Simulation of Processes with Multiple Regeneration Sequences

James M. Calvin and Marvin K. Nakayama

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

December 20, 1999

Abstract

The classical 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 single 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, and the variance of the resulting estimator for certain performance measures (e.g., the time-average variance constant) can vary greatly depending on the particular regeneration sequence chosen. In a previous paper we introduced an efficiency-improvement technique for regenerative simulation of processes having two sequences of regeneration times based on permuting regenerative cycles associated with the second sequence of regeneration points. In this paper we show how 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, and the results seem to indicate that the permuted estimator for the time-average variance constant can have variance that is independent of the primary regeneration sequence used to run the simulation.

Key words: Output analysis, regenerative simulation, variance reduction, efficiency improvement.

1 Introduction

The goal of a discrete-event simulation is typically to estimate some steady-state or transient performance measure of a stochastic system. For example, one may be interested in determining the expected steady-state waiting times in a queueing system or the time-average variance of the number of customers in the system.

The classical regenerative method (Crane and Iglehart 1975) 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 single sequence of regeneration times. One can then construct asymptotically valid confidence intervals.

Many stochastic processes arising in practice are regenerative with respect to more than one sequence of regeneration times. For example, if the process is an irreducible Markov chain on a finite state space, then the successive hitting times to any fixed state form a regeneration sequence. However, the classical regenerative method does not exploit the additional structure in such a model, and the particular choice of the regeneration sequence can significantly affect the variability of the regenerative estimator of certain performance measures (e.g., the time-average variance constant). Unfortunately, there are no reliable rules that can *a priori* select a "good" regeneration sequence.

In a previous paper (Calvin and Nakayama 1998a) we introduced a new class of estimators for simulations of processes with two sequences of regeneration points. The basic idea of the method is to run a simulation of a fixed number of regeneration 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. The estimators were proven to have the same bias and no greater (and at times significantly smaller) variance than the standard estimator. However, when simulating systems with more than two regeneration sequences, we empirically found that the variance of the resulting estimator can vary greatly depending on the particular pair of sequences chosen. For some pairs, the variance was reduced by a factor of up to 8, but for other pairs, there was no variance reduction. Since it may be difficult to identify a "good" pair of regeneration sequences prior to running the simulation, one may not be able to apply the method as a "black box."

In this paper we show how to exploit all regeneration sequences simultaneously when

there are more than two such sequences. We begin by simulating a fixed number of cycles with respect to one regeneration sequence, which we call sequence T_1. The idea is that simultaneously permuting cycles for all 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, which reduces variance even more than the previous method based on only two sequences. We derive closed-form expressions for the average, so the computational overhead is small; typically the variance reduction far outweighs the additional computational costs.

We present numerical results that show that permuting all possible regeneration sequences can yield large variance reductions (up to a factor of 66 in our experiments). The amount of variance reduction is at least as much as that obtained from our previous method using any pair of sequences. Moreover (and perhaps more significantly), the results seem to indicate that when estimating the time-average variance constant, the method yields an estimator whose variability is affected little by the choice of sequence T_1. Thus, the method can apparently be used as a "black box" since it does not rely on the user to specify difficult-to-choose parameters (i.e., which regeneration sequences to use).

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 then describe how to exploit more than two sequences under certain assumptions. In Section 3, we modify our permuted estimator based on only two regeneration sequences from Calvin and Nakayama (1998a) to derive a new estimator. We then use this to construct the permuted estimators for multiple regeneration sequences in Section 4. In Section 5 we extend some ideas developed in Calvin and Nakayama (1997) to prove a strong law of large numbers and a central limit theorem for the new permuted estimator. Thus, we can construct confidence intervals for the estimator. Section 6 presents the results of numerical experiments. Some of the results of this paper appear (without proofs) in Calvin and Nakayama (1998c).

2 Basic Idea

We first review the basic idea of the approach using two distinct regeneration sequences and then describe how to exploit more than two sequences; see Calvin and Nakayama (1998a) for details on the technique in the setting of two regeneration 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 the following:

A1 *There are $s \geq 2$ disjoint increasing sequences of nonnegative finite stopping times, $T_1 = \{T_1(i) : i = 0, 1, 2, \ldots\}$ with $T_1(0) = 0$, $T_2 = \{T_2(i) : i = 0, 1, 2, \ldots\}$, ..., $T_s = \{T_s(i) : i = 0, 1, 2, \ldots\}$, 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 *T_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 *T_j-regeneration*.

We will develop estimators under the following assumption regarding the regeneration sequences:

A2 *Any T_j-regeneration and any T_{j+2}-regeneration are separated by a T_{j+1}-regeneration, and any T_{j+2}-regeneration and any T_j-regeneration are separated by a T_{j+1}-regeneration.*

In other words, no "cycles" are allowed in regeneration sequences. For example, the assumption disallows a process in which a T_1-regeneration is followed by a T_2-regeneration, which is followed by a T_3-regeneration, which is followed by a T_1-regeneration. Our assumption is satisfied by the following examples:

Example 1 Let X be a birth-death Markov chain on the state space $S = \{1, 2, 3, \ldots\}$, and define the T_j-regeneration sequence to be the hitting times to state j, $j \in S$. Then Assumption 2 holds. ∎

Example 2 If the regeneration sequences of a process do not satisfy Assumption 2, they can be redefined in such a way that the assumption does hold. This can be accomplished by redefining the regeneration sequences so that we "ignore" those regenerations that would cause Assumption 2 to be violated. We illustrate this point with the example of the M/G/1 queue with first-come-first-served service. We can define the jth regeneration sequence as the successive times that a departure leaves $j-1$ jobs in the system for $j = 1, 2, 3, \ldots$. With this set of sequences, Assumption 2 does not hold. (This is because after a regeneration from the jth sequence, there can be 2 customer arrivals before the next service completion, in which case the next regeneration would be from the $(j+2)$th sequence, and so there was no intervening regeneration from the $(j+1)$th sequence.) Instead, we modify our definition of the regeneration sequences so that after a regeneration from the jth sequence, we ignore the original regenerations until the next regeneration from either the $(j-1)$th, jth, or $(j+1)$th sequence. With this change, Assumption 2 holds. For more details, see the appendix. ∎

We start by reviewing the technique that uses two regeneration sequences. The simplest way of explaining the method 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 denote by $\widehat{\alpha}_m(\vec{X})$, based on a sample path $\vec{X}$ of a fixed number m of T_1-cycles. Choose a regeneration sequence T_j, $2 \le j \le s$. Construct a new sample path $\vec{X}'$ from $\vec{X}$ by permuting the T_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 T_1-cycles in the path $\vec{X}$. Note that if $M_j = 0$ or 1, then the path $\vec{X}$ has no T_j-cycles, and if $M_j = 2$, then there is only one T_j-cycle. Assume now that $M_j \ge 3$. Then for the given path $\vec{X}$, we can look at the $(M_j - 1)$ T_j-cycles in the path. We generate a uniform random permutation of the $(M_j - 1)$ T_j-cycles within the path $\vec{X}$, and this gives us our new sample path $\vec{X}'$, which also has m T_1-cycles. Calvin and Nakayama (1998a) prove 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 by $N(\vec{X})$, 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 show how 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 T_1-cycles, then the T_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 Two Regeneration Sequences

We consider estimating

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

where

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

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

for $k \geq 1$, and $f_U, f_V : S \to \Re$ are some "reward" functions. The standard estimator of α is

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

where the "1" in $\widehat{\alpha}_m(1, \vec{X})$ denotes that the T_1-cycles are used to construct the estimator. We want to derive a formula for the permuted estimator of α. In Calvin and Nakayama

(1998b) we describe a class of performance-measure estimators that fall within this framework, including likelihood-ratio derivative estimators (e.g., see Glynn 1990), Tin estimators for low-bias estimation of ratios of means (e.g., see p. 104 of Shedler 1993), and estimators of the second moment of the cumulative reward over a cycle and the time-average variance constant (which appears in the central limit theorem for the time-average reward; e.g., see p. 100 of Shedler 1993).

For $Z = U, V$, let $Z_{i,i}^{\uparrow}(k)$ be the kth instance of the Z functional over an excursion from an T_i-regeneration back to an T_i-regeneration with no intervening T_j-regenerations for $j < i$, but with regenerations from the other sequences possible in between. For example, if the first six regenerations (from any sequences) on a particular sample path are $T_1(0)$, $T_2(0)$, $T_1(1)$, $T_2(1)$, $T_3(0)$, and $T_2(2)$, then $U_{2,2}^{\uparrow}(1) = \int_{T_2(1)}^{T_2(2)} f_U(X(t))\,dt$.

