Using Permutations in Regenerative Simulations to Reduce Variance

James M. Calvin
and
Marvin K. Nakayama
New Jersey Institute of Technology

Simulation and Modeling Laboratory, Department of Computer and Information Science, New Jersey Institute of Technology, Newark, NJ 07102-1982, U.S.A., Tel: (973) 596-3398; Fax: (973) 596-5777, e-mail: marvin@cis.njit.edu

We propose a new estimator for a large class of performance measures obtained from a regenerative simulation of a system having two distinct sequences of regeneration times. To construct our new estimator, we first generate a sample path of a fixed number of cycles based on one sequence of regeneration times, divide up the path into segments based on the second sequence of regeneration times, permute the segments, and calculate the performance on the new path using the first sequence of regeneration times. We average over all possible permutations to construct the new estimator. This strictly reduces variance when the original estimator is not simply an additive functional of the sample path. To use the new estimator in practice, the extra computational effort is not large since one does not actually have to compute all permutations as we derive explicit formulas for our new estimators. We examine the small-sample behavior of our estimators. In particular, we prove that for any fixed number of cycles from the first regenerative sequence, our new estimator has smaller mean squared error than the standard estimator. We show explicitly that our method can be used to derive new estimators for the expected cumulative reward until a certain set of states is hit and the time-average variance parameter of a regenerative simulation.

Categories and Subject Descriptors: G.3 [**Mathematics of Comuting**]: Probability and Statistics—*probabilistic algorithms (including Monte Carlo); statistical computing*; I.6.6 [**Computing Methodologies**]: Simulation and Modeling—*Simulation Output Analysis*

General Terms: Algorithms, Experimentation, Performance, Theory

Additional Key Words and Phrases: Regenerative simulation, permutations, efficiency improvement, variance reduction.

The work of the first author was supported by the National Science Foundation under grant DMI-9500173 and by NJIT SBR Grant 421860. The work of the second author was supported by the National Science Foundation under grant DMI-9624469 and by NJIT SBR Grant 421180.
Address: Simulation and Modeling Laboratory, Department of Computer and Information Science, New Jersey Institute of Technology, Newark, NJ 07102-1982, U.S.A., e-mail: calvin@cis.njit.edu, marvin@cis.njit.edu

1. INTRODUCTION

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

In this paper we propose a new simulation estimator for a performance measure of a regenerative process having two different sequences of regeneration times, and study its small-sample behavior. The idea of our approach is as follows. First simulate a fixed number of regenerative cycles from the first sequence of regeneration times, and compute one estimate. We construct another estimator by dividing up the original sample path into segments with endpoints given by the second sequence of regeneration times, and creating a new sample path by permuting the segments (except for the initial and final segments). We then compute a second estimate of α from the new permuted path. We show that this estimate has the same distribution as the original one. Our new estimator is finally constructed as the average of the estimates over all possible permutations. This strictly reduces variance when the estimator is not a purely additive function of the sample path. We show that to compute our new estimators, one does not have to actually calculate all permutations and the average over all of them. Instead, we derive formulas for the new estimators, where the expressions can be easily computed by accumulating some extra quantities during the simulation. The storage requirements of our methods are fixed and do not grow as the simulation run length increases. Hence, there is little extra computational effort or storage needed to construct our new estimators.

For a run length of any fixed number of cycles from the first regenerative sequences, the new estimator has the same expected value as the standard estimator and lower variance; thus, it has lower mean squared error. While it turns out that 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, and in this paper, we consider three specific examples.

First, we derive a new estimator for the second moment of the cumulative reward during a regenerative cycle. We show that the standard regenerative variance estimator fits into this framework. 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 an 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.

We also construct a new estimator for the cumulative reward until some set of states is first hit, which includes the mean time to failure as a special case. Here, the performance measure can be expressed as a ratio of expectations, and we apply our technique to the numerator and denominator separately.

In some sense our method reuses the collected data to construct a new estimator, and as such, it is related to other statistical techniques. For example, the bootstrap [Efron 1979] takes a given sample and resamples the data with replacement. In contrast, one can think of our approach as resampling the data without replacement (i.e., permuting the data), and then averaging over all possible resamples. Other related methods include U-statistics (Chapter 5 of [Serfling 1980]), V-statistics [Sen 1977], and permutation tests (e.g., [Conover 1980]).

The rest of the paper is organized as follows. In Section 2, we discuss our assumptions and the standard estimator of a generic performance measure α. We present the basic idea of how to construct our new estimator using a simple example in Section 3. Section 4 contains a more formal description of our method. Section 5 describes the new estimator for the second moment of the cumulative reward over a regenerative cycle and shows how these results can be used to derive a new estimator of the variance parameter arising in a regenerative simulation. We also discuss here the special case of continuous-time Markov chains. In Section 6 we derive new estimators for the expected cumulative reward until some set of states is hit. We analyze the storage and computational costs of our new estimator in Section 7. We present in Section 8 the results of some simulation experiments comparing our new estimators with the standard ones. Section 9 discusses directions for future research. Most of the proofs are collected in the appendix. Also, we give pseudo-code for one of our estimators in the appendix. (Calvin and Nakayama [1997b] present the basic ideas of our approach, without proofs, in the setting of discrete-time Markov chains.)

2. GENERAL FRAMEWORK

Let X be a continuous-time stochastic process having sample paths that are right continuous with left limits on a state space $S \subset \Re^d$. Note that 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 *regenerative process* (in the classic sense) if

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

(ii). for any $i \geq 0$, $\theta_{T(i)}(X, T)$ 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 more details. This definition allows for so-called delayed regenerative processes (e.g., Section 2.6 of [Kingman 1972]).

Let $T_1 = \{T_1(i) : i = 0, 1, 2, \ldots\}$ with $T_1(0) = 0$ and $T_2 = \{T_2(i) : i = 0, 1, 2, \ldots\}$ be two distinct increasing sequences of nonnegative finite stopping times such that (X, T_1) and (X, T_2) are both regenerative processes. For example, if X is an irreducible, positive-recurrent, discrete-time or continuous-time Markov chain on a countable state space S, then we can define T_1 and T_2 to be the sequences of hitting times to the states $v \in S$ and $w \in S$, respectively, where we assume that $X(0) = v$ and $w \neq v$.

Our goal is to estimate some performance measure α, which we will do by generating a sample path segment $\vec{X} = \{X(t) : 0 \leq t < T_1(m_1)\}$ of a fixed number m_1 of regenerative 1-cycles of our regenerative process. Here, we use the terminology "1-cycles" to denote cycles determined by the sequence T_1; i.e., the ith 1-cycle is the path segment $\{X(t) : T_1(i-1) \leq t < T_1(i)\}$. We similarly define "2-cycles" relative to the sequence T_2. Now we define the standard estimator of α based on the sample path $\vec{X}$ of m_1 1-cycles to be

$$\widehat{\alpha}(\vec{X}) = h(\vec{X}) \tag{1}$$

where $h \equiv h_{m_1}$ is some function. This general framework includes many performance measures of interest.

Example 1. Suppose

$$\alpha = E\left[\left(\int_{T_1(0)}^{T_1(1)} g(X(t))\,dt\right)^p\right]$$

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

$$h(\vec{X}) = \frac{1}{m_1}\sum_{k=1}^{m_1} Y(g;k)^p = \widehat{\alpha}(\vec{X}),$$

where

$$Y(g;k) \equiv Y_1(g;k) = \int_{T_1(k-1)}^{T_1(k)} g(X(t))\,dt \tag{2}$$

for $k \geq 1$. Note that here $\widehat{\alpha}(\vec{X})$ is an unbiased estimator of α. We will examine this example with $p = 2$ in Section 5.

Example 2. Suppose that

$$\alpha = \sigma^2 = \frac{E[Z(f;1)^2]}{E[\tau(1)]}, \tag{3}$$

where for $k \geq 1$,

$$\begin{aligned} \tau(k) \;&\equiv\; \tau_1(k) \;=\; T_1(k) - T_1(k-1), \\ Z(f;k) \;&=\; Y(f;k) - r\tau(k), \end{aligned}$$

$f : S \to \Re$ is some "cost" function, and $r = E[Y(f;1)]/E[\tau(1)]$. Observe that α now is the variance parameter arising from a regenerative simulation. (More details are given in Section 5.1.) Then we can define $h(\vec{X})$ by

$$h(\vec{X}) = \frac{\sum_{k=1}^{m_1}(Y(f;k) - \widehat{r}\tau(k))^2}{\sum_{k=1}^{m_1}\tau(k)} = \widehat{\alpha}(\vec{X}), \tag{4}$$

where

$$\widehat{r} = \frac{\sum_{k=1}^{m_1} Y(f;k)}{\sum_{k=1}^{m_1}\tau(k)}.$$

Note that $\widehat{\alpha}(\vec{X})$ is the standard regenerative estimator of σ^2. We will return to this example in Section 5.1.

Example 3. Suppose we are interested in computing

$$\eta = E\left[\int_0^{T_F} g(X(t))\,dt\right], \tag{5}$$

where $T_F = \inf\{t > 0 : X(t) \in F\}$ for some set of states $F \subset S$. Thus, η is the expected cumulative reward until hitting F conditional on $T_1(0) = 0$, and the mean time to failure is a special case. It can be shown that

$$\eta = \frac{\xi}{\gamma}, \tag{6}$$

where

$$\xi = E\left[\int_0^{T_F \wedge T_1(1)} g(X(t))\,dt\right]$$

and

$$\gamma = E\left[1\{T_F < T_1(1)\}\right],$$

with $a \wedge b = \min(a, b)$; e.g., see [Goyal et al. 1992]. To estimate η, we generate sample paths $\vec{X}_1$ and $\vec{X}_2$, each consisting of m_1 1-cycles, and we use $\vec{X}_1$ to estimate ξ and $\vec{X}_2$ to estimate γ. We can either let $\vec{X}_1 = \vec{X}_2$ or take $\vec{X}_1$ independent of $\vec{X}_2$.

We examine the estimation of the numerator and denominator in (6) separately. First, if we want to estimate $\alpha = \xi$, then we define the function h by

$$h(\vec{X}_1) = \frac{1}{m_1}\sum_{k=1}^{m_1} D(k) = \widehat{\xi}(\vec{X}_1),$$

where

$$D(k) = \int_{T_1(k-1)}^{T_1(k)\wedge T'_F(k)} g(X(t))\,dt$$

with $T'_F(k) = \inf\{t > T_1(k-1) : X(t) \in F\}$. On the other hand, if we want to estimate $\alpha = \gamma$, then we define the function h by

$$h(\vec{X}_2) = \frac{1}{m_1}\sum_{k=1}^{m_1} I(k) = \widehat{\gamma}(\vec{X}_2),$$

where

$$I(k) = 1\{T'_F(k) < T_1(k)\}$$

and $1\{\cdot\}$ is the indicator function of the event $\{\cdot\}$. Thus, the standard estimator of η is

$$\widehat{\eta}(\vec{X}_1, \vec{X}_2) = \frac{\widehat{\xi}(\vec{X}_1)}{\widehat{\gamma}(\vec{X}_2)}. \tag{7}$$

We will return to this example in Section 6.

