on Modeling and Computer Simulation*, to appear.

Calvin, J. M. and M. K. Nakayama. 1998b. Permuted estimators for regenerative simulations. Working paper.

Calvin, J. M. and M. K. Nakayama. 1998c. Exploiting multiple regeneration sequences in simulation output analysis. Working paper.

Glynn, P. W. 1990. Likelihood ratio derivative estimators for stochastic systems. *Comm. ACM* 33: 75–84.

Glynn, P. W. and W. Whitt 1992. The asymptotic efficiency of simulation estimators. *Operations Research* 40: 505–520.

Kalashnikov, V. 1994. *Topics on Regenerative Processes*. Boca Raton: CRC Press.

Reiman, M. I. and A. Weiss. 1989. Sensitivity analysis for simulations via likelihood ratios. *Operations Research* 37: 830–844.

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 the Ph.D. in Operations Research from Stanford University. His research interests include global optimization and simulation output analysis.

MARVIN K. NAKAYAMA is an assistant professor in the Department of Computer and Information Science at the New Jersey Institute of Technology. He received the 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 interests include applied probability, simulation output analysis, variance-reduction techniques, gradient estimation, and rare-event simulation.