For $i, j \in S$, let $h_{i,j}$ be the number of times on the sample path $\vec{X}$ that a stopping time from the T_i-sequence is followed next by a stopping time from the T_j-sequence, and number the instances of the Z functional between these times

$$Z_{i,j}(1), Z_{i,j}(2), \ldots, Z_{i,j}(h_{i,j}),$$

and similarly number the $Z_{i,i}^{\uparrow}$'s

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

In Figure 1 we illustrate some of our notation with an example of a sample path $\vec{X}$ of a continuous-time process X on a continuous state space S. Regeneration sequence T_i, $i = 1, 2, 3, 4$, corresponds to successive hits to state i. The number of T_1-cycles is $m = 4$, and $h_{4,4} = 0$, $h_{1,1} = h_{2,2} = h_{3,3} = h_{3,4} = h_{4,3} = 1$, $h_{2,3} = h_{3,2} = 2$, and $h_{1,2} = h_{2,1} = 3$.

In general, $\widetilde{\alpha}_m(1, \vec{X})$ will denote the expectation of $\widehat{\alpha}_m(1, \vec{X})$ with respect to a random ordering of all of the regeneration cycles. We first specialize to the case $s = 2$. The following theorem generalizes a similar result that appears in Calvin and Nakayama (1998b).

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

$$\widetilde{\alpha}_m(1, \vec{X}) \quad = \quad \frac{1}{h_{1,1} + h_{1,2}} \left[\sum_{k=1}^{h_{1,1}} U_{1,1}(k) V_{1,1}(k) + \sum_{k=1}^{h_{1,2}} \left(U_{1,2}(k) V_{1,2}(k) + U_{2,1}(k) V_{2,1}(k) \right) \right.$$

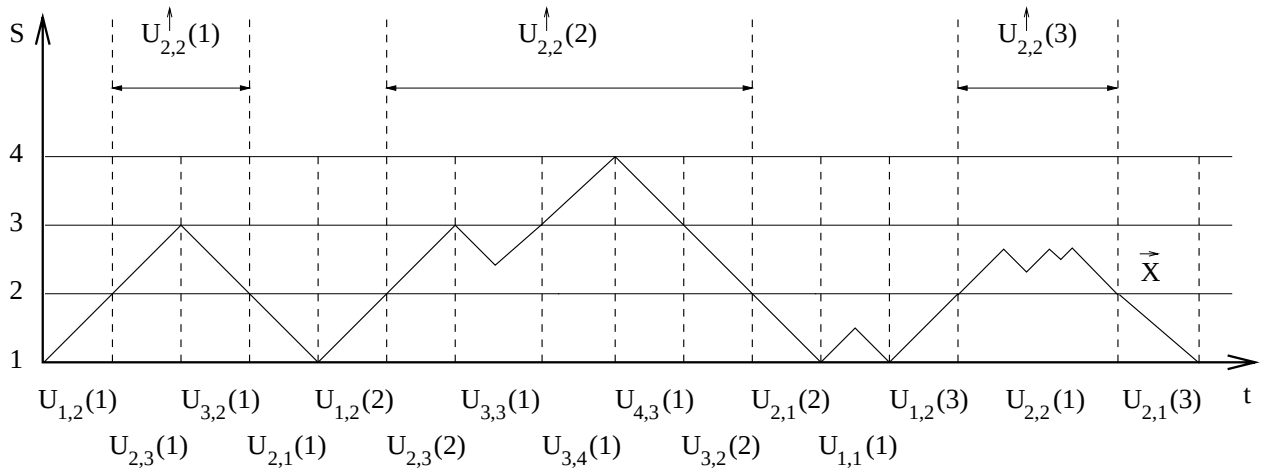

Figure 1: A sample path of a process with 4 regeneration sequences.

$$
\begin{aligned}
&+\frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\sum_{j=1}^{h_{1,2}} V_{2,1}(j) + \frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} V_{1,2}(k)\sum_{j=1}^{h_{1,2}} U_{2,1}(j) \\
&+\frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} (U_{1,2}(k)+U_{2,1}(k))\sum_{j=1}^{h_{2,2}} V^{\uparrow}_{2,2}(j) \\
&+\frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} (V_{1,2}(k)+V_{2,1}(k))\sum_{j=1}^{h_{2,2}} U^{\uparrow}_{2,2}(j) \\
&+\frac{2}{1+h_{1,2}}\sum_{\substack{i,j\le h_{2,2},\\ i\neq j}} U^{\uparrow}_{2,2}(i)V^{\uparrow}_{2,2}(j) + \sum_{l=1}^{h_{2,2}} U^{\uparrow}_{2,2}(l)V^{\uparrow}_{2,2}(l)\Bigg]
\end{aligned}
$$

if $M_2 \geq 3$*, and* $\widetilde{\alpha}_m(1,\vec{X}) = \widehat{\alpha}_m(1,\vec{X})$ *otherwise, where* $\widehat{\alpha}_m(1,\vec{X})$ *is the standard estimator of* α *as defined in (5). For all* $m \geq 1$*, the estimator satisfies* $E[\widetilde{\alpha}_m(1,\vec{X})] = \alpha$ *and*

$$Var(\widetilde{\alpha}_m(1,\vec{X})) \leq Var(\widehat{\alpha}_m(1,\vec{X})).$$

Proof. Let E_i denote expectation with respect to a uniform random permutation of the T_i-cycles, given the original path $\vec{X}$. Calvin and Nakayama (1998a) derive the following formula for the expectation with respect to permuting only the T_2-cycles:

$$
\begin{aligned}
E_2\left[\widehat{\alpha}_m(1,\vec{X})\right] = \frac{1}{h_{1,1}+h_{1,2}}\Bigg[&\sum_{k=1}^{h_{1,1}} U_{1,1}(k)V_{1,1}(k) + \sum_{k=1}^{h_{1,2}} (U_{1,2}(k)V_{1,2}(k)+U_{2,1}(k)V_{2,1}(k)) \\
&+\frac{1}{h_{1,2}-1}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\left(\sum_{j=1}^{h_{1,2}} V_{2,1}(j) - V_{2,1}(\psi(k))\right) + \frac{1}{h_{1,2}-1}\sum_{k=1}^{h_{1,2}} V_{1,2}(k)\left(\sum_{j=1}^{h_{1,2}} U_{2,1}(j) - U_{2,1}(\psi(k))\right) \\
&+\frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} (U_{1,2}(k)+U_{2,1}(k))\sum_{j=1}^{h_{2,2}} V^{\uparrow}_{2,2}(j) + \frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} (V_{1,2}(k)+V_{2,1}(k))\sum_{j=1}^{h_{2,2}} U^{\uparrow}_{2,2}(j)
\end{aligned}
$$

$$+\sum_{l=1}^{h_{2,2}} U_{2,2}^{\uparrow}(l)V_{2,2}^{\uparrow}(l) + \frac{2}{1+h_{1,2}} \sum_{\substack{i,j\le h_{2,2},\\ i\neq j}} U_{2,2}^{\uparrow}(i)V_{2,2}^{\uparrow}(j)\Bigg],$$

where $\psi(1) = h_{1,2}$ and $\psi(k) = k-1$ for $k > 1$. Now consider the effect of permuting the T_1-cycles; only the expressions in the second line of the above formula are affected. Taking the expectation of the first of those expressions with respect to a random permutation of the T_1-cycles gives

$$\begin{aligned}
&E_1\left[\frac{1}{h_{1,2}-1}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\left(\sum_{j=1}^{h_{1,2}} V_{2,1}(j) - V_{2,1}(\psi(k))\right)\right] \\
&= \frac{1}{h_{1,2}-1}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\sum_{j=1}^{h_{1,2}} V_{2,1}(j) - \frac{1}{h_{1,2}-1}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\, E_1\left[V_{2,1}(\psi(k))\right] \\
&= \frac{1}{h_{1,2}-1}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\sum_{j=1}^{h_{1,2}} V_{2,1}(j) - \frac{1}{h_{1,2}-1}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\frac{1}{h_{1,2}}\sum_{j=1}^{h_{1,2}} V_{2,1}(j) \\
&= \frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\sum_{j=1}^{h_{1,2}} V_{2,1}(j).
\end{aligned}$$

The expectation of the second expression simplifies similarly, thus proving the validity of the expression for $\widetilde{\alpha}_m(1,\vec{X})$.