3. BASIC IDEA

Our goal now is to create a new estimator for α. We begin by giving a heuristic explanation of how it is constructed by considering the simple example illustrated in Figure 1. For simplicity, we depict a continuous sample path on a continuous state space S. The T_1 sequence corresponds to hits to the state v, and the T_2 sequence corresponds to hits to state w. The top graph shows the original sample path generated having $m_1 = 5$ regenerative 1-cycles. For this path, there are $M_2 = 4$ occurrences of stopping times from sequence T_2. To make it easier to see the individual 2-cycles, each is depicted using a different line style. Now we can construct new sample paths from the original path by permuting the 2-cycles, resulting in $(M_2 - 1)! = 6$ total possible paths. The second graph shows one such permuted path. Here, the original third 2-cycle is now first, the original first 2-cycle is now second, and the original second 2-cycle is now third. The new 1-cycle times are $T_1'(i)$, $i = 0, 1, \ldots, 5$, and the new 2-cycle times are $T_2'(j)$, $j = 0, 1, 2, 3$. The third graph contains another permuted path, in which the original second 2-cycle is now first, the original third 2-cycle is now second, and the original first 2-cycle is now third. The 1-cycle times are now $T_1''(i)$, $i = 0, 1, \ldots, 5$, and the new 2-cycle times are $T_2''(j)$, $j = 0, 1, 2, 3$. Note that for each new path, the number of 1-cycles is the same as in the original path, but the paths of some of the 1-cycles have changed. We show in Theorem 1 that all of the paths have the same distribution. For each possible path, we can compute an estimator of α based on the m_1 new 1-cycles by applying the performance function $h \equiv h_{m_1}$ to it. Our new estimator is then the average over all $(M_2 - 1)!$ estimators constructed.

It turns out that we do not actually have to construct all permuted paths to calculate the value of our new estimator. The basic reason for this is that we can break up any sample path into a collection of segments of different types. After any permutation, the path changes, but the collection of segments does not. To calculate our new estimator for a given (original) path, we need to determine the different ways the segments can be put together when 2-cycles are permuted. In particular, since we form an estimator based on the 1-cycles for every permutation, we want to understand how 1-cycles are formed from the segments.

Another key factor that will allow us to explicitly compute our new estimator without actually computing all permutations is that for the performance measures we consider, the contribution of each 1-cycle to the overall estimator can be expressed as a function of the segments in the cycle. For instance, in Example 1 with $p = 2$, we can express the contribution from each 1-cycle as the square of the sum of contributions of the segments in the cycle.

We now examine more closely the four different types of path segments that can arise. We focus on the example in Figure 1.

(1) The first type is a 1-cycle that does not contain a hit to w. The segments of this type in the original path of the figure are the first and third 1-cycles; i.e., the segment from $T_1(0)$ to $T_1(1)$ and the segment from $T_1(2)$ to $T_1(3)$. Segments of this type never change under permutation, although they may occur at different times. For example, the third 1-cycle in the original path appears as the fourth 1-cycle in the second permuted path. This segment is the third 1-cycle in the first permuted path, but it occurs at a different time. The first 1-cycle in our

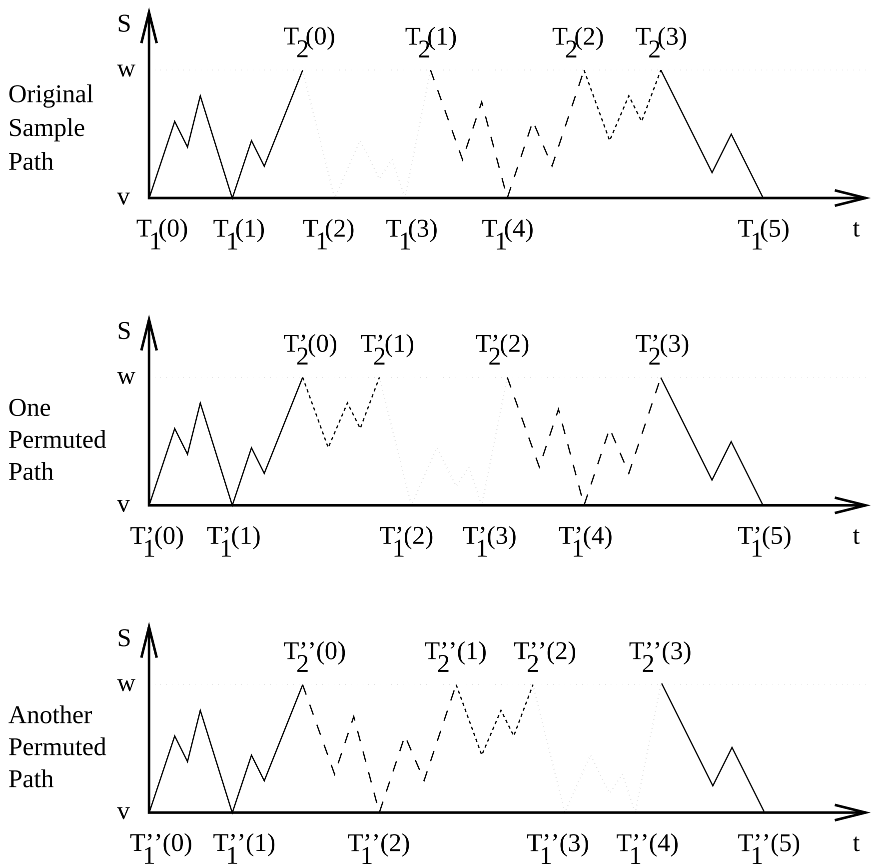

Fig. 1. A sample path and some corresponding permuted paths.

example always appears in the same place in all permutations.

(2) Now consider any 2-cycle in which state v is not hit, such as the third 2-cycle in the original path in the figure. After any permutation this 2-cycle will be in the interior of some 1-cycle. For example, the third 2-cycle in the original path is in the interior of the fifth 1-cycle in the original path, and in the interior of the second (resp., third) 1-cycle in the first (resp., second) permuted path.

(3) The next type of segment goes from w to v before hitting w again. No matter how the 2-cycles are permuted, this type of segment is always the end of some 1-cycle. For example, consider the path segment from $T_2(0)$ to $T_1(2)$ in the original sample path. In this path, the segment is the end of the second 1-cycle. In the first permuted path, this segment is again the end of the second 1-cycle, but this new second 1-cycle is different from that in the original path. On the other hand, this segment in the second permuted path is the end of the third 1-cycle. In general, any segment that goes from w to v before hitting w again will be the end of some 1-cycle in any permuted path.

(4) The final type of segment goes from v to w before hitting v again. In any permutation, this segment will be the beginning of a 1-cycle. For example, consider the path segment from $T_1(3)$ to $T_2(1)$ in the original sample path. In this path, the segment is the beginning of the fourth 1-cycle. In the second permuted path this segment is the beginning of the fifth 1-cycle. In the first permuted path, the segment is again the beginning of the fourth 1-cycle. In general, any segment that goes from v to w before hitting v again will be the beginning of some 1-cycle in any permuted path.

Note that the original sample path in Figure 1 consists of segments of types appearing in the following order: 1, 4, 3, 1, 4, 3, 4, 2, 3. In any permutation, the segments will appear in a different order, but the collection of segments never changes.

Recall that for each permuted path, we compute an estimate of the performance measure α based on the m_1 1-cycles. So now we examine how 1-cycles can be constructed from the permutations. For the original path, we divide up the 1-cycles into those that hit state w and those that do not. The ones that do not hit w are type-1 segments and are unaffected by permutations.

Now we examine how permutations affect 1-cycles that hit w. These cycles always start with a type-4 segment, followed by some number (possibly zero) of segments of type 2, and end with a type-3 segment. For example, the fifth 1-cycle in the original path starts with the type-4 segment from $T_1(4)$ to $T_2(2)$, is followed by the type-2 segment from $T_2(2)$ to $T_3(3)$, and concludes with the type-3 segment from $T_2(3)$ to $T_1(5)$. Also, the fourth 1-cycle in the original path begins with a type-4 segment from $T_1(3)$ to $T_2(1)$ and terminates with a type-3 segment from $T_2(1)$ to $T_1(4)$. This characterization of 1-cycles holds not only for the original sample path, but also for any permuted path. Moreover, for any 2-cycle that hits v in the original path, the type-3 segment and type-4 segment in it will always be in the same 2-cycle in any permutation, and so these two segments can never be in the same 1-cycle since the type-3 segment will always be the end of one 1-cycle and the type-4 segment will always be the beginning of the following 1-cycle. For example, the type-3 segment from $T_2(0)$ to $T_1(2)$ and the type-4 segment from $T_1(3)$ to $T_2(1)$

are always in the same 2-cycle in any permutation, and as such, they are always in successive 1-cycles. Also, in our example the first type-4 segment from $T_1(1)$ to $T_2(0)$ and the last type-3 segment from $T_2(3)$ to $T_1(5)$ will never be in the same 1-cycle. Any other pair of type-3 segment and type-4 segment will be in the same 1-cycle in some permutation. Thus, to construct all 1-cycles that hit w that are possible under permutations of 2-cycles, we have to consider all valid pairings of the type-4 and type-3 segments, and allocate the type-2 segments among the pairs. The proofs in Sections A.2 and A.3 basically use this reasoning.

4. FORMAL DEVELOPMENT OF GENERAL METHOD

Now we more formally show how to construct our new estimator. We begin with some new notation. Let $\mathcal{X}$ be the space of paths $x : [0, \zeta(x)) \subset [0, \infty) \to S$ that are right continuous with left limits on $[0, \zeta(x))$. For $x_1, x_2 \in \mathcal{X}$ define a new element $x_1 \circ x_2$ by

$$\zeta(x_1 \circ x_2) = \zeta(x_1) + \zeta(x_2),$$

and

$$x_1 \circ x_2(t) = \begin{cases} x_1(t) & \text{if } 0 \le t < \zeta(x_1) \\ x_2(t - \zeta(x_1)) & \text{if } \zeta(x_1) \le t < \zeta(x_1) + \zeta(x_2) \end{cases}.$$

Thus, the new path $x_1 \circ x_2$ is obtained by concatenating x_2 on to the end of x_1.

Given the original sample path $\vec{X}$, which consists of m_1 1-cycles, 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 determining the number of times M_2 that the stopping times from the sequence T_2 occur during the m_1 1-cycles. Note that if $M_2 = 0$ or 1, then the path $\vec{X}$ has no 2-cycles. If $M_2 = 2$, then there is only one 2-cycle. Assume now that $M_2 \ge 3$. Then for the given path $\vec{X}$, we can now look at the $(M_2 - 1)$ 2-cycles in the path. We generate a uniform random permutation of the $(M_2 - 1)$ 2-cycles within the path $\vec{X}$, and this gives us our new sample path $\vec{X}'$, which also has m_1 1-cycles. More specifically, define $M_2 = 1 + \sup\{k : T_2(k) \le T_1(m_1)\}$. If $M_2 \le 2$, then let $\vec{X}' = \vec{X}$. If $M_2 \ge 3$, then we break up the path $\vec{X}$ into

$$\vec{X} = \vec{X}^0 \circ \vec{X}^1 \circ \vec{X}^2 \circ \cdots \circ \vec{X}^{M_2 - 1} \circ \vec{X}^{M_2},$$

where

$$\vec{X}^0 = \{X(t) : 0 \le t < T_2(0) \equiv \zeta(\vec{X}^0)\} \tag{8}$$

is the initial path segment until the first time a stopping time from sequence T_2 occurs,

$$\vec{X}^{M_2} = \{X(T_2(M_2 - 1) + t) : 0 \le t < T_1(m_1) - T_2(M_2 - 1) \equiv \zeta(\vec{X}^{M_2})\} \tag{9}$$

is the final path segment from the last time a stopping time from sequence T_2 occurs until the end of the path, and

$$\vec{X}^k = \{X(T_2(k - 1) + t) : 0 \le t < T_2(k) - T_2(k - 1) \equiv \zeta(\vec{X}^k)\}, \tag{10}$$

is the kth 2-cycle of the original path $\vec{X}$, $k = 1, 2, \ldots, M_2 - 1$. Let $(\pi(1), \pi(2), \ldots, \pi(M_2 - 1))$ be a uniform random permutation of $1, 2, \ldots, M_2 - 1$. Then we define our new path $\vec{X}' = \{X'(t) : 0 \le t < T_1(m_1)\}$ to be

$$\vec{X}' = \vec{X}^0 \circ \vec{X}^{\pi(1)} \circ \vec{X}^{\pi(2)} \circ \cdots \circ \vec{X}^{\pi(M_2-1)} \circ \vec{X}^{M_2},$$

which is the original path $\vec{X}$ with the 2-cycles permuted. Note that $\vec{X}'$ and $\vec{X}$ have the same number m_1 of 1-cycles, and we prove in Section A.1 that $\vec{X} \stackrel{\mathcal{D}}{=} \vec{X}'$.