Calvin and Nakayama (1998a) proved that when permuting only T_2-cycles, each permuted path $\vec{X}^{(i)}$ in (1) has the same distribution as the original path $\vec{X}$. We can apply the same argument to show that if $\vec{X}'$ is constructed from $\vec{X}$ by permuting both T_1-cycles and T_2-cycles, then $\vec{X}'$ has the same distribution as $\vec{X}$. Hence, if we let the $\vec{X}^{(i)}$, $i = 1, 2, \ldots, N(\vec{X})$, denote the paths that can be constructed by permuting both T_1-cycles and T_2-cycles, then it follows from (1) that $E[\widetilde{\alpha}_m(1,\vec{X})] = \alpha$ and $\mathrm{Var}(\widetilde{\alpha}_m(1,\vec{X})) \le \mathrm{Var}(\widehat{\alpha}_m(1,\vec{X}))$. ∎

For what follows it will be convenient to rewrite the expression for $\widetilde{\alpha}_m(1,\vec{X})$ as

$$\begin{aligned}
\widetilde{\alpha}_m(1,\vec{X}) \;=\; & \frac{1}{h_{1,1}+h_{1,2}}\Bigg[\sum_{k=1}^{h_{1,1}} U_{1,1}(k)V_{1,1}(k) + \sum_{k=1}^{h_{1,2}}\left(U_{1,2}(k)V_{1,2}(k) + U_{2,1}(k)V_{2,1}(k)\right) \\
&+\frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\sum_{j=1}^{h_{1,2}} V_{2,1}(j) + \frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} V_{1,2}(k)\sum_{j=1}^{h_{1,2}} U_{2,1}(j) \\
&+\frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}}\left(U_{1,2}(k)+U_{2,1}(k)\right)\sum_{j=1}^{h_{2,2}} V_{2,2}^{\uparrow}(j) \\
&+\frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}}\left(V_{1,2}(k)+V_{2,1}(k)\right)\sum_{j=1}^{h_{2,2}} U_{2,2}^{\uparrow}(j)
\end{aligned}$$

$$+\frac{h_{1,2}-1}{h_{1,2}+1}\sum_{l=1}^{h_{2,2}} U_{2,2}^{\uparrow}(l)V_{2,2}^{\uparrow}(l) + \frac{2}{1+h_{1,2}}\left(\sum_{l=1}^{h_{2,2}} U_{2,2}^{\uparrow}(l)\right)\left(\sum_{l=1}^{h_{2,2}} V_{2,2}^{\uparrow}(l)\right)\Bigg]. \quad (6)$$

4 Multiple Regeneration 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 $s \geq 2$ distinct regeneration sequences.

Theorem 2 *Suppose that there are $s \geq 2$ regeneration sequences. For $1 \leq i \leq s$, let*

$$\begin{aligned}
\beta(i) \;=\; & \sum_{k=1}^{h_{i,i}} U_{i,i}(k)V_{i,i}(k) + \sum_{k=1}^{h_{i,i+1}} \left(U_{i,i+1}(k)V_{i,i+1}(k) + U_{i+1,i}(k)V_{i+1,i}(k)\right) \\
& + \frac{1}{h_{i,i+1}} \sum_{k=1}^{h_{i,i+1}} U_{i,i+1}(k) \sum_{j=1}^{h_{i,i+1}} V_{i+1,i}(j) + \frac{1}{h_{i,i+1}} \sum_{k=1}^{h_{i,i+1}} V_{i,i+1}(k) \sum_{j=1}^{h_{i,i+1}} U_{i+1,i}(j) \\
& + \frac{1}{h_{i,i+1}} \sum_{k=1}^{h_{i,i+1}} \left(U_{i,i+1}(k) + U_{i+1,i}(k)\right) \sum_{j=1}^{h_{i+1,i+1}+h_{i+1,i+2}} V_{i+1,i+1}^{\uparrow}(j) \\
& + \frac{1}{h_{i,i+1}} \sum_{k=1}^{h_{i,i+1}} \left(V_{i,i+1}(k) + V_{i+1,i}(k)\right) \sum_{j=1}^{h_{i+1,i+1}+h_{i+1,i+2}} U_{i+1,i+1}^{\uparrow}(j) \\
& + \frac{2}{1+h_{i,i+1}} \left(\sum_{l=1}^{h_{i+1,i+1}+h_{i+1,i+2}} U_{i+1,i+1}^{\uparrow}(l)\right)\left(\sum_{l=1}^{h_{i+1,i+1}+h_{i+1,i+2}} V_{i+1,i+1}^{\uparrow}(l)\right),
\end{aligned}$$

where $h_{i,j} = U_{i,j}(k) = V_{i,j}(k) = U_{i,i}^{\uparrow}(k) = V_{i,i}^{\uparrow}(k) = 0$ if i or j is greater than s. Then

$$\widetilde{\alpha}_m(1,\vec{X}) = \frac{1}{h_{1,1}+h_{1,2}} \sum_{i=1}^{s} \beta(i) \prod_{j=1}^{i-1} \left(\frac{h_{j,j+1}-1}{h_{j,j+1}+1}\right), \quad (7)$$

where we take $\prod_{j=1}^{0} = 1$. For all $m \geq 1$, the estimator satisfies $E[\widetilde{\alpha}_m(1,\vec{X})] = \alpha$ and

$$Var(\widetilde{\alpha}_m(1,\vec{X})) \leq Var(\widehat{\alpha}_m(1,\vec{X})).$$

Proof. When $s = 2$, equation (6) is the expectation of the standard estimator with respect to a random ordering of the T_1-cycles and T_2-cycles, and we can think of $h_{2,2}$ in (6) as the number of T_2-regenerations that are *not* followed by a T_1-regeneration. If $s > 2$, then $h_{2,2} + h_{2,3}$ is the number of T_2-regenerations that are *not* followed by a T_1-regeneration, and so in (6) we replace $h_{2,2}$ with $h_{2,2} + h_{2,3}$, which is the total number of the $U_{2,2}^{\uparrow}(j)$'s. Hence, with this modification to (6), we can now interpret (6) as the conditional expectation from

randomly permuting the T_1-cycles and T_2-cycles, given the original orderings of the T_3-cycles through the T_s-cycles. In other words, letting $E_{i_1,i_2,\ldots,i_p}$ denote expectation with respect to uniform random permutations of the T_{i_1}-cycles through the T_{i_p}-cycles, we get that (6) becomes

$$
\begin{aligned}
&E_{1,2}\left[\widehat{\alpha}_m(1,\vec{X})\right] \;= \\
&\quad \frac{1}{h_{1,1}+h_{1,2}}\left[\sum_{k=1}^{h_{1,1}} U_{1,1}(k)V_{1,1}(k) + \sum_{k=1}^{h_{1,2}} \left(U_{1,2}(k)V_{1,2}(k) + U_{2,1}(k)V_{2,1}(k)\right)\right. \\
&\quad + \frac{1}{h_{1,2}} \sum_{k=1}^{h_{1,2}} U_{1,2}(k) \sum_{j=1}^{h_{1,2}} V_{2,1}(j) + \frac{1}{h_{1,2}} \sum_{k=1}^{h_{1,2}} V_{1,2}(k) \sum_{j=1}^{h_{1,2}} U_{2,1}(j) \\
&\quad + \frac{1}{h_{1,2}} \sum_{k=1}^{h_{1,2}} \left(U_{1,2}(k) + U_{2,1}(k)\right) \sum_{j=1}^{h_{2,2}+h_{2,3}} V^{\uparrow}_{2,2}(j) + \frac{1}{h_{1,2}} \sum_{k=1}^{h_{1,2}} \left(V_{1,2}(k) + V_{2,1}(k)\right) \sum_{j=1}^{h_{2,2}+h_{2,3}} U^{\uparrow}_{2,2}(j) \\
&\quad \left. + \frac{2}{1+h_{1,2}} \left(\sum_{l=1}^{h_{2,2}+h_{2,3}} U^{\uparrow}_{2,2}(l)\right)\left(\sum_{l=1}^{h_{2,2}+h_{2,3}} V^{\uparrow}_{2,2}(l)\right) + \frac{h_{1,2}-1}{h_{1,2}+1} \sum_{l=1}^{h_{2,2}+h_{2,3}} U^{\uparrow}_{2,2}(l)V^{\uparrow}_{2,2}(l)\right].
\end{aligned}
$$

Our permuted estimator $\widetilde{\alpha}_m(1,\vec{X})$ is then the expectation of this expression with respect to random permutations on the T_r-cycles for $r = 3, 4, \ldots, s$. However, the only part of the formula that changes under those permutations is the term

$$\frac{h_{1,2}-1}{h_{1,2}+1} \sum_{l=1}^{h_{2,2}+h_{2,3}} U^{\uparrow}_{2,2}(l)V^{\uparrow}_{2,2}(l),$$

which can be expressed as

$$\frac{h_{1,2}-1}{h_{1,2}+1} \left(h_{2,2}+h_{2,3}\right) \widehat{\alpha}(2,\vec{X}_2),$$

where $\vec{X}_2$ is the original path $\vec{X}$ but with all the $1 \to 1$, $1 \to 2$, and $2 \to 1$ segments deleted. (Here, we use the terminology "$i \to j$ segment" to denote a segment of the sample path from an T_i-regeneration to a T_j-regeneration, with no other regenerations in between.) In Figure 2 we demonstrate how this is done for the sample path given in Figure 1. The top picture in Figure 2 shows the original path in Figure 1 with the $1 \to 1$, $1 \to 2$, and $2 \to 1$ segments removed, and the bottom picture shows the sample path $\vec{X}_2$ obtained by concatenating the remaining segments.