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

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

which is just the estimator obtained from the new sample path $\vec{X}'$ and is based on m_1 1-cycles (recall that $h \equiv h_{m_1}$). The number of possible paths $\vec{X}'$ we can construct from $\vec{X}$ is $N(\vec{X}) = (M_2 - 1)!$, which depends on $\vec{X}$ and is therefore random. We label these paths $\vec{X}^{(1)} = \vec{X}, \vec{X}^{(2)}, \ldots, \vec{X}^{(N)}$, each of which has the same distribution as $\vec{X}$, and for each one we construct $\widehat{\alpha}(\vec{X}^{(i)})$. We finally define our new estimator for α to be

$$\widetilde{\alpha}(\vec{X}) = \frac{1}{N(\vec{X})} \sum_{i=1}^{N(\vec{X})} h(\vec{X}^{(i)}) = \frac{1}{N(\vec{X})} \sum_{i=1}^{N(\vec{X})} \widehat{\alpha}(\vec{X}^{(i)}). \qquad (11)$$

Another way of looking at our new estimator is as follows. We first generate the original path $\vec{X}$ and use it to construct the $N(\vec{X})$ new paths $\vec{X}^{(1)}, \ldots, \vec{X}^{(N)}$. We then choose one of the new paths at random uniformly from $\vec{X}^{(1)}, \ldots, \vec{X}^{(N)}$, and let this be $\vec{X}'$. Since $\vec{X}' \stackrel{\mathcal{D}}{=} \vec{X}$, we can think of $\widehat{\alpha}(\vec{X}')$ as a standard estimator of α since it has the same distribution as $\widehat{\alpha}(\vec{X})$. Then we construct our new estimator $\widetilde{\alpha}(\vec{X})$ to be the conditional expectation of $\widehat{\alpha}(\vec{X}')$ with respect to the uniform random choice of $\vec{X}'$ given the original path $\vec{X}$. That is, if E_* denotes expectation with respect to choosing $\vec{X}'$ from the uniform distribution on $\vec{X}^{(i)}$, $1 \le i \le N$, given $\vec{X}$, then we write

$$\widetilde{\alpha}(\vec{X}) = E_* \left[\widehat{\alpha}(\vec{X}')\right].$$

Assuming that $E[\hat{\alpha}(\vec{X})^2] < \infty$, the new estimator has the same mean as the original since

$$E[\widetilde{\alpha}(\vec{X})] = E\left[E_*[\widehat{\alpha}(\vec{X}')]\right] = E[\hat{\alpha}(\vec{X}')] = E[\hat{\alpha}(\vec{X})]$$

because $\vec{X}' \stackrel{\mathcal{D}}{=} \vec{X}$. Moreover, decomposing the variance by conditioning on $\vec{X}$ gives us

$$\mathrm{Var}\left(\hat{\alpha}(\vec{X})\right) = \mathrm{Var}\left(\hat{\alpha}(\vec{X}')\right) = \mathrm{Var}\left(E_*[\hat{\alpha}(\vec{X}')]\right) + E\left[\mathrm{Var}(\,\hat{\alpha}(\vec{X}') \,|\, \vec{X})\right].$$

Thus, $E\left[\mathrm{Var}(\,\hat{\alpha}(\vec{X}') \,|\, \vec{X})\right] \ge 0$ implies that the variance of the new estimator $\widetilde{\alpha}(\vec{X}) = E_*\left[\widehat{\alpha}(\vec{X}')\right]$ is no greater than that of the original estimator $\hat{\alpha}(\vec{X})$. This calculation, combined with the fact that $\vec{X} \stackrel{\mathcal{D}}{=} \vec{X}'$ (which will be proved in Section A.1), establishes the following theorem.

THEOREM 1. *Let T_1 and T_2 be two distinct sequences of stopping times, and construct the estimator $\widetilde{\alpha}(\vec{X})$ defined by (11). Assume that $E[\widehat{\alpha}(\vec{X})^2] < \infty$. Then*

$E[\widetilde{\alpha}(\vec{X})] = E[\widehat{\alpha}(\vec{X})]$, and

$$Var[\widetilde{\alpha}(\vec{X})] \leq Var[\widehat{\alpha}(\vec{X})], \tag{12}$$

and so the mean squared error of our new estimator $\widetilde{\alpha}(\vec{X})$ is no greater than that of the original estimator $\widehat{\alpha}(\vec{X})$. Strict inequality is obtained in (12) unless $Var(\,\hat{\alpha}(\vec{X}') \,|\, \vec{X}) = 0$ with probability one.

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 2-cycles. For example, this is the case in Example 1 with $p = 1$ since

$$\begin{aligned} h(\vec{X}) &= \frac{1}{m_1} \int_0^{T_1(m_1)} g(X(t))\,dt \\ &= \frac{1}{m_1} \left(\int_0^{T_2(0)} g(X(t))\,dt + \sum_{k=1}^{M_2-1} \int_{T_2(k-1)}^{T_2(k)} g(X(t))\,dt + \int_{T_2(M_2-1)}^{T_1(m_1)} g(X(t))\,dt \right) \\ &= \frac{1}{m_1} \left(\int_0^{T_2(0)} g(X(t))\,dt + \sum_{k=1}^{M_2-1} \int_{T_2(\pi(k)-1)}^{T_2(\pi(k))} g(X(t))\,dt + \int_{T_2(M_2-1)}^{T_1(m_1)} g(X(t))\,dt \right) \\ &= h(\vec{X}'), \end{aligned}$$

and so $\widetilde{\alpha}(\vec{X}) = \widehat{\alpha}(\vec{X})$. Similarly, by choosing $g(x) \equiv 1$, we see that permuting 2-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 $p > 1$ in Example 1, we have in general that $h(\vec{X}) \neq h(\vec{X}')$, and so typically $\widetilde{\alpha}(\vec{X}) \neq \widehat{\alpha}(\vec{X})$. 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 (11). Finally, applying the above idea separately to the numerator and denominator in the ratio expression for the mean cumulative reward until hitting some set of states F as in Example 3 will result in a new estimator.

5. ESTIMATING THE SECOND MOMENT OF CUMULATIVE CYCLE REWARD

For our new estimator $\widetilde{\alpha}(\vec{X})$ to be computationally efficient, we need to calculate explicitly the conditional expectation in (11) without having to construct all possible permutations. We first do this for Example 1 with $p = 2$; i.e., when

$$\alpha = E[Y(g;1)^2] \tag{13}$$

and our standard estimator of α is

$$\widehat{\alpha}(\vec{X}) = \frac{1}{m_1} \sum_{k=1}^{m_1} Y(k)^2, \tag{14}$$

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

$$\widetilde{\alpha}(\vec{X}) = E_* \left[\frac{1}{m_1} \sum_{k=1}^{m_1} Y'(k)^2 \right], \tag{15}$$

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 (15) in this particular setting, we will divide up the original path into segments using the approach described in Sections 3 and 4. We need some new notation to do this. For our two sequences of stopping times T_1 and T_2, let $H(1;2) \subset \{1, 2, \ldots, m_1\}$ denote the set of indices of the 1-cycles in which a T_2 stopping time occurs, and define the complementary set $J(1;2) = \{1, 2, \ldots, m_1\} - H(1;2)$. More specifically, $H(1;2) = \{i \leq m_1 : T_1(i-1) \leq T_2(j) < T_1(i)$ for some $j\}$. We analogously define the set $H(2;1)$ with the roles of T_1 and T_2 reversed. Let $h_{12} = |H(1;2)|$. For $k \in H(1;2)$, define $T_2'(k) = \inf\{t > T_1(k-1) : T_2(l) = t$ for some $l\}$, which is the first occurrence of a stopping time from sequence T_2 after the $(k-1)$st stopping time from the sequence T_1. Similarly define $\widetilde{T}_2(k) = \sup\{t < T_1(k) : T_2(l) = t$ for some $l\}$, which is the last occurrence of a stopping time from sequence T_2 before the kth occurrence of the stopping-time sequence T_1. Then, for $k \in H(1;2)$, we let

$$Y_{12}(k) = \int_{T_1(k-1)}^{T_2'(k)} g(X(t))\, dt,$$

which is the contribution to $Y(k)$ until a stopping time from sequence T_2 occurs, and let

$$Y_{21}(k) = \int_{\widetilde{T}_2(k)}^{T_1(k)} g(X(t))\, dt,$$

which is the contribution to $Y(k)$ from the last occurrence of a stopping time from sequence T_2 in the kth 1-cycle until the end of the cycle. Also, for $l \in J(2;1)$, let

$$Y_{22}(l) = \int_{T_2(l-1)}^{T_2(l)} g(X(t))\, dt,$$

which is the integral of $g(X(t))$ over the lth 2-cycle in which there is no occurrence of a stopping time from sequence T_1. We now define $B_k \subset J(2;1)$ to be the set of indices of those 2-cycles that do not contain any occurrences of the stopping times from the sequence T_1 and that are between $T_1(k-1)$ and $T_1(k)$. It then follows that for $k \in H(1;2)$,

$$Y(k) = Y_{12}(k) + \sum_{l \in B_k} Y_{22}(l) + Y_{21}(k).$$

Hence,

$$\begin{aligned}
\widehat{\alpha}(\vec{X}) &= \frac{1}{m_1}\left(\sum_{k \in J(1;2)} Y(k)^2 + \sum_{k \in H(1;2)} Y(k)^2\right) \\
&= \frac{1}{m_1}\left(\sum_{k \in J(1;2)} Y(k)^2 + \sum_{k \in H(1;2)} \left[Y_{12}(k) + \sum_{l \in B_k} Y_{22}(l) + Y_{21}(k)\right]^2\right) \\
&= \frac{1}{m_1}\left(\sum_{k \in J(1;2)} Y(k)^2 + \sum_{k \in H(1;2)} \left[Y_{12}(k)^2 + Y_{21}(k)^2\right]\right)
\end{aligned}$$

$$+ \frac{1}{m_1} \sum_{k \in H(1;2)} \left(2Y_{12}(k)Y_{21}(k) + 2[Y_{12}(k) + Y_{21}(k)] \sum_{l \in B_k} Y_{22}(l) + \left(\sum_{l \in B_k} Y_{22}(l) \right)^2 \right). \quad (16)$$

In the last expression, the first term does not change if we replace the original sample path $\vec{X}$ with the new sample path $\vec{X}'$, whereas the last term does change. In Section A.2, we compute explicitly the conditional expectation of (16), when $\vec{X}$ is replaced with $\vec{X}'$, with respect to a random permutation given the original path $\vec{X}$.

The expression for this involves some more notation. Define

$$\bar{Y}_{12} = \frac{1}{h_{12}} \sum_{k \in H(1;2)} Y_{12}(k)$$

and

$$\bar{Y}_{21} = \frac{1}{h_{12}} \sum_{k \in H(1;2)} Y_{21}(k).$$

Finally, define β_l to be the lth smallest element of the set $H(1;2)$ for $l = 1, 2, \ldots, h_{12}$, and define $\beta_0 = \beta_{h_{12}}$. For $k = \beta_l \in H(1;2)$ for some $l = 1, 2, \ldots, h_{12}$, define $\psi(k) = \beta_{l-1}$; i.e., $\psi(k)$ is the index in $H(1;2)$ that occurs just before k if k is not the first index and is the last element in $H(1;2)$ if k is the first element. The following theorem is proved in Section A.2. (Pseudo-code for our estimator is given in the appendix.)

THEOREM 2. *Suppose we want to estimate α defined in (13), and assume that $E[Y(g;1)^4] < \infty$. Then, our new estimator is given by $\widetilde{\alpha}(\vec{X}) = \widehat{\alpha}(\vec{X})$ if $M_2 < 3$, and otherwise by*

$$\begin{aligned} \widetilde{\alpha}(\vec{X}) \;=\; & \frac{1}{m_1} \left(\sum_{k \in J(1;2)} Y(k)^2 + \sum_{k \in H(1;2)} \left[Y_{12}(k)^2 + Y_{21}(k)^2 \right] \right. \\ & + \frac{2}{h_{12} - 1} \sum_{k \in H(1;2)} Y_{12}(k) \left(\sum_{j \in H(1;2)} Y_{21}(j) - Y_{21}(\psi(k)) \right) + \sum_{k \in J(2;1)} Y_{22}(k)^2 \\ & \left. + 2(\bar{Y}_{12} + \bar{Y}_{21}) \sum_{k \in J(2;1)} Y_{22}(k) + \frac{2}{1 + h_{12}} \sum_{\substack{l,m \in J(2;1) \\ l \neq m}} Y_{22}(l) Y_{22}(m) \right). \end{aligned} \quad (17)$$

The estimator satisfies $E[\widetilde{\alpha}(\vec{X})] = \alpha$ and $Var(\widetilde{\alpha}(\vec{X})) \leq Var(\widehat{\alpha}(\vec{X}))$ when $\widehat{\alpha}(\vec{X})$ is the standard estimator of α as defined in (14).

5.1 A New Estimator for the Time-Average Variance

We can use Theorem 2 to construct a new estimator for the variance parameter in a regenerative simulation of the process X. We start by first giving a more complete explanation of Example 2 in Section 2.

Let $f : S \to \Re$ be some cost function. Define

$$r_t = \frac{1}{t} \int_0^t f(X(s))\, ds.$$

Since X is a regenerative process, there exists some constant r such that $r_t \to r$ as $t \to \infty$ with probability 1 (see Theorem 2.2 of [Shedler 1993]). Also, r satisfies the ratio formula $r = E[Y(f;1)]/E[\tau(1)]$. Assuming that $E[Z(f;1)^2] < \infty$, there exists a finite positive constant σ such that

$$t^{1/2}\frac{(r_t - r)}{\sigma} \xrightarrow{\mathcal{D}} N(0,1) \tag{18}$$

as $t \to \infty$. The constant σ^2 is called the *time-average variance* of X and is given in (3). Given the central limit theorem described by (18), 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 $\widehat{\sigma}^2(\vec{X}) = \widehat{\alpha}(\vec{X})$ defined in (4). Note that $\widehat{\sigma}^2(\vec{X})$ can be expressed as

$$\widehat{\sigma}^2(\vec{X}) = \frac{\sum_{k=1}^{m_1} Y(f-\widehat{r};k)^2}{\sum_{k=1}^{m_1} \tau(k)}.$$

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

$$\widetilde{\sigma}^2(\vec{X}) = E_*\left[\frac{\sum_{k=1}^{m_1} Y'(f-\widehat{r}';k)^2}{\sum_{k=1}^{m_1} \tau'(k)}\right] = \frac{E_*\left[\sum_{k=1}^{m_1} Y'(f-\widehat{r}';k)^2\right]}{\sum_{k=1}^{m_1} \tau(k)}$$

since $\sum_{k=1}^{m_1} \tau'(k) = T_1(m_1) = \sum_{k=1}^{m_1} \tau(k)$ is independent of the permutation of 2-cycles. Also, observe that

$$\widehat{r}' = \frac{1}{T_1(m_1)} \int_0^{T_1(m_1)} f(X(s))\, ds = \widehat{r}$$

is independent of the permutation of 2-cycles, so

$$\widetilde{\sigma}^2(\vec{X}) = \frac{E_*\left[\sum_{k=1}^{m_1} Y'(f-\widehat{r};k)^2\right]}{\sum_{k=1}^{m_1} \tau(k)}; \tag{19}$$

i.e., we can replace $\widehat{r}'$ with $\widehat{r}$. The following is a direct consequence of Theorem 2.

COROLLARY 3. *Suppose we want to estimate σ^2 defined in (3), and assume that $E[Z(f;1)^4] < \infty$. Then, our new estimator $\widetilde{\sigma}^2(\vec{X})$ is given by (19), where the numerator is as in (17) with the function $g = f - \widehat{r}$. The estimator satisfies $E[\widetilde{\sigma}^2(\vec{X})] = E[\widehat{\sigma}^2(\vec{X})]$ and $Var[\widetilde{\sigma}^2(\vec{X})] \le Var[\widehat{\sigma}^2(\vec{X})]$.*

5.2 Continuous-Time Markov Chains

We now consider the special case of an irreducible, positive-recurrent, continuous-time Markov chain $X = \{X(s) : s \ge 0\}$ on a countable state space S having generator matrix $Q = \{q(x,y) : x, y \in S\}$. Define $\lambda(x) = -q(x,x)$ for $x \in S$. Let $V = \{V_n : n = 0,1,2,\ldots\}$ be the embedded discrete-time Markov chain, and $W = \{W_n : n = 0,1,2,\ldots\}$ be the sequence of random holding times of the continuous-time Markov chain; i.e., W_n is the time between the nth and $(n+1)$st transitions of X. Define $A_0 = 0$, and for $n \geq 1$, let $A_n = \sum_{k=0}^{n-1} W_k$, which is the time of the nth transition. It is well known that conditional on V, the holding time in state V_n is exponentially distributed with mean $1/\lambda(V_n)$ and that W_i and W_j are (conditionally) independent for $i \neq j$.

Assume that the sequences of stopping times T_1 and T_2 correspond to hitting times to fixed states $v \in S$ and $w \in S$, respectively, with $w \neq v$, and assume that $X(0) = v$. More specifically, define $\xi_1(0) = 0$ and $\xi_1(k) = \inf\{k > \xi_1(k-1) : V_k = v\}$ for $k = 1, 2, \ldots$, which is the sequence of hitting times to state v for the discrete-time Markov chain. Similarly, define $\xi_2(0) = \inf\{k > 0 : V_k = w\}$ and $\xi_2(k) = \inf\{k > \xi_2(k-1) : V_k = w\}$ for $k = 1, 2, \ldots$. Then, we define $T_1(k) = A_{\xi_1(k)}$ and $T_2(k) = A_{\xi_2(k)}$ for $k = 0, 1, 2, \ldots$.

Suppose that we want to estimate α as defined in (13), and our standard estimator of α is given in (14). Now note that

$$Y(g;k) = \int_{T_1(k-1)}^{T_1(k)} g(X(s))\, ds = \sum_{i=\xi_1(k-1)}^{\xi_1(k)-1} g(V_i) W_i.$$

Using discrete-time conversion [Hordijk et al. 1976; Fox and Glynn 1990] gives us

$$\begin{aligned} E[\,Y(g;k)^2 \,|\, V\,] \;=\; & E\left[\sum_{i=\xi_1(k-1)}^{\xi_1(k)-1} g(V_i)^2 W_i^2 + \sum_{\substack{i,j=\xi_1(k-1)\\ i\neq j}}^{\xi_1(k)-1} g(V_i) g(V_j) W_i W_j \,\middle|\, V \right] \\ = & \sum_{i=\xi_1(k-1)}^{\xi_1(k)-1} g(V_i)^2\, E\left[W_i^2 \,\middle|\, V \right] + \sum_{\substack{i,j=\xi_1(k-1)\\ i\neq j}}^{\xi_1(k)-1} g(V_i) g(V_j)\, E\left[\, W_i W_j \mid V \,\right] \\ = & \sum_{i=\xi_1(k-1)}^{\xi_1(k)-1} \frac{2g(V_i)^2}{\lambda(V_i)^2} + \sum_{\substack{i,j=\xi_1(k-1)\\ i\neq j}}^{\xi_1(k)-1} \frac{g(V_i)}{\lambda(V_i)} \frac{g(V_j)}{\lambda(V_j)} \end{aligned}$$

since W_i, W_j, $i \neq j$, are conditionally independent given V. For a function $f : S \to \Re$, let

$$\bar{Y}(f;k) = \sum_{i=\xi_1(k-1)}^{\xi_1(k)-1} f(V_i),$$

which is the cumulative reward over the kth cycle for the discrete-time chain V, and define the functions $g_1 : S \to \Re$ and $g_2 : S \to \Re$ to be $g_1(x) = g(x)/\lambda(x)$ and $g_2(x) = g_1(x)^2$. Therefore, we get

$$E[\,Y(g;k)^2 \,|\, V\,] = \bar{Y}(g_1;k)^2 + \bar{Y}(g_2;k)$$

and

$$\bar{\alpha}(\vec{X}) \equiv E\left[\, \widehat{\alpha}(\vec{X}) \,|\, V \,\right] = \frac{1}{m_1} \sum_{k=1}^{m_1} \bar{Y}(g_1;k)^2 + \frac{1}{m_1} \sum_{k=1}^{m_1} \bar{Y}(g_2;k). \qquad (20)$$

To create our new estimator of α, we then compute the conditional expectation of $\bar{\alpha}(\vec{X})$ with respect to a random uniform permutation of 2-cycles given the original path $\vec{X}$. Define

$$\widetilde{\widetilde{\alpha}}(g_1;\vec{X}) = E_*\left[\frac{1}{m_1}\sum_{k=1}^{m_1}\bar{Y}(g_1;k)^2\right]. \tag{21}$$

The last term in (20) is independent of permutations of 2-cycles, and so we get the following expression for $\widetilde{\widetilde{\alpha}}(\vec{X}) = E_*\left[\bar{\alpha}(\vec{X})\right]$, which follows from Theorems 1 and 2.

Theorem 4. *Suppose X is an irreducible, positive-recurrent, continuous-time Markov chain on a countable state space S, and we want to estimate α defined in (13). Assume that T_1 and T_2 correspond to the hitting times to states v and w, respectively, with $w \neq v$. Assume that $E[Y(g_1;1)^4] < \infty$. Then, our new estimator is given by $\widetilde{\widetilde{\alpha}}(\vec{X}) = \bar{\alpha}(\vec{X})$ if $M_2 < 3$, and*

$$\widetilde{\widetilde{\alpha}}(\vec{X}) = \widetilde{\widetilde{\alpha}}(g_1;\vec{X}) + \frac{1}{m_1}\sum_{k=1}^{m_1}\bar{Y}(g_2;k),$$

where $\widetilde{\widetilde{\alpha}}(g_1;\vec{X})$ is defined by (21), which can be computed from (17) with the function g_1. The estimator satisfies $E[\widetilde{\widetilde{\alpha}}(\vec{X})] = \alpha$ and $Var(\widetilde{\widetilde{\alpha}}(\vec{X})) \leq Var(\widehat{\alpha}(\vec{X}))$ when $\widehat{\alpha}(\vec{X})$ is the standard estimator of α as defined in (14).

If we had instead first converted to discrete time and then computed $\bar{Y}(g;1)^2$ for the discrete-time Markov chain and its conditional expectation with respect to the permutation, we would have obtained $\widetilde{\widetilde{\alpha}}(g_1;\vec{X})$ as our estimator for α. However, since $E[\bar{Y}(g_2;k)] > 0$ for any function $g \neq 0$, $E[\widetilde{\widetilde{\alpha}}(g_1;\vec{X})] \neq \alpha$, so $\widetilde{\widetilde{\alpha}}(g_1;\vec{X})$ is biased. On the other hand, our estimator $\widetilde{\widetilde{\alpha}}(\vec{X})$ is unbiased.

6. EXPECTED CUMULATIVE REWARD UNTIL HITTING A SET

Recall that we can express the expected cumulative reward until a hitting time given in (5) as the ratio $\eta = \xi/\gamma$ in (6), and the standard estimator of η is defined in (7). Also, recall that the numerator ξ is estimated using the sample path $\vec{X}_1$ and the denominator γ is estimated from path $\vec{X}_2$. We will examine both the cases when $\vec{X}_1 = \vec{X}_2$ and when $\vec{X}_1$ and $\vec{X}_2$ are independent.