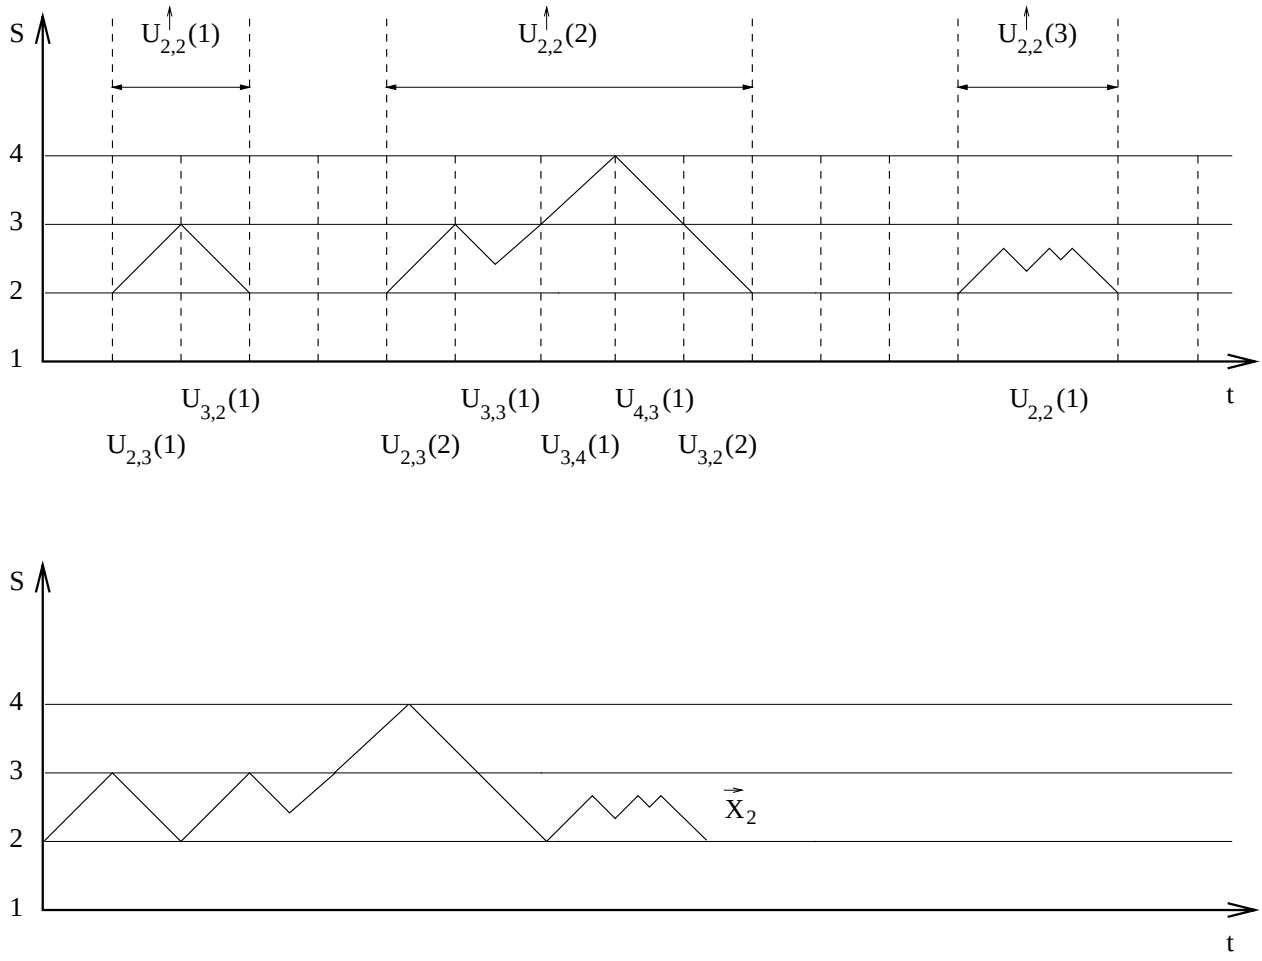

Figure 2: Constructing the new path $\vec{X}_2$.

Taking expectations with respect to uniform random permutations of the T_r-cycles for $r = 3, 4, \ldots, s$, therefore gives the recursive formula

$$\begin{aligned}
\widetilde{\alpha}_m(1,\vec{X}) &= \frac{1}{h_{1,1}+h_{1,2}}\left[\sum_{k=1}^{h_{1,1}} U_{1,1}(k)V_{1,1}(k) + \sum_{k=1}^{h_{1,2}} (U_{1,2}(k)V_{1,2}(k) + U_{2,1}(k)V_{2,1}(k))\right. \\
&\quad + \frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} U_{1,2}(k)\sum_{j=1}^{h_{1,2}} V_{2,1}(j) + \frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} V_{1,2}(k)\sum_{j=1}^{h_{1,2}} U_{2,1}(j) \\
&\quad + \frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} (U_{1,2}(k)+U_{2,1}(k))\sum_{j=1}^{h_{2,2}+h_{2,3}} V_{2,2}^{\uparrow}(j) + \frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} (V_{1,2}(k)+V_{2,1}(k))\sum_{j=1}^{h_{2,2}+h_{2,3}} U_{2,2}^{\uparrow}(j) \\
&\quad \left. + \frac{2}{1+h_{1,2}}\left(\sum_{l=1}^{h_{2,2}+h_{2,3}} U_{2,2}^{\uparrow}(l)\right)\left(\sum_{l=1}^{h_{2,2}+h_{2,3}} V_{2,2}^{\uparrow}(l)\right) + \frac{h_{1,2}-1}{h_{1,2}+1}\,(h_{2,2}+h_{2,3})\,\widetilde{\alpha}_m(2,\vec{X}_2)\right] \\
&= \frac{1}{h_{1,1}+h_{1,2}}\left(\beta(1) + \frac{h_{1,2}-1}{h_{1,2}+1}\,(h_{2,2}+h_{2,3})\,\widetilde{\alpha}_m(2,\vec{X}_2)\right).
\end{aligned}$$

To derive a non-recursive version of the estimator, set $\gamma(0) = 1$ and for $1 \le i < s$, set

$$\gamma(i) = \left(\frac{h_{i,i+1}-1}{h_{i,i+1}+1}\right)$$

and $\widetilde{\alpha'}(i,\vec{X}) = (h_{i,i}+h_{i,i+1})\,\widetilde{\alpha}_m(i,\vec{X})$. Then observe that

$$\begin{aligned}
\widetilde{\alpha'}(1,\vec{X}) &= \beta(1) + \gamma(1)\widetilde{\alpha'}(2,\vec{X}_2) \\
&= \beta(1) + \gamma(1)\beta(2) + \gamma(1)\gamma(2)\widetilde{\alpha'}(3,\vec{X}_3) \\
&\cdots \\
&= \beta(1) + \gamma(1)\beta(2) + \gamma(1)\gamma(2)\beta(3) + \cdots + \gamma(1)\gamma(2)\cdots\gamma(s-1)\beta(s) \\
&= \sum_{i=1}^{s} \beta(i)\prod_{j=0}^{i-1}\gamma(j), \qquad (8)
\end{aligned}$$

which establishes (7).

The proof of $E[\widetilde{\alpha}_m(1,\vec{X})] = \alpha$ and $\mathrm{Var}(\widetilde{\alpha}_m(1,\vec{X})) \le \mathrm{Var}(\widehat{\alpha}_m(1,\vec{X}))$ is analogous to that in the proof of Theorem 1. ∎

The following pseudocode outlines the calculations for the estimator based on m T_1-cycles. The pseudocode assumes that regenerations can only occur at state-transition epochs.