In the context of estimating the mean time to failure of highly reliable Markovian systems, Goyal, Shahabuddin, Heidelberger, Nicola, and Glynn [1992] and Shahabuddin [1994] estimate ξ and γ independently; i.e., $\vec{X}_1$ and $\vec{X}_2$ are independent. This is useful because then different sampling techniques can be applied to estimate the two quantities. In particular, γ is the probability of a rare event and so it is estimated using importance sampling. On the other hand, we can efficiently estimate ξ using naive simulation (i.e., no importance sampling). Below, we do not apply importance sampling to estimate γ, but one can also derive a new estimator of γ when using importance sampling.

Our new estimator of η is defined as

$$\widetilde{\eta}(\vec{X}_1,\vec{X}_2) = \frac{\widetilde{\xi}(\vec{X}_1)}{\widetilde{\gamma}(\vec{X}_2)}, \tag{22}$$

where

$$\begin{aligned}\widetilde{\xi}(\vec{X}_1) &= E_*[\widehat{\xi}(\vec{X}_1)],\\ \widetilde{\gamma}(\vec{X}_2) &= E_*[\widehat{\gamma}(\vec{X}_2)].\end{aligned}$$

Now to explicitly calculate the numerator and denominator, we will divide up the original path into segments using the approach described in Sections 3 and 4. We need some new notation to do this. For $k \in H(1;2)$, let

$$\begin{aligned}I_{12}(k) &= 1\{T'_F(k) < T'_2(k)\},\\ I_{22}(l) &= 1\{T_F^{(2)}(l) < T_2(l)\},\\ I_{21}(k) &= 1\{T_F^{(1)}(k) < T_1(k)\},\end{aligned}$$

with $T'_2(k) = \inf\{t > T_1(k) : T_2(i) = t \text{ for some } i\}$, $T_F^{(2)}(l) = \inf\{t > T_2(l-1) : X(t) \in F\}$ and $T_F^{(1)}(k) = \inf\{t > \widetilde{T}_1(k) : X(t) \in F\}$. Hence, $I_{12}(k)$ (resp., $I_{21}(k)$) is the indicator of whether the set F is hit in the initial 1-2 segment (resp., final 2-1 segment) of the 1-cycle with index $k \in H(1;2)$. Also, $I_{22}(l)$ is the indicator whether the set F is hit in the 2-cycle with index $l \in J(2;1)$.

We first consider the denominator γ. To derive the new estimator of γ from permuting the 2-cycles, we first write

$$\widehat{\gamma}(\vec{X}_2) = \frac{1}{m_1}\left(\sum_{k\in J(1;2)} I(k) + \sum_{k\in H(1;2)} I(k)\right).$$

The first term on the right-hand side is independent of permutations of the 2-cycles. For the second term we note that for $k \in H(1;2)$,

$$I(k) = \max\left(I_{12}(k),\ \max_{l\in B_k} I_{22}(l),\ I_{21}(k)\right).$$

Thus,

$$\widetilde{\gamma}(\vec{X}_2) = \frac{1}{m_1}\left(\sum_{k\in J(1;2)} I(k) + E_*\left[\sum_{k\in H(1;2)} \max\left(I_{12}(k),\ \max_{l\in B'_{\rho(k)}} I_{22}(l),\ I_{21}(\rho(k))\right)\right]\right), \quad (23)$$

where $\rho(k)$ is the index of the I_{21} variable that follows the $I_{12}(k)$ variable after a permutation of the 2-cycles, and B'_l is the same as B_l except that B'_l is after a permutation. We work out in Section A.3 the conditional expectation appearing above.

We now examine the estimation of ξ. Note that the standard estimator of ξ satisfies

$$\widehat{\xi}(\vec{X}_1) = \frac{1}{m_1}\left(\sum_{k\in J(1;2)} D(k) + \sum_{k\in H(1;2)} D(k)\right). \quad (24)$$

The first term is not affected by permuting the 2-cycles, but the second term is.

For $k \in H(1;2)$,

$$D(k) = D_{12}(k) + \sum_{l \in B_k} D_{22}(l)(1 - I_{12}(k)) \prod_{\substack{j \in B_k \\ j<l}} (1 - I_{22}(j)) + D_{21}(k)(1 - I_{12}(k)) \prod_{l \in B_k} (1 - I_{22}(l)), \tag{25}$$

where

$$\begin{aligned} D_{12}(k) &= \int_{T_1(k-1)}^{T'_F(k) \wedge T'_2(k)} g(X(t))\, dt, \\ D_{22}(l) &= \int_{T_2(l-1)}^{T_F^{(2)}(l) \wedge T_2(l)} g(X(t))\, dt, \\ D_{21}(k) &= \int_{\widetilde{T}_2(k)}^{T_F^{(1)}(k) \wedge T_1(k)} g(X(t))\, dt. \end{aligned}$$

Hence, to compute the new estimator, we need to compute the conditional expectation of the second term in (24), which we can do by using the representation for $D(k)$ given in (25); this is done in Section A.3.

To present what the new estimator actually works out to, we need some more notation. Define

$$\begin{aligned} k_{21} &= |\{k \in H(1;2) : I_{21}(k) = 0\}|, \\ k_{12} &= |\{k \in H(1;2) : I_{12}(k) = 0\}|, \\ k_c &= |\{k \in H(1;2) : I_{12}(k) = 0,\, I_{21}(\psi(k)) = 0\}|. \end{aligned}$$

Also, define

$$\begin{aligned} r &= \sum_{l \in J(2;1)} I_{22}(l), \\ p &= \sum_{l \in J(2;1)} (1 - I_{22}(l)), \\ d_0 &= \frac{1}{p} \sum_{l \in J(2;1)} D_{22}(l)(1 - I_{22}(l)), \\ d_1 &= \frac{1}{r} \sum_{l \in J(2;1)} D_{22}(l) I_{22}(l). \end{aligned}$$

Then we have the following result, whose proof is given in Section A.3.

THEOREM 5. *Suppose we want to estimate η in (6), and assume that $E[D(1)^2] < \infty$. Then, our new estimator is given by (22) where* *(i).* $\widetilde{\xi}(\vec{X}_1) = \widehat{\xi}(\vec{X}_1)$ *if* $M_2(\vec{X}_1) < 3$*, and otherwise by*

$$\begin{aligned}\widetilde{\xi}(\vec{X}_1) &= \frac{1}{m_1}\left(\sum_{k\in J(1;2)} D(k) + \sum_{k\in H(1;2)} D_{12}(k)\right.\\ &+ \frac{1}{h_{12}-1+r}\left(k_{12}\sum_{j\in H(1;2)} D_{21}(j) - \sum_{k\in H(1;2)}(1-I_{12}(k))D_{21}(\psi(k))\right)\\ &+ \left. k_{12}\left(d_0\frac{p}{r+h_{12}} + d_1\frac{r}{r+h_{12}-1}\right)\right);\end{aligned}$$

(ii). $\widetilde{\gamma}(\vec{X}_2) = \widehat{\gamma}(\vec{X}_2)$ *if* $M_2(\vec{X}_2) < 3$*, and otherwise by*

$$\widetilde{\gamma}(\vec{X}_2) = \frac{1}{m_1}\left(\sum_{k\in J(1;2)} I(k) + h_{12} - \frac{k_{12}k_{21}-k_c}{h_{12}-1+r}\right).$$

The estimators $\widetilde{\xi}(\vec{X}_1)$ *and* $\widetilde{\gamma}(\vec{X}_2)$ *satisfy* $E[\widetilde{\xi}(\vec{X}_1)] = \xi$, $E[\widetilde{\gamma}(\vec{X}_2)] = \gamma$, *Var*$(\widetilde{\xi}(\vec{X}_1)) \le$ *Var*$(\widehat{\xi}(\vec{X}_1))$*, and Var*$(\widetilde{\gamma}(\vec{X}_2)) \le$ *Var*$(\widehat{\gamma}(\vec{X}_2))$*, where* $\widehat{\xi}(\vec{X}_1)$ *and* $\widehat{\gamma}(\vec{X}_2)$ *are the standard estimators of* ξ *and* γ*, respectively.*

In Theorem 5 the variables used in part (i) are defined for the sample path $\vec{X}_1$, and the variables in part (ii) are for the sample path $\vec{X}_2$. For example, h_{12} in part (i) is the cardinality of the set $H(1;2)$ for the path $\vec{X}_1$, and in part (ii) it is the same but instead for the path $\vec{X}_2$.

Theorem 5 shows that our new estimator for η has unbiased and lower-variance estimators for both the numerator and denominator, but the effect on the resulting ratio estimator is more difficult to analyze rigorously. Instead, we now heuristically examine the bias and variance of the ratio estimator.

To do this, we generically let $\bar{\xi}$, $\bar{\gamma}$, and $\bar{\eta} = \bar{\xi}/\bar{\gamma}$ be estimators of ξ, γ, and η, respectively. Then using first- and second-order Taylor series expansions, we have the following approximations for the bias and variance of $\bar{\eta}$:

$$\mathrm{Bias}(\bar{\eta}) \approx \frac{E[\bar{\xi}]\,\mathrm{Var}(\bar{\gamma})}{(E[\bar{\gamma}])^3} - \frac{\mathrm{Cov}\left(\bar{\xi},\bar{\gamma}\right)}{(E[\bar{\gamma}])^2} \tag{26}$$

and

$$\mathrm{Var}(\bar{\eta}) \approx \frac{\mathrm{Var}(\bar{\xi})}{(E[\bar{\gamma}])^2} + \frac{(E[\bar{\xi}])^2\,\mathrm{Var}(\bar{\gamma})}{(E[\bar{\gamma}])^4} - \frac{2\,E[\bar{\xi}]\,\mathrm{Cov}\left(\bar{\xi},\bar{\gamma}\right)}{(E[\bar{\gamma}])^3}; \tag{27}$$

see p. 181 of [Mood et al. 1974].

We now use these approximations to analyze the standard and new estimators for η. First, consider the case when $\vec{X}_1$ and $\vec{X}_2$ are independent. Then $\widehat{\xi}(\vec{X}_1)$ and $\widehat{\gamma}(\vec{X}_2)$ are independent, so $\mathrm{Cov}\left(\widehat{\xi}(\vec{X}_1),\widehat{\gamma}(\vec{X}_2)\right) = 0$. Similarly, $\widetilde{\xi}(\vec{X}_1)$ and $\widetilde{\gamma}(\vec{X}_2)$ are independent, so $\mathrm{Cov}\left(\widetilde{\xi}(\vec{X}_1),\widetilde{\gamma}(\vec{X}_2)\right) = 0$. Hence, it follows from Theorem 5 and (26) and (27) that

$$\mathrm{Bias}(\widetilde{\eta}(\vec{X}_1,\vec{X}_2)) \preceq \mathrm{Bias}(\widehat{\eta}(\vec{X}_1,\vec{X}_2))$$

and

$$\mathrm{Var}(\widetilde{\eta}(\vec{X}_1, \vec{X}_2)) \preceq \mathrm{Var}(\widehat{\eta}(\vec{X}_1, \vec{X}_2)),$$

where we use the notation $a \preceq b$ to mean that a is approximately no greater than b. Hence, the mean square error of $\widetilde{\eta}(\vec{X}_1, \vec{X}_2)$ is approximately no greater than that of $\widehat{\eta}(\vec{X}_1, \vec{X}_2)$.

In the case when $\vec{X}_1 = \vec{X}_2$, $\mathrm{Cov}\left(\widehat{\xi}(\vec{X}_1), \widehat{\gamma}(\vec{X}_2)\right) \neq 0$ and $\mathrm{Cov}\left(\widetilde{\xi}(\vec{X}_1), \widetilde{\gamma}(\vec{X}_2)\right) \neq 0$. Also, we have that

$$\mathrm{Cov}\left(\widetilde{\xi}(\vec{X}_1), \widetilde{\gamma}(\vec{X}_2)\right) \leq \mathrm{Cov}\left(\widehat{\xi}(\vec{X}_1), \widehat{\gamma}(\vec{X}_2)\right)$$

by Lemma 2.1.1 of [Bratley et al. 1987]. But since the variances of the new estimators of the numerator and denominator are smaller than those for the original estimators, we cannot compare the biases and variances of $\widehat{\eta}(\vec{X}_1, \vec{X}_2)$ and $\widetilde{\eta}(\vec{X}_1, \vec{X}_2)$. However, we examine this case empirically in Section 8 and find that there is a variance reduction and smaller mean squared error.