`Define arrays` $h(i,j), SU(i,j), SV(i,j), SUV(i,j)$, $1 \le i \le s$, $|i-j| \le 1$,

```
and accumulators AU, AV, CYCLES, all initialized to 0.
i = 1;
do {
   simulate sample path until next state transition, updating AU, AV;
   for (j = 1; j <= s; ++j)
      if (T_j-regeneration occurred) {
         h(i,j) = h(i,j) + 1;
         SU(i,j) = SU(i,j) + AU;
         SV(i,j) = SV(i,j) + AV;
         SUV(i,j) = SUV(i,j) + AU * AV;
         i = j;
         AU = 0;
         AV = 0;
         if (j == 1) CYCLES = CYCLES + 1;
      };
} while (CYCLES < m);
i = s;
U^↑ = 0;
V^↑ = 0;
for (i = s; i > 0; --i) {
   U^↑ = U^↑ + SU(i+1, i+1);
   V^↑ = V^↑ + SV(i+1, i+1);
   β(i) = SUV(i,i) + SUV(i,i+1) + SUV(i+1,i)
      + SU(i,i+1) * SV(i+1,i)/h(i,i+1)
      + SV(i,i+1) * SU(i+1,i)/h(i,i+1)
      + (SU(i,i+1) + SU(i+1,i)) * V^↑/h(i,i+1)
      + (SV(i,i+1) + SV(i+1,i)) * U^↑/h(i,i+1)
      + (2/(1 + h(i,i+1))) * U^↑ * V^↑;
}
prod = 1;
α̃_m(1, X⃗) = 0;
for (i = 1; i <= s; ++i) {
```

$\widetilde{\alpha}_m(1,\vec{X}) = \widetilde{\alpha}_m(1,\vec{X}) + \beta(i) * prod;$

$prod = prod * (h(i,i+1) - 1)/(h(i,i+1) + 1);$

}

$\widetilde{\alpha}_m(1,\vec{X}) = \widetilde{\alpha}_m(1,\vec{X})/m;$

The total storage required is proportional to the number of regeneration sequences, s. The work of the do-while-loop is proportional to the sample-path length. The calculation after the do-while-loop requires $O(s)$ work, independent of the simulation run length.

For simplicity we derived the permuted estimator based on simulating a fixed number of T_v-cycles for $v = 1$, but similar estimators can be constructed for other choices of $v > 1$. For example, suppose that in addition to the regeneration sequences indexed by $\{1, 2, \ldots, s\}$, as above, there are additional regeneration sequences indexed by $\{2', 3', \ldots, s'\}$, where the new regeneration sequences have the same acyclic property A2 assumed of the original set; i.e., any $T_{i'}$-regeneration and any $T_{i'+2}$-regeneration are separated by an $T_{i'+1}$-regeneration, and any $T_{i'+2}$-regeneration and any $T_{i'}$-regeneration are separated by an $T_{i'+1}$-regeneration. Then

$$\begin{aligned}
\widehat{\alpha}_m(1,\vec{X}) &= \frac{1}{m}\left(\sum_{k=1}^{h_{1,2}} U(k)V(k) + \sum_{k=1}^{h_{1,2'}} U'(k)V'(k)\right) \\
&= \frac{h_{1,2}}{m}\frac{1}{h_{1,2}}\sum_{k=1}^{h_{1,2}} U(k)V(k) + \frac{h_{1,2'}}{m}\frac{1}{h_{1,2'}}\sum_{k=1}^{h_{1,2'}} U'(k)V'(k),
\end{aligned}$$

where the U', V' variables are the analogs of U, V defined for the regeneration sequences $T_{2'}, \ldots, T_{s'}$. Now construct the permuted estimator, as before, for each of the components separately.

To use Theorem 2 based on T_v-cycles for $v > 1$, renumber the states as follows: $v - 1$ becomes $2'$, $v - 2$ becomes $3'$, $\ldots$, 1 becomes s', $v + 1$ becomes 2, $v + 2$ becomes 3, and so on.

5 Central Limit Theorem

We now prove a central limit theorem for our estimator $\widetilde{\alpha}_m \equiv \widetilde{\alpha}_m(1,\vec{X})$. To establish this result, we will use the following key observation: For any i and k, $U_{i,i}(k)$ is defined over a

single T_1-cycle. Similarly, $U_{i,i+1}(k)$, $U_{i+1,i}(k)$, and $U^{\uparrow}_{i+1,i+1}(l)$ are each defined over a single T_1-cycle. The same holds for the corresponding V variables. This will allow us to derive an expression for the estimator $\widetilde{\alpha}_m$ as a function of the sample mean of i.i.d. vectors defined over T_1-cycles, plus a remainder term that appropriately vanishes. Our approach generalizes a technique developed in Calvin and Nakayama (1997).

To do this, we need more notation. Define $H_{i,i}(n)$ to be the set of k-indices for which $U_{i,i}(k)$ is defined over the n-th T_1-cycle; i.e., $H_{i,i}(n) = \{k \geq 1 : T_1(n-1) \leq T_i(k-1) < T_i(k) \leq T_1(n)$ with no regenerations from any sequence occurring between $T_i(k-1)$ and $T_i(k)\}$. Observe that $H_{i,i}(n)$ is also the set of k-indices for which $V_{i,i}(k)$ is defined over the n-th T_1-cycle. Similarly, define $H_{i,i+1}(n)$ to be the set of k-indices for which $U_{i,i+1}(k)$ is defined over the n-th T_1-cycle. Analogously define the sets $H_{i+1,i}(n)$ and $H^{\uparrow}_{i+1,i+1}(n)$ relative to the variables $U_{i+1,i}(k)$ and $U^{\uparrow}_{i+1,i+1}(k)$, respectively. For example in Figure 1, $H_{2,2}(1) = H_{2,2}(2) = H_{2,2}(3) = \emptyset$ and $H_{2,2}(4) = \{1\}$. Also, $H^{\uparrow}_{2,2}(1) = \{1\}$, $H^{\uparrow}_{2,2}(2) = \{2\}$, $H^{\uparrow}_{2,2}(3) = \emptyset$, and $H^{\uparrow}_{2,2}(4) = \{3\}$,

For each $i = 1, 2, \ldots, s$, define the following variables:

$$
\begin{aligned}
A_{i,1}(n) &= |H_{i,i+1}(n)|, \\
A_{i,2}(n) &= \sum_{k \in H_{i,i}(n)} U_{i,i}(k) V_{i,i}(k), \\
A_{i,3}(n) &= \sum_{k \in H_{i,i+1}(n)} \left(U_{i,i+1}(k) V_{i,i+1}(k) + U_{i+1,i}(k) V_{i+1,i}(k) \right), \\
A_{i,4}(n) &= \sum_{k \in H_{i,i+1}(n)} U_{i,i+1}(k), \\
A_{i,5}(n) &= \sum_{k \in H_{i,i+1}(n)} V_{i+1,i}(k), \\
A_{i,6}(n) &= \sum_{k \in H_{i,i+1}(n)} V_{i,i+1}(k), \\
A_{i,7}(n) &= \sum_{k \in H_{i,i+1}(n)} U_{i+1,i}(k), \\
A_{i,8}(n) &= \sum_{k \in H^{\uparrow}_{i+1,i+1}(n)} V^{\uparrow}_{i+1,i+1}(k), \\
A_{i,9}(n) &= \sum_{k \in H^{\uparrow}_{i+1,i+1}(n)} U^{\uparrow}_{i+1,i+1}(k).
\end{aligned}
$$

Also, define $A(n) = (A_{i,j}(n) : i = 1, 2, \ldots, s;\ j = 1, 2, \ldots, 9)$, $n = 1, 2, \ldots$, and note that $A(n)$ is a random vector defined over the n-th T_1-cycle. (Throughout this paper, all vectors are column vectors.) Hence, $A(1), A(2), \ldots$ are i.i.d. Let $\mu = (\mu_{i,j} : i = 1, 2, \ldots, s;\ j = 1, 2, \ldots, 9) =$

$E[A(1)]$, and let $\Sigma = (\sigma_{(i,j),(k,l)} : i,k = 1,2,\ldots,s;\, j,l = 1,2,\ldots,9)$ be the covariance matrix of $A(1)$, where $\sigma_{(i,j),(k,l)} = \mathrm{Cov}(A_{i,j}(1), A_{k,l}(1))$. For $i = 1,2,\ldots,s$, and $j = 1,2,\ldots,9$, define the sample means $\bar{A}_{i,j}(m) = \frac{1}{m}\sum_{n=1}^{m} A_{i,j}(n)$, and define $\bar{A}(m) = (\bar{A}_{i,j}(m) : i = 1,2,\ldots,s;\, j = 1,2,\ldots,9)$.

Let $a = (a_{i,j} : i = 1,2,\ldots,s;\, j = 1,2,\ldots,9)$ be a vector in $\Re^{9s}$, and we define the function $f : \Re^{9s} \to \Re$ as

$$f(a) = \sum_{i=1}^{s} \left[a_{i,2} + a_{i,3} + \frac{a_{i,4}\, a_{i,5}}{a_{i,1}} + \frac{a_{i,6}\, a_{i,7}}{a_{i,1}} + \frac{(a_{i,4} + a_{i,7}) a_{i,8}}{a_{i,1}} + \frac{(a_{i,6} + a_{i,5}) a_{i,9}}{a_{i,1}} + \frac{2\, a_{i,9}\, a_{i,8}}{a_{i,1}} \right].$$

Let $\nabla f(a)$ be the vector of partial derivatives of f evaluated at a.

In the proof of Theorem 3 below, we show that the asymptotic variance of $\sqrt{m}(\widetilde{\alpha}_m - \alpha)$ is

$$\sigma_\alpha^2 = \nabla f(\mu)^T\, \Sigma\, \nabla f(\mu), \tag{9}$$

where the superscript T denotes transpose. Thus, applying the following approach, we can estimate σ_α from the same simulation that is used to construct the permuted estimator $\widetilde{\alpha}_m$. First generate a sample path of m T_1-cycles and compute $\widetilde{\alpha}_m$. From the same sample path, compute the $\bar{A}_{i,j}(m)$, $i = 1,2,\ldots,s$, $j = 1,2,\ldots,9$. Then compute $\nabla f(\bar{A}(m))$. Also, we estimate the covariance $\sigma_{(i,j),(k,l)}$, $i,k = 1,2,\ldots,s$, $j,l = 1,2,\ldots,9$, by

$$\hat{\sigma}_{(i,j),(k,l)}(m) = \frac{1}{m-1} \sum_{n=1}^{m} \left(A_{i,j}(n) - \bar{A}_{i,j}(m) \right) \left(A_{k,l}(n) - \bar{A}_{k,l}(m) \right).$$

Let $\hat{\Sigma}(m) = (\hat{\sigma}_{(i,j),(k,l)} : i,k = 1,2,\ldots,s;\, j,l = 1,2,\ldots,9)$ be the sample covariance matrix. Then the estimator of σ_α is $\hat{\sigma}_\alpha(m) = \left(\nabla f(\bar{A}(m))^T \hat{\Sigma}(m)\, \nabla f(\bar{A}(m)) \right)^{1/2}$. An asymptotically valid $100(1-\delta)\%$ confidence interval for α is

$$\left[\widetilde{\alpha}_m - \frac{\kappa_\delta \hat{\sigma}_\alpha(m)}{\sqrt{m}},\ \widetilde{\alpha}_m + \frac{\kappa_\delta \hat{\sigma}_\alpha(m)}{\sqrt{m}} \right],$$

where κ_δ is the upper $\delta/2$ critical point of a standard normal distribution (i.e., $P\{N(0,1) \le \kappa_\delta\} = 1 - \delta/2$) and $N(x,y)$ denotes a normal random variable with mean x and variance y.

For $k \ge 1$, let $U(|f_U|,k) = \int_{T_1(k-1)}^{T_1(k)} |f_U(X(t))| dt$ and $V(|f_V|,k) = \int_{T_1(k-1)}^{T_1(k)} |f_V(X(t))| dt$. Let "$\Rightarrow$" denote convergence in distribution. Then the following holds.

Theorem 3 *Assume that* $E\left[U(|f_U|,1)^2 + V(|f_V|,1)^2 + U(|f_U|,1)^2\, V(|f_V|,1)^2\right] < \infty$. *Then as* $m \to \infty$,

(i) $\widetilde{\alpha}_m \to \alpha$ *a.s.;*

(ii) $\sqrt{m}\,(\widetilde{\alpha}_m - \alpha)\,/\hat{\sigma}_\alpha(m) \Rightarrow N(0,1)$.

Proof. For the sake of simplicity, we only prove the result when the functions f_U and f_V defined in (3) and (4), respectively, are non-negative. The results can easily be generalized for arbitrary functions. Note that

$$\begin{aligned} U(n)\,V(n) &= \left(\sum_{i=1}^{s}\left(\sum_{j\in H_{i,i}(n)} U_{i,i}(j) + \sum_{k\in H_{i,i+1}(n)} U_{i,i+1}(k) + \sum_{l\in H_{i+1,i}(n)} U_{i+1,i}(l)\right)\right) \\ &\quad \times \left(\sum_{i=1}^{s}\left(\sum_{j\in H_{i,i}(n)} V_{i,i}(j) + \sum_{k\in H_{i,i+1}(n)} V_{i,i+1}(k) + \sum_{l\in H_{i+1,i}(n)} V_{i+1,i}(l)\right)\right) \\ &\geq A_{i,2}(n). \end{aligned}$$

Thus, our assumption implies that $\mu_{i,2} < \infty$. Similar arguments can be made to establish that $\mu_{i,j} < \infty$ and that $E[A_{i,j}(1)^2] < \infty$ for $j = 3, 4, \ldots, 9$. Note that $A_{i,1}(n)$ is bounded by a geometric random variable, and so it has finite moments of all orders. It then follows that μ and Σ are finite.

Note that $h_{1,1} + h_{1,2} = m$, and so in (7), we can bring the $1/m$ inside the summation and multiply it through to get that

$$\widetilde{\alpha}_m = f(\bar{A}(m)) + R(m), \tag{10}$$

where

$$R(m) = \frac{1}{m}\sum_{i=1}^{s} D_i(m)\left(\sum_{l=1}^{h_{i+1,i+1}+h_{i+1,i+2}} U^{\uparrow}_{i+1,i+1}(l)\right)\left(\sum_{l=1}^{h_{i+1,i+1}+h_{i+1,i+2}} V^{\uparrow}_{i+1,i+1}(l)\right)$$

and

$$D_i(m) = \left(\frac{2}{1+h_{i,i+1}}\prod_{j=1}^{i-1}\left(\frac{h_{j,j+1}-1}{h_{j,j+1}+1}\right)\right) - \frac{2}{h_{i,i+1}}.$$

We will prove that as $m \to \infty$, $f\left(\bar{A}(m)\right) \to \alpha$ a.s., $\sqrt{m}\,\left(f\left(\bar{A}(m)\right) - \alpha\right)/\sigma_\alpha \Rightarrow N(0,1)$, and $\sqrt{m}\,R(m) \to 0$ a.s.

Note that for $i = 1, 2, \ldots, s$, we must have that $P\{T_i(0) < T_1(1)\} > 0$. (If this were not true, then $P\{T_i(0) > T_1(1)\} = 1$ implies that $P\{T_i(0) > T_1(k)\} = 1$ for $k \geq 1$ by the regenerative property. It would then follow that $P\{T_i(0) = \infty\} = 1$, so T_i would not be a sequence of finite regeneration points, contradicting our assumption A1.) Hence, $\mu_{i,1} > 0$. Also, we previously proved that $\mu < \infty$, and so $\bar{A}(m) \to \mu$ a.s., as $m \to \infty$. Since f is

continuous at μ, $f(\bar{A}(m)) \to f(\mu)$ a.s., as $m \to \infty$. Moreover, we can easily show that the function f has a non-zero differential at μ, and so using the corollary on p. 124 of Serfling (1980), we get that

$$\frac{\sqrt{m}}{\sigma_\alpha}\left(f\left(\bar{A}(m)\right) - f(\mu)\right) \Rightarrow N(0,1), \tag{11}$$

as $m \to \infty$, where σ_α is defined in (9).

We now show that $\sqrt{m}\,R(m) \to 0$ a.s., as $m \to \infty$. For each i, j, the strong law of large numbers implies that $\bar{A}_{i,j}(m) \to \mu_{i,j}$ a.s. since $\mu_{i,j} < \infty$ for all i, j. Note that

$$\sqrt{m}\,R(m) = \sum_{i=1}^{s} \sqrt{m}\left[\left(\frac{2m}{1+h_{i,i+1}}\prod_{j=1}^{i-1}\left(\frac{h_{j,j+1}-1}{h_{j,j+1}+1}\right)\right) - \frac{2m}{h_{i,i+1}}\right]\bar{A}_{i,9}(m)\,\bar{A}_{i,8}(m),$$

and we consider separately the summands for $i = 1$ and for $i > 1$. A little algebra shows that the summand for $i = 1$ equals

$$\left[\frac{-2}{\sqrt{m}}\frac{1}{\bar{A}_{i,1}(m)}\frac{m}{h_{i,i+1}+1}\right]\bar{A}_{i,9}(m)\,\bar{A}_{i,8}(m) \;\to\; 0 \cdot \frac{1}{\mu_{i,2}^2}\,\mu_{i,9}\,\mu_{i,8} \;=\; 0$$

a.s., since $\mu_{i,1} = P\{T_i(1) < T_1(1)\} > 0$ by assumption. Now for $i > 1$, the summand equals (after a little algebra)

$$\left[\frac{2m}{1+h_{i,i+1}}\sqrt{m}\left(\frac{\prod_{j=1}^{i-1}(h_{j,j+1}-1) - \prod_{j=1}^{i-1}(h_{j,j+1}+1)}{\prod_{j=1}^{i-1}(h_{j,j+1}+1)}\right) - \frac{2m/(h_{i,i+1}+1)}{\sqrt{m}\bar{A}_{i,1}(m)}\right]\bar{A}_{i,9}(m)\,\bar{A}_{i,8}(m).$$