7. STORAGE AND COMPUTATION COSTS

We now discuss the implementation issues associated with constructing our new estimator $\widetilde{\alpha}(\vec{X})$ given in (17) for the case when α is defined in (13). First note that the first term in the second line of (17), excluding the factor $2/(h_{12} - 1)$, satisfies

$$\begin{aligned} &\sum_{k \in H(1;2)} Y_{12}(k) \left(\sum_{j \in H(1;2)} Y_{21}(j) - Y_{21}(\psi(k)) \right) \\ &= \sum_{k \in H(1;2)} Y_{12}(k) \sum_{j \in H(1;2)} Y_{21}(j) - \sum_{k \in H(1;2)} Y_{12}(k) Y_{21}(\psi(k)). \end{aligned}$$

Also, the last term in the last line of (17) satisfies

$$\sum_{\substack{l,m \in J(2;1) \\ l \neq m}} Y_{22}(l) Y_{22}(m) = \left(\sum_{k \in J(2;1)} Y_{22}(k) \right)^2 - \left(\sum_{k \in J(2;1)} Y_{22}(k)^2 \right).$$

Hence, to construct our estimator $\widetilde{\alpha}(\vec{X})$, we need to calculate the following quantities:

—the sum of the $Y(k)^2$ over the 1-cycles $k \in J(1;2)$;
—the sums of the $Y_{12}(k)$, $Y_{21}(k)$, $Y_{12}(k)^2$, and $Y_{21}(k)^2$ over the 1-cycles $k \in H(1;2)$;
—the sum of the $Y_{12}(k) Y_{21}(\psi(k))$ over the 1-cycles $k \in H(1;2)$;
—the sums of the $Y_{22}(k)$ and $Y_{22}(k)^2$ over the 2-cycles $k \in J(2;1)$.

To compute these quantities in a simulation, we do not have to store the entire sample path but rather only need to keep track of the various cumulative sums as the simulation progresses. Also, the amount of storage required is fixed and does not increase with the simulation run length. Therefore, compared to the standard estimator, the new estimator can be constructed with little additional computational effort and storage. (Pseudo-code for this estimator is given in the appendix.) A similar situation holds when estimating η using the estimator defined in Theorem 5.

We conclude this section with a rough comparison of the work required for the new estimator with that for the standard regenerative method when estimating α given in (13). Let W_s be the (random) amount of work to generate a particular sample path of m_1 1-cycles in a discrete-event simulation, where W_s includes the work for the random-variate generation, determining transitions, and appropriately updating data structures needed in the sample-path generation. This quantity is the same for our new method and the standard method.

We now study the work needed for the output analysis required by the standard regenerative method. After every transition in the simulation, we need to update the value of the current cycle-quantity $Y(g;k)$; see (2). Let θ_1 denote this (deterministic) amount of work, and if there are N total transitions in the sample path, then the total work for updating the $Y(g;k)$ during the entire simulation is $N\theta_1$. At the end of every 1-cycle, we have to square the current cycle quantity $Y(g;k)$ and add it to its accumulator; see (14). Let θ_2 denote this (deterministic) amount of work required at the end of each 1-cycle, and since there are m_1 1-cycles along the path, the total work for accumulating the sum of the $Y(g;k)^2$ is $m_1\theta_2$. Therefore, the cumulative work (including sample-path generation and output analysis) for the standard regenerative method is

$$W_r = W_s + N\theta_1 + m_1\theta_2.$$

Now we determine the amount of work needed for the output analysis of our new permutation method. By examining the pseudo-code in the appendix, we see that after every transition, a single accumulator is updated. Every time a stopping time from either sequence T_1 or T_2 occurs, we compute either a square or a product of two terms and update at most three accumulators, and the amount of work for this is essentially at most $3\theta_2$. Since the number of times this needs to be done is $m_1 + M_2$, the cumulative work for our permutation method is

$$W_p = W_s + N\theta_1 + (m_1 + 3M_2)\theta_2 = W_r + 3M_2\theta_2.$$

Therefore, the ratio of cumulative work of our permutation method relative to the standard regenerative method is

$$R_W = 1 + \frac{3M_2\theta_2}{W_s + N\theta_1 + m_1\theta_2}.$$

Typically in a regenerative simulation, the amount of time W_s required to generate the sample path is much greater than the time needed to perform the output analysis, and so R_W will usually be close to 1. Hence, the overhead of using our method in a simulation will most likely be very small, and this is what we observed in our experimental results in the next section.

8. EXPERIMENTAL RESULTS

Our example is based on the Ehrenfest urn model. The transition probabilities for this discrete-time Markov chain are given by $P_{0,1} = P_{s,s-1} = 1$, and

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

In our experiments we take $s = 8$. The stopping-time sequences T_1 and T_2 for our regenerative simulation correspond to hitting times to the states v and w, respectively, and so state v is the return state for the regenerative simulation. We ran several simulations of this system to estimate two different performance measures: σ^2, which is the time-average variance constant from Section 5.1, and η, which is the mean hitting time to a set F from Section 6. For each performance measure, we ran our experiments with several different choices for v, and for each v, we examined all possible choices for w. Choosing $w = v$ has no effect on the resulting estimator, so this corresponds the standard estimator. We ran 1,000 independent replications for each choice of v and w. Tables 1–3 and 5–6 present the results from estimating the two performance measures, giving the sample average and sample variance of our new estimator over the 1,000 replications. The average cycle lengths change with different choices of v; in order to make the results somewhat comparable across the tables, we changed the number of simulated cycles for each case so that the total expected number of simulated transitions remains the same. For Table 1, corresponding to $v = 1$, we simulated 1,000 cycles, and a greater number for the other tables. For example, the expected cycle length is 3.5 times as long for state 1 as for state 2, so in Table 2, we simulated 3,500 cycles. Since our new estimator reduces the variance but at the cost of extra computational effort, we also compare the efficiencies (inverse of the product of the variance and the time to generate the estimator) of our new estimator and the standard one, as suggested by Hammersley and Handscomb [1964] and Glynn and Whitt [1992].

8.1 Results from Estimating Variance

We first examine the results from estimating the time-average variance σ^2 with cost function $f(x) = x$. We performed 3 experiments, corresponding to return states $v = 1, 2$ and 4, and these results are given in Tables 1–3, respectively.

The transition probabilities are symmetric around state 4 (the mean of the binomial stationary distribution), so our first choice of return state $v = 1$ in Table 1 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 Table 3 we examine the same model, but now with our return state v chosen to be the stationary mean, 4. 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. (Unfortunately, there are no reliable rules for choosing *a priori* a good return state.) Comparing the three tables, we see that the minimum

Table 1. Estimating variance with $v = 1$.

w	Avg. of $\widetilde{\sigma}^2$	Sample Var.
0	13.92	1.28
1	13.92	1.28
2	13.97	0.46
3	13.98	0.25
4	13.99	0.18
5	13.98	0.21
6	13.99	0.35
7	13.99	0.65
8	13.96	1.11

Table 2. Estimating variance with $v = 2$.

w	Avg. of $\widetilde{\sigma}^2$	Sample Var.
0	13.96	0.48
1	13.96	0.48
2	13.96	0.48
3	13.97	0.25
4	14.00	0.19
5	13.99	0.18
6	13.99	0.24
7	13.97	0.35
8	13.97	0.46

variability does not change much across tables (0.17 for Table 3, and 0.18 for the other tables). This example suggests that it may be possible to compensate for a bad choice of v by an appropriate choice of w.

Finally, we illustrate the computational burden of our method. Table 4 shows the work required for the results obtained in Table 1. The first column gives, for each choice of w, the relative work (CPU time for generating the sample path and output analysis) required for our new estimator, that is, the CPU time with our new estimator divided by the CPU time for the standard regenerative method. (The row with $w = 1$ corresponds to the standard regenerative method, so all entries are 1; the other entries are normalized with respect to these.) The second column gives the relative variance; that is, the sample variance of our new estimator divided by the sample variance for the standard regenerative estimator. The last column gives

Table 3. Estimating variance with $v = 4$.

w	Avg. of $\widetilde{\sigma}^2$	Sample Var.
0	13.99	0.20
1	14.00	0.20
2	14.00	0.18
3	14.00	0.18
4	13.99	0.20
5	13.99	0.17
6	13.99	0.18
7	13.99	0.19
8	13.99	0.20

Table 4. Comparison of efficiency, $v = 1$.

w	Relative Work	Relative Var.	Relative Efficiency
0	1.001	1.000	0.999
1	1.000	1.000	1.000
2	1.021	0.370	2.646
3	1.036	0.201	4.808
4	1.028	0.125	7.752
5	1.020	0.153	6.410
6	1.012	0.297	3.333
7	1.010	0.476	2.079
8	1.003	0.859	1.161

the relative efficiency; that is, the inverse of the product of the relative work from column 1 and the relative variance from the second column.

Notice that in the cases where the variance reduction is small, the increase in work is also small. More work is needed when there is a larger variance reduction, but this is still no more than a few percent increase. Note that in the best case ($w = 4$), the efficiency was improved by nearly a factor of eight. Because little additional work is needed by our method, within each of the other tables, the run times are approximately the same for the different values of w. Therefore, one can roughly approximate the relative efficiencies for Tables 2 and 3 as the ratio of the sample variances for states v and w.

It should also be noted that the Ehrenfest model considered here is very simple compared with typical simulation models. The additional work required to compute our new estimators is independent of the model, and so if the work to generate the sample path is much larger than for the Ehrenfest model, the relative increase in work would be correspondingly small.

8.2 Results from Estimating Hitting Times to a Set

We now consider estimating η, which is the mean hitting time to a set of states F starting from a state v for our Ehrenfest model with reward function $g(x) \equiv 1$. We take $F = \{7\}$, which is hit infrequently. Tables 5 and 6 show our results from generating 1,000 independent replications for $v = 4$ and $v = 2$, respectively. In each replication, we generated a sample path which we used to compute the new estimators for both the numerator and denominator. Hence, using the terminology of Section 6, we let $\vec{X}_1 = \vec{X}_2$. For each path, we generated 1,000 cycles for $v = 2$ and 2,500 cycles for $v = 4$, so that the expected sample path length is the same in each case. We calculated (i.e., not using simulation) the theoretical values to be $\eta = 50.3714$ for $v = 4$, and $\eta = 54.7905$ for $v = 2$, so in addition to examining the variance of our new estimator, we can also study the mean squared error. Note that the theoretical value of η depends on the starting state v.

First observe that there is no change in the estimates of η and their variances for certain choices of w. This is due to the fact that permuting w-cycles in these cases has no effect on the estimators of either the numerator or denominator. For the other values of w, the (relative) variance reduction is significantly greater when $v = 2$ (Table 6) than when $v = 4$ (Table 5). In addition, although the absolute bias is the greatest for $w = 6$ for both choices of v, its magnitude is quite small,

Table 5. Estimating expected hitting time to state 7 with $v = 4$.

w	Avg. of $\widetilde{\eta}$	Sample Var.	MSE
0	49.83	14.66	14.96
1	49.83	14.66	14.96
2	49.83	14.66	14.96
3	49.83	14.66	14.96
4	49.83	14.66	14.96
5	49.76	13.18	13.55
6	49.71	11.00	11.44
7	49.83	14.66	14.96
8	49.83	14.66	14.96

Table 6. Estimating expected hitting time to state 7 with $v = 2$.

w	Avg. of $\widetilde{\eta}$	Sample Var.	MSE
0	54.32	21.01	21.24
1	54.32	21.01	21.24
2	54.32	21.01	21.24
3	54.25	17.22	17.50
4	54.17	15.34	15.73
5	54.10	13.55	14.03
6	54.02	11.79	12.38
7	54.32	21.01	21.24
8	54.32	21.01	21.24

and when examining the mean squared error, the variance reduction overwhelms the effect of the bias.