With probability 1, $2m/(1+h_{i,i+1}) \to 2/\mu_{i,1}$, $(2m/(h_{i,i+1}+1))/(\sqrt{m}\bar{A}_{i,1}(m)) \to 0$, $\bar{A}_{i,9}(m) \to \mu_{i,9}$, and $\bar{A}_{i,8}(m) \to \mu_{i,8}$. Moreover,

$$\begin{aligned}
&\sqrt{m}\left(\frac{\prod_{j=1}^{i-1}(h_{j,j+1}-1) - \prod_{j=1}^{i-1}(h_{j,j+1}+1)}{\prod_{j=1}^{i-1}(h_{j,j+1}+1)}\right) \\
&\quad= \frac{\prod_{j=1}^{i-1}(h_{j,j+1}-1) - \prod_{j=1}^{i-1}(h_{j,j+1}+1)}{m^{i-2}} \times \frac{m^{i-1}}{\prod_{j=1}^{i-1}(h_{j,j+1}+1)} \times \frac{\sqrt{m}}{m}.
\end{aligned} \tag{12}$$

For the first term on the right-hand side of (12), note that

$$\frac{1}{m^{i-2}}\left(\prod_{j=1}^{i-1}(h_{j,j+1}-1) - \prod_{j=1}^{i-1}(h_{j,j+1}+1)\right) = \sum_{k=1}^{i-1}\frac{1}{m^{i-2}}\prod_{\substack{j=1\\ j\neq k}}^{i-1} h_{j,j+1} + \frac{D(m)}{m^{i-2}},$$

where $D(m)$ is a sum of a finite number of terms, each of which is the product of at most $i-3$ of the $h_{j,j+1}$. Hence, the first term in (12) converges to $\sum_{k=1}^{i-1}\prod_{\substack{j=1\\ j\neq k}}^{i-1}\mu_{j,1}$ a.s. The second term in (12) converges to $1/(\prod_{j=1}^{i-1}\mu_{j,1})$ a.s., and the third term in (12) vanishes as $m \to \infty$.

Therefore, (12) goes to 0 a.s., and putting this all together proves that $\sqrt{m}R(m) \to 0$ a.s., as $m \to \infty$.

It then follows that $\sqrt{m}\ (\widetilde{\alpha}_m - f(\mu)) \Rightarrow N(0, \sigma_\alpha^2)$ as $m \to \infty$ by (10), (11), Theorem 4.4 of Billingsley (1968), and the continuous mapping theorem (Theorem 5.1 of Billingsley 1968). Now we prove that $f(\mu) = \alpha$ and $\widetilde{\alpha}_m \to \alpha$ a.s., as $m \to \infty$. We previously showed that $f(\bar{A}(m)) \to f(\mu)$ a.s., and $\sqrt{m}\, R(m) \to 0$ a.s., as $m \to \infty$. Hence, it follows from (10) that $\widetilde{\alpha}_m \to f(\mu)$ a.s., as $m \to \infty$. Since $E[\widetilde{\alpha}_m] = \alpha$ and $\mathrm{Var}[\widetilde{\alpha}_m] \leq \mathrm{Var}[\widehat{\alpha}_m]$ for every $m > 0$ by Theorem 2, it follows that

$$E[\widetilde{\alpha}_m^2] = \mathrm{Var}[\widetilde{\alpha}_m] + \alpha^2 \leq \mathrm{Var}[\widehat{\alpha}_m] + \alpha^2 = \frac{\mathrm{Var}[U(1)V(1)]}{m} + \alpha^2 \leq \mathrm{Var}[U(1)V(1)] + \alpha^2 < \infty$$

for all $m > 0$. It then follows that $\{\widetilde{\alpha}_m : m = 1, 2, \ldots\}$ is uniformly integrable, and so $E[\widetilde{\alpha}_m] \to f(\mu)$ as $m \to \infty$; see the corollary to Theorem 25.12 of Billingsley (1986). Therefore, $f(\mu) = \alpha$ since $E[\widetilde{\alpha}_m] = \alpha$ for all $m > 0$. Hence, $\sqrt{m}\ (\widetilde{\alpha}_m - \alpha)\,/\sigma_\alpha \Rightarrow N(0,1)$ as $m \to \infty$.

Finally, since the mean μ and covariance matrix Σ are finite and since $\mu_{i,1} > 0$, we have that $\nabla f(\bar{A}(m)) \to \nabla f(\mu)$ a.s. and $\hat{\Sigma}(m) \to \Sigma$ a.s., as $m \to \infty$. Therefore, $\hat{\sigma}_\alpha(m)/\sigma_\alpha \to 1$ a.s., as $m \to \infty$, and so $\sqrt{m}\,(\widetilde{\alpha}_m - \alpha)\,/\hat{\sigma}_\alpha(m) \Rightarrow N\,(0,1)$, as $m \to \infty$, by Theorems 4.4 and 5.1 of Billingsley (1968). ∎

6 Numerical Experiments

We now present some numerical results from simulating discrete-time and continuous-time Markov chains 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 one regeneration sequence (i.e., the standard estimators) and the permuted estimators using all of the sequences. We ran different experiments with different choices of return state v; i.e., we let the T_1-sequence correspond to hits to state v for different choices of v. In all experiments we defined the reward function f as $f(i) = i/s$.

The first 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

$$R(i, i+1) = \frac{s/2}{i + s/2} = 1 - R(i, i-1)$$

	Sample variance		
v	Standard	Permuted	Ratio
25	0.0028	0.0022	1.3
20	0.0045	0.0022	2.1
15	0.0136	0.0020	6.7

Table 1: Discrete-time Erlang model, $s = 50$.

	Sample variance		
v	Standard	Permuted	Ratio
50	0.0058	0.0044	1.3
40	0.0115	0.0039	2.9
30	0.2341	0.0035	66.3

Table 2: Discrete-time Erlang model, $s = 100$.

for $0 < i < s$ and $R(0,1) = R(s, s-1) = 1$. This chain is a discrete-time version of the Erlang loss system.

Tables 1–3 contain the results from running three different sets of simulation experiments to estimate the time-average variance constant σ_X^2 with $s = 50$, $s = 100$, and $s = 200$. (We also computed numerically, i.e., not using simulation, the theoretical values of σ_X^2 to be 1.014 when $s = 50$, 1.010 when $s = 100$, and 1.005 when $s = 200$.) Each simulation experiment consisted of 1,000 independent replications. In each table the second (resp., third) column gives the sample variance over the 1,000 replications of the standard (resp., permuted) estimators of σ_X^2. The last column gives the ratio of the sample variances. The number of cycles was adjusted for each choice of return state so that on average about 200,000 transitions were generated in each replication. Thus, the different results are comparable.

In each table the first row corresponds to experiments with v equal to the steady-state mean of the process X. Notice that when v is the steady-state mean, the relative improvement from the permuted estimator over the standard estimator is roughly the same as the

	Sample variance		
v	Standard	Permuted	Ratio
100	0.011	0.0079	1.4
80	0.043	0.0082	5.2

Table 3: Discrete-time Erlang model, $s = 200$.

model size increases. The other rows correspond to experiments with v increasingly farther from the mean. In these cases, the variance of the standard estimator increases with the distance of v from the mean, whereas the variance of the permuted estimator stays about the same. Thus, the amount of variance reduction from permuting increases with the distance of v from the mean (with the maximum variance reduction of a factor of 66).

As noted before, the sample variance of the permuted estimator is relatively insensitive to the choice of return state v, and always results in a variance reduction. This is desirable since when one is simulating a system with more than one regeneration sequence, the variability of estimators of certain performance measures (e.g., the time-average variance constant) can differ greatly depending on the choice of the return state when applying the classical regenerative method. Unfortunately, there are no reliable rules that can *a priori* choose a "good" sequence for the classical regenerative method when more than one sequence is available. Also, if we use the estimator from Calvin and Nakayama (1998a) that exploits only two regeneration sequences, then the amount of variance reduction varies significantly depending on the particular pair of sequences used. On the other hand, it seems that one can avoid this problem by applying our method, which uses all sequences simultaneously. Thus, it appears that our method can be applied as a "black box."

We also simulated a continuous-time Erlang loss system. In this model there are s servers, each having exponentially distributed service times with mean 1. Jobs arrive according to a Poisson process with rate $s/2$. The embedded discrete-time Markov chain of this model is the same as that of the previous example, but now the holding time in state i is exponentially distributed with parameter $i + s/2$; thus, the chain tends to spend more time in the lower states. The simulation results for the continuous-time model with $s = 50$ are shown in Table 4. We again computed the sample variances of the estimates of the time-average variance constant over 1,000 independent replications. The sample variances in Table 4 are less than the corresponding values in Table 1, since the holding times tend to be much less than 1. (If all the holding times are exponentially distributed with mean 1, then the sample variances of the estimators of σ_X^2 are similar to those in Table 1.) However, the amounts of variance reduction obtained by permuting in Tables 1 and 4 are about the same. Thus, it seems that our method is equally effective for simulations of discrete-time and continuous-time processes.