We now explain why the choice of w that results in the most variance reduction is $w = 6$. In the original sample path without permuting the w-cycles, the set F is hit a certain number of times. By permuting the w-cycles, we get a variance reduction if there are some v-cycles that hit w but not F and if within a particular v-cycle that hits F, there is more than one w-cycle that hits F and does not hit v. Permuting the w-cycles then can distribute the hits to F to more of the v-cycles. The amount of variance reduction in estimating γ is largely determined by the difference between the maximum and minimum number of v-cycles that hit F from permuting the w-cycles. Choosing $w = 7$ or $w = 8$ results in no variance reduction because we are working with a birth-death process and so the process always hits F no later than it hits w within a v-cycle, and so permuting w-cycles has no effect. Of the remaining choices for w, $w = 6$ maximizes the number of w-cycles that hit F, hence the largest variance reduction. Therefore, in general, we suggest that the state w should be chosen so that $w \notin F$ and it is as "close" as possible to the set F to maximize the number of w-cycles that hit F.

9. DIRECTIONS FOR FUTURE RESEARCH

We are currently investigating how to contruct confidence intervals for our permuted estimators; see [Calvin and Nakayama 1997a]. This entails proving central limit theorems for our estimators. Another area on which we are working is determining how to choose the two sequences of regenerative times T_1 and T_2 when there are more than two possibilities. For example, this arises when simulating a Markov chain, since successive hits to any fixed state form a regenerative sequence. We explored this to some degree experimentally in Section 8, but further study is needed.

ACKNOWLEDGMENTS

The authors would like to thank the Editor in Chief and the two anonymous referees for their helpful comments on the paper.

APPENDIX

A. PROOFS

A.1 Proof of Theorem 1

We need only prove that $\vec{X} \stackrel{\mathcal{D}}{=} \vec{X}'$; that is, the paths have the same distribution when 2-cycles are permuted. Recall M_2 is the number of times that stopping times in the sequence T_2 occur by time $T_1(m_1)$, and if $M_2 = n+1$, define $\vec{X}^0, \vec{X}^1, \ldots, \vec{X}^{n+1}$ as in (8)–(10). For the path $\vec{X}^k \in \mathcal{X}$, define $R(\vec{X}^k)$ to be the number of times a stopping time from sequence T_1 occurs in $\vec{X}^k$. Now define $L = R(\vec{X}^0) + R(\vec{X}^{n+1})$ to be the number of times that stopping times from the sequence T_1 occur outside of the 2-cycles (we do not need to bother with the case $M_2 = 0$ since then we do not change the path).

If f is a (measurable) function mapping sample paths to nonnegative real values, then

$$\begin{aligned} E\left[f(\vec{X})\right] &= \sum_{i \le m_1, n} E\left[f(\vec{X}); M_2 = n+1, L = m_1 - i\right] \\ &= \sum_{i \le m_1, n} E\left[f_{n,i}(\vec{X}^0, \vec{X}^1, \ldots, \vec{X}^{n+1}); M_2 = n+1, L = m_1 - i\right] \end{aligned}$$

for some (measurable) functions $f_{n,i}$, and so it suffices to show that

$$P\left(\vec{X}^0 \in A_0,\ \vec{X}^1 \in A_1,\ \vec{X}^2 \in A_2, \ldots, \vec{X}^n \in A_n,\ \vec{X}^{n+1} \in A_{n+1},\ M_2 = n+1,\ L = m_1 - i\right)$$

is invariant under 2-cycle permutations, where the A_i are (measurable) sets. Note that

$$\begin{aligned} &P\left(\vec{X}^0 \in A_0,\ \vec{X}^1 \in A_1,\ \vec{X}^2 \in A_2, \ldots, \vec{X}^n \in A_n,\ \vec{X}^{n+1} \in A_{n+1},\ M_2 = n+1,\ L = m_1 - i\right) \\ &= P\left(\vec{X}^0 \in A_0,\ \vec{X}^{n+1} \in A_{n+1},\ \vec{X}^1 \in A_1,\ \vec{X}^2 \in A_2, \ldots, \vec{X}^n \in A_n \mid M_2 = n+1,\ L = m_1 - i\right) \\ &\quad \times P\left(M_2 = n+1,\ L = m_1 - i\right). \end{aligned}$$

Given M_2 and L, the initial 1-2 segment (i.e., $\vec{X}^0$) and final 2-1 segment (i.e., $\vec{X}^{n+1}$) are conditionally independent of the 2-cycles, so the last probability can be written

$$\begin{aligned} &= P\left(\vec{X}^0 \in A_0, \vec{X}^{n+1} \in A_{n+1} \mid M_2 = n+1, L = m_1 - i\right) P\left(M_2 = n+1, L = m_1 - i\right) \\ &\quad \times P\left(\vec{X}^1 \in A_1, \vec{X}^2 \in A_2, \ldots, \vec{X}^n \in A_n \mid M_2 = n+1, L = m_1 - i\right). \end{aligned}$$

In examining the effect of permutations of the 2-cycles, we need consider only the last probability, which we rewrite

$$\begin{aligned}
&P\left(\vec{X}^1 \in A_1, \vec{X}^2 \in A_2, \dots, \vec{X}^n \in A_n \,\middle|\, M_2 = n+1, L = m_1 - i\right) \\
&= P\left(\vec{X}^1 \in A_1, \vec{X}^2 \in A_2, \dots, \vec{X}^n \in A_n \,\middle|\, \sum_{j=1}^{n} R(\vec{X}^j) = i, R(\vec{X}^{n+1}) > 0, L = m_1 - i\right) \\
&= \frac{P\left(\vec{X}^1 \in A_1, \vec{X}^2 \in A_2, \dots, \vec{X}^n \in A_n, \sum_{j=1}^{n} R(\vec{X}^j) = i, R(\vec{X}^{n+1}) > 0, L = m_1 - i\right)}{P\left(\sum_{j=1}^{n} R(\vec{X}^j) = i, R(\vec{X}^{n+1}) > 0, L = m_1 - i\right)}.
\end{aligned}$$

Since we are interested in the effect of permutations, we only look at the numerator of the last expression:

$$\begin{aligned}
&= P\left(\vec{X}^1 \in A_1, \vec{X}^2 \in A_2, \dots, \vec{X}^n \in A_n, \sum_{j=1}^{n} R(\vec{X}^j) = i, R(\vec{X}^{n+1}) > 0, L = m_1 - i\right) \\
&= P\left(\vec{X}^1 \in A_1, \vec{X}^2 \in A_2, \dots, \vec{X}^n \in A_n, \sum_{j=1}^{n} R(\vec{X}^j) = i\right) P\left(R(\vec{X}^{n+1}) > 0, L = m_1 - i\right)
\end{aligned}$$

and we are finally left with the task of showing that for any permutation σ,

$$\begin{aligned}
&P\left(\vec{X}^1 \in A_1, \vec{X}^2 \in A_2, \dots, \vec{X}^n \in A_n, \sum_{j=1}^{n} R(\vec{X}^j) = i\right) \\
&\quad = P\left(\vec{X}^{\sigma(1)} \in A_1, \vec{X}^{\sigma(2)} \in A_2, \dots, \vec{X}^{\sigma(n)} \in A_n, \sum_{j=1}^{n} R(\vec{X}^j) = i\right).
\end{aligned}$$

But this follows from the fact that

$$\left(\vec{X}^1, \vec{X}^2, \dots, \vec{X}^n\right) \stackrel{\mathcal{D}}{=} \left(\vec{X}^{\sigma(1)}, \vec{X}^{\sigma(2)}, \dots, \vec{X}^{\sigma(n)}\right).$$

Therefore, $\vec{X}' \stackrel{\mathcal{D}}{=} \vec{X}$, and the theorem is proved. Notice that we only used the conditional exchangeability of the cycles, and not the full independence.

A.2 Proof of Theorem 2

Recall that in Section 5 we defined $H(1;2)$, $J(1;2)$, h_{12}, $Y_{12}(k)$, $Y_{21}(k)$, $Y_{22}(l)$, and B_k for the original sample path $\vec{X}$. Using a permuted path $\vec{X}'$ instead of the original path $\vec{X}$ in (16), we get

$$\begin{aligned}
\widehat{\alpha}(\vec{X}') &= \frac{1}{m_1}\left(\sum_{k \in J'(1;2)} Y'(k)^2 + \sum_{k \in H'(1;2)} \left[Y'_{12}(k)^2 + Y'_{21}(k)^2\right]\right) \\
&+ \frac{1}{m_1}\sum_{k \in H'(1;2)}\left(2Y'_{12}(k)Y'_{21}(k) + 2[Y'_{12}(k) + Y'_{21}(k)]\sum_{l \in B'_k} Y'_{22}(l) + \left(\sum_{l \in B'_k} Y'_{22}(l)\right)^2\right), \quad (28)
\end{aligned}$$

where the Y', H', J', B' variables and sets are the same as the Y, H, J, B variables and sets, respectively, with $\vec{X}'$ replacing $\vec{X}$ in (16). Recall that E_* is the conditional expectation operator corresponding to a random (uniform) permutation of 2-cycles (as was done when constructing the path $\vec{X}'$ from $\vec{X}$) given the original sample path $\vec{X}$. Also, recall that we define our new estimator to be

$$\widetilde{\alpha}(\vec{X}) = E_*[\widehat{\alpha}(\vec{X}')],$$

which we will now show is equivalent to (17).

First note that by our construction of the path $\vec{X}'$ from $\vec{X}$, the first term in (28) does not change when replacing $\vec{X}$ with $\vec{X}'$. Hence,

$$\begin{aligned} &E_*\left[\frac{1}{m_1}\left(\sum_{k\in J'(1;2)} Y'(k)^2 + \sum_{k\in H'(1;2)} \left[Y_{12}'(k)^2 + Y_{21}'(k)^2\right]\right)\right] \\ &\quad = \frac{1}{m_1}\left(\sum_{k\in J(1;2)} Y(k)^2 + \sum_{k\in H(1;2)} \left[Y_{12}(k)^2 + Y_{21}(k)^2\right]\right). \end{aligned} \tag{29}$$

Now we compute the conditional expectation of the second term of (28). We can assume that $h_{12} \geq 3$. Let $\rho(k)$ be the index of the Y_{21} segment that follows the $Y_{12}(k)$ segment after a permutation of the 2-cycles. Note that $\rho(k) \neq \psi(k)$ since $Y_{12}(k)$ and $Y_{21}(\psi(k))$ are always in the same 2-cycle, no matter how the 2-cycles are permuted. Any of the other $h_{12} - 1$ indices from $H(1;2)$ are equally likely, however, so that

$$\begin{aligned} E_*\left[\sum_{k\in H'(1;2)} 2\,Y_{12}'(k)\,Y_{21}'(k)\right] &= E_*\left[\sum_{k\in H(1;2)} 2\,Y_{12}(k)\,Y_{21}(\rho(k))\right] \\ &= 2\sum_{k\in H(1;2)} Y_{12}(k)\,E_*\left[Y_{21}(\rho(k))\right] \\ &= 2\sum_{k\in H(1;2)} Y_{12}(k)\frac{1}{h_{12}-1}\left(\sum_{j\in H(1;2)} Y_{21}(j) - Y_{21}(\psi(k))\right). \end{aligned} \tag{30}$$