	Sample variance		
v	Standard	Permuted	Ratio
25	1.37×10^{-6}	9.80×10^{-7}	1.40
20	1.39×10^{-6}	6.26×10^{-7}	2.21
15	4.37×10^{-6}	6.14×10^{-7}	7.11

Table 4: Continuous-time Erlang model, $s = 50$.

Acknowledgements

The authors would like to thank Phil Heidelberger and Paul Glasserman for interesting discussions related to this paper. The work of the first author was supported in part by the National Science Foundation under grants DMI-9500173 and DMI-9900117. The work of the second author was supported in part by the National Science Foundation under grants DMI-9624469 and DMI-9900117. The second author performed part of the work while visiting the Management Science Division of the Columbia Business School.

Appendix

Here we explain Example 2 more fully. Let the process X correspond to the number of total customers in an $M/G/1$ queue with first-come-first-served service discipline and traffic intensity $\rho < 1$. We define the regeneration sequences $T_1, T_2, T_3, \ldots$, in the following manner. Assume $X(0) = 0$. First define the sequence $A_1 = \{A_1(i) : i = 0, 1, 2, \ldots\}$ where $A_1(0) = 0$, and $A_1(i)$ is the time of the i-th transition from state 1 to state 0. Also, define the sequences $A_j = \{A_j(i) : i = 0, 1, 2, \ldots\}$, $j = 2, 3, 4, \ldots$, where $A_j(i)$ is the time of the $(i+1)$st transition from state j to state $j-1$. Note that each sequence A_j is a regeneration sequence. For example, see Figure 3.

Now we define the sequences $B = \{B(i) : i = 0, 1, 2, \ldots\}$ and $C = \{C(i) : i = 0, 1, 2, \ldots\}$ as follows. Let $B(0) = 0$ and $C(0) = 1$. Then for $i \geq 1$, recursively define $B(i) = \inf\{t > B(i-1) : t = A_j(k)$ for some $k \geq 0$ and $j \geq 1$ such that $j \in \{C(i-1)-1, C(i-1), C(i-1)+1\}\}$, and $C(i) = j$ if and only if $B(i) = A_j(k)$ for some $k \geq 0$. Thus, if $B(i-1)$ corresponds to a time from the sequence A_l (i.e., if $C(i-1) = l$), then $B(i)$ is the first time after time $B(i-1)$ that comes from either A_{l-1}, A_l, or A_{l+1}, and $C(i)$ keeps track of which A_j-sequence

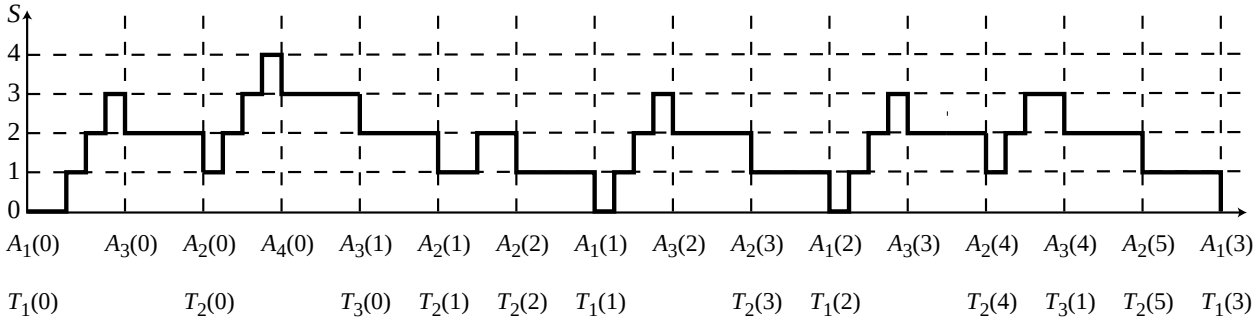

Figure 3: A sample path of the total number of customers in an $M/G/1$ queue.

the time $B(i)$ came from. Now we define the sequence $D_j = \{D_j(i) : i = 0, 1, 2, \ldots\}$ such that $D_j(0) = \inf\{l \geq 0 : C(l) = j\}$ and recursively let $D_j(i) = \inf\{l > D_j(i-1) : C(l) = j\}$ for $i \geq 1$. Hence, $D_j(i)$ is the index l such that $C(l)$ is the $(i+1)$th time that the C-sequence takes on the value j. Finally we define the sequences T_j, $j = 1, 2, \ldots$, with $T_j(i) = B(D_j(i))$, $i = 0, 1, 2, \ldots$. Thus, the sequence T_j contains those times from the sequence A_j that correspond to some time in the sequence B.

In Figure 3, the first regeneration from any of the A_j-sequences, $j = 1, 2, 3, \ldots$, corresponds to time $A_1(0) = 0$, and we define $B(0) = A_1(0) = 0$. Since $B(0)$ corresponds to a time from the A_1-sequence, we define $C(0) = 1$. Also, $D_1(0) = 0$ because 0 is the index of the first time that the C-sequence equals 1. Thus, $T_1(0) = B(D_1(0)) = A_1(0)$. Next, to define $B(1)$ and $C(1)$, since $C(0) = 1$, we look for the next A_j-regeneration for which $j = 1$ or $j = 2$ (there is no A_j-sequence for $j = 0$). Thus, we skip over the time $A_3(0)$, and we use the time corresponding to $A_2(0)$. We define $B(1) = A_2(0)$, and let $C(1) = 2$ since we use a regeneration from the A_2-sequence to define $B(1)$. Also, let $D_2(0) = 1$ since 1 is the index of the first time that the C-sequence equals 2. Thus, $T_2(0) = B(D_2(0)) = A_2(0)$. Next, to define $B(2)$ and $C(2)$, since $C(1) = 2$, we look for the next A_j-regeneration for which $j = 1$, 2, or 3. Hence, we skip over $A_4(0)$, and use $A_3(1)$. We define $B(2) = A_3(1)$, and let $C(2) = 3$ since we use a regeneration from the A_3-sequence to define $B(2)$. Also, let $D_3(0) = 2$ since 2 is the index of the first time that the C-sequence equals 3. Thus, $T_3(0) = B(D_3(0)) = A_3(1)$. Next, to define $B(3)$ and $C(3)$, since $C(2) = 3$, we look for the next A_j-regeneration for which $j = 2$, 3, or 4. Hence, we use $A_2(1)$. We define $B(3) = A_2(1)$, and let $C(3) = 2$ since we use a regeneration from the A_2-sequence to define $B(3)$. Also, let $D_2(1) = 3$ since 3 is the index of the second time that the C-sequence equals 2. Thus, $T_2(1) = B(D_2(1)) = A_2(1)$.

Note that $T_j \subset A_j$ for all $j = 1, 2, \ldots$, so each T_j-sequence is a thinned version of the A_j-sequence. It can be shown that each T_j-sequence is a regeneration sequence, and that $T_1 = A_1$ and $T_2 = A_2$. Also, because of the way we constructed the sequences T_j by thinning the A_j-sequences, Assumption 2 will hold.

On the other hand, Assumption 2 does *not* hold if we define each sequence $T_j = A_j$ since the process can have an A_j-regeneration followed immediately by an A_{j+2}-regeneration without an intervening A_{j+1}-regeneration. For example, at the beginning of Figure 3, an A_1-regeneration is immediately followed by an A_3-regeneration without an intervening A_2-regeneration.

References

Billingsley, P. 1968. *Convergence of Probability Measures*. New York: John Wiley.

Billingsley, P. 1986. *Probability and Measure, Second Edition*. New York: John Wiley.

Calvin, J. M. and M. K. Nakayama. 1997. Central limit theorems for permuted regenerative estimators. *Operations Research*, to appear. Available from `http://www.cis.njit.edu/marvin`

Calvin, J. M. and M. K. Nakayama. 1998a. 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. 1998b. Permuted product and importance-sampling estimators for regenerative simulations. Submitted for publication. Available from `http://www.cis.njit.edu/marvin`

Calvin, J. M. and M. K. Nakayama. 1998c. Exploiting multiple regeneration sequences in simulation output analysis. *Proceedings of the 1998 Winter Simulation Conference*, Washington, D.C., (D. J. Medeiros, E. F. Watson, J. S. Carson, M. S. Manivannan, eds.), IEEE Press, 695–700.

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

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

Serfling, R. J. 1980. *Approximation Theorems of Mathematical Statistics*. New York: John Wiley.

Shedler, G. S. 1993. *Regenerative Stochastic Simulation*. San Diego: Academic Press.