For the second summand in the second term of (28), we have

$$\begin{aligned} &E_*\left[\sum_{k\in H(1;2)} 2\left[Y_{12}(k) + Y_{21}(\rho(k))\right] \sum_{l\in B'_{\rho(k)}} Y_{22}(l)\right] \\ &= 2\,E_*\left[\sum_{k\in H(1;2)} Y_{12}(k) \sum_{l\in B'_{\rho(k)}} Y_{22}(l)\right] + 2\,E_*\left[\sum_{k\in H(1;2)} Y_{21}(\rho(k)) \sum_{l\in B'_{\rho(k)}} Y_{22}(l)\right] \\ &= 2\sum_{k\in H(1;2)} Y_{12}(k)\,E_*\left[\sum_{l\in B'_{\rho(k)}} Y_{22}(l)\right] + 2\sum_{k\in H(1;2)} Y_{21}(k)\,E_*\left[\sum_{l\in B'_k} Y_{22}(l)\right] \end{aligned}$$

$$= \; 2 \sum_{k\in H(1;2)} Y_{12}(k)\, E_* \left[\sum_{l\in B'_k} Y_{22}(l)\right] + 2 \sum_{k\in H(1;2)} Y_{21}(k)\, E_* \left[\sum_{l\in B'_k} Y_{22}(l)\right]$$
$$= \; 2\,[\bar{Y}_{12} + \bar{Y}_{21}] \sum_{l\in J(2;1)} Y_{22}(l). \tag{31}$$

For the third summand in the second term of (28), we have

$$E_* \left[\sum_{k\in H(1;2)} \left(\sum_{l\in B_k} Y_{22}(l)\right)^2\right] \;=\; E_* \left[\sum_{k\in H(1;2)} \left(\sum_{l\in J(2;1)} Y_{22}(l)\, 1_{\{l\in B_k\}}\right)^2\right]$$
$$= \; E_* \left[\sum_{k\in H(1;2)} \left(\sum_{l\in J(2;1)} Y_{22}(l)^2\, 1_{\{l\in B_k\}} + \sum_{\substack{l,m\in J(2;1)\\ l\neq m}} Y_{22}(l) Y_{22}(m)\, 1_{\{l\in B_k, m\in B_k\}}\right)\right].$$

Now note that

$$E_* \left[\sum_{k\in H(1;2)} \sum_{l\in J(2;1)} Y_{22}(l)^2\, 1_{\{l\in B_k\}}\right] = \sum_{l\in J(2;1)} Y_{22}(l)^2 \sum_{k\in H(1;2)} P_*\left(l \in B_k\right) = \sum_{l\in J(2;1)} Y_{22}(l)^2.$$

Also,

$$E_* \left[\sum_{k\in H(1;2)} \sum_{\substack{l,m\in J(2;1)\\ l\neq m}} Y_{22}(l) Y_{22}(m)\, 1_{\{l\in B_k, m\in B_k\}}\right]$$
$$= \sum_{\substack{l,m\in J(2;1)\\ l\neq m}} Y_{22}(l) Y_{22}(m) \sum_{k\in H(1;2)} P_*\left(l \in B_k, m \in B_k\right)$$
$$= \sum_{\substack{l,m\in J(2;1)\\ l\neq m}} Y_{22}(l) Y_{22}(m) \sum_{k\in H(1;2)} P_*\left(1 \in B_k, 2 \in B_k\right).$$

Now use the fact that

$$\sum_{k\in H(1;2)} P_*\left(1 \in B_k, 2 \in B_k\right) = \frac{2}{1 + h_{12}}, \tag{32}$$

which follows from

LEMMA 6. *Suppose that p white balls, numbered 1 to p, are placed along with q black balls into $p+q$ boxes arranged in a line, with each box getting exactly one ball. Apply a uniformly chosen random permutation to the balls. Then the probability that ball 1 and ball 2 are not separated by a black ball is $2/(2+q)$.*

To apply this lemma to (32), we let $q = h_{12} - 1$ be the number of 2-cycles that include one of the T_1 stopping times, and p be the number of remaining 2-cycles.

PROOF. Let D be the number of boxes in between the boxes containing ball 1 and ball 2, and let L_i be the box number containing ball i ($i = 1, 2$). Then

$$\begin{aligned}
P(D = k) &= \sum_{l=1}^{p+q} P(D = k \mid L_1 = l) \frac{1}{p+q} \\
&= \sum_{l=1}^{p+q-k-1} P(L_2 = l + k + 1 \mid L_1 = l) \frac{1}{p+q} + \sum_{l=k+2}^{p+q} P(L_2 = l - k - 1 \mid L_1 = l) \frac{1}{p+q} \\
&= \sum_{l=1}^{p+q-k-1} \frac{1}{p+q-1} \frac{1}{p+q} + \sum_{l=k+2}^{p+q} \frac{1}{p+q-1} \frac{1}{p+q} \\
&= \frac{2(p+q-k-1)}{(p+q)(p+q-1)}.
\end{aligned}$$

Given that $D = k$, the probability that balls 1 and 2 are not separated by at least 1 black ball is the probability that all q black balls are chosen from the $p+q-2-k$ boxes that are not between ball 1 and ball 2, which is

$$\frac{\binom{p+q-k-2}{q}}{\binom{p+q-2}{q}}.$$

Thus the desired probability is

$$\begin{aligned}
&\sum_{k=0}^{p-2} \frac{\binom{p+q-k-2}{q}}{\binom{p+q-2}{q}} \frac{2(p+q-k-1)}{(p+q)(p+q-1)} \\
&= \frac{2}{(p+q)(p+q-1)} \frac{1}{\binom{p+q-2}{q}} \sum_{k=0}^{p-2} \binom{p+q-k-2}{q} (p+q-k-1) \\
&= \frac{2}{(p+q)(p+q-1)} \frac{1}{\binom{p+q-2}{q}} (q+1) \sum_{k=0}^{p-2} \binom{q+1+k}{q+1} \\
&= \frac{2}{(p+q)(p+q-1)} \frac{1}{\binom{p+q-2}{q}} (q+1) \binom{q+p}{q+2} \\
&= \frac{2}{q+2}.
\end{aligned}$$

□

Hence,

$$\begin{aligned}
&E_* \left[\sum_{k \in H(1;2)} \left(\sum_{l \in B_k} Y_{22}(l) \right)^2 \right] \\
&= \sum_{l \in J(2;1)} Y_{22}(l)^2 + \frac{2}{1 + h_{12}} \sum_{\substack{l, m \in J(2;1) \\ l \neq m}} Y_{22}(l) Y_{22}(m). \qquad (33)
\end{aligned}$$

Finally, putting together (29), (30), (31), and (33), we get that $\widetilde{\alpha}(\vec{X})$ is as in (17). The unbiasedness and variance reduction follow directly from Theorem 1.

A.3 Proof of Theorem 5

We first need the following result.

LEMMA 7. *Suppose that q black balls, r white balls, and p red balls are placed at random in $q+p+r$ boxes arranged in a line (one ball per box). The probability that there are no white balls in an interval formed by two particular black balls (or the start or end of the boxes) is $q/(q+r)$.*

PROOF. Count boxes from the left until a non-red ball is encountered. The desired probability is the probability that the first non-red ball is a black ball. Since of the non-red balls q are black and r are white, this probability is $q/(q+r)$. □

Now we prove Theorem 5. First recall our definitions of $I_{12}(k)$, $I_{22}(l)$, $I_{21}(k)$, k_{21}, k_{12}, k_c, r, p, d_0, d_1, $D_{12}(k)$, $D_{22}(l)$, and $D_{21}(k)$ given in Section 6. Also recall (23). The first term in (23) is independent of the permutations, and so we now consider the second term.

Note that $I_{12}(i)$, $I_{22}(j)$, and $I_{21}(k)$ are independent for any i, j, k. Then

$$
\begin{aligned}
& E_* \left[\sum_{k \in H(1;2)} \max \left(I_{12}(k),\ \max_{l \in B'_{\rho(k)}} I_{22}(l),\ I_{21}(\rho(k)) \right) \right] \\
&= \sum_{k \in H(1;2)} E_* \left[\max \left(I_{12}(k),\ \max_{l \in B'_{\rho(k)}} I_{22}(l),\ I_{21}(\rho(k)) \right) \right] \\
&= \sum_{k \in H(1;2)} \left(1 - P_* \left\{ \max \left(I_{12}(k),\ \max_{l \in B'_{\rho(k)}} I_{22}(l),\ I_{21}(\rho(k)) \right) = 0 \right\} \right) \\
&= \sum_{k \in H(1;2)} \left(1 - P_* \{ I_{12}(k) = 0 \}\ P_* \{ I_{21}(\rho(k)) = 0 \}\ P_* \left\{ \max_{l \in B'_{\rho(k)}} I_{22}(l) = 0 \right\} \right) \\
&= \sum_{k \in H(1;2)} \left(1 - (1 - I_{12}(k)) \frac{k_{21} + I_{21}(\psi(k)) - 1}{h_{12} - 1} P_* \left\{ I_{22}(l) = 0,\ l \in B'_{\rho(k)} \right\} \right) \\
&= h_{12} - P_* \{ I_{22}(l) = 0,\ l \in B'_1 \} \sum_{k \in H(1;2)} (1 - I_{12}(k)) \frac{k_{21} - (1 - I_{21}(\psi(k)))}{h_{12} - 1} \\
&= h_{12} - \frac{P_* \{ I_{22}(l) = 0,\ l \in B'_1 \}}{h_{12} - 1} (k_{12} k_{21} - k_c) \\
&= h_{12} - \frac{k_{12} k_{21} - k_c}{h_{12} - 1 + r},
\end{aligned}
$$

where we have used Lemma 7.

We now examine the estimation of ξ. Recall (24) and (25). The first term in (24) is not affected by permuting the 2-cycles, but the second term is, and so we now examine the second term. First note that

$$E_* \left[\sum_{k \in H'(1;2)} D'_{12}(k) \right] = \sum_{k \in H(1;2)} D_{12}(k) \tag{34}$$

since the $D_{12}(k)$ are unaffected by permutations of 2-cycles. Also,

$$\begin{aligned}
& E_* \left[\sum_{k \in H(1;2)} D_{21}(\rho(k))(1 - I_{12}(k)) \prod_{l \in B'_{\rho(k)}} (1 - I_{22}(l)) \right] \\
&= \sum_{k \in H(1;2)} (1 - I_{12}(k))\, E_* \left[D_{21}(\rho(k)) \prod_{l \in B'_{\rho(k)}} (1 - I_{22}(l)) \right] \\
&= \sum_{k \in H(1;2)} (1 - I_{12}(k))\, E_* \left[D_{21}(\rho(k)) \right] P_* \left\{ \prod_{l \in B'_{\rho(k)}} (1 - I_{22}(l)) = 1 \right\} \\
&= \sum_{k \in H(1;2)} (1 - I_{12}(k)) \frac{1}{h_{12} - 1} \left(\sum_{j \in H(1;2)} D_{21}(j) - D_{21}(\psi(k)) \right) P_* \left\{ I_{22}(l) = 0,\ l \in B'_{\rho(k)} \right\} \\
&= P_* \left\{ I_{22}(l) = 0,\ l \in B'_1 \right\} \sum_{k \in H(1;2)} (1 - I_{12}(k)) \frac{1}{h_{12} - 1} \left(\sum_{j \in H(1;2)} D_{21}(j) - D_{21}(\psi(k)) \right) \\
&= \frac{1}{h_{12} - 1 + r} \left(k_{12} \sum_{j \in H(1;2)} D_{21}(j) - \sum_{k \in H(1;2)} (1 - I_{12}(k)) D_{21}(\psi(k)) \right).
\end{aligned} \tag{35}$$

Finally, after the permutation of the 2-cycles, define $R(k)$ to be the number of 2-cycles in $J(2;1)$ that immediately follow the path segment corresponding to $D_{12}(k)$, and let $\delta_1(k), \delta_2(k), \ldots, \delta_{R(k)}(k)$ be the indices of those 2-cycles in $J(2;1)$ that immediately follow the path segment corresponding to $D_{12}(k)$ in the order they appear. Then

$$\begin{aligned}
& E_* \left[\sum_{k \in H(1;2)} \sum_{l=1}^{R(k)} D_{22}(\delta_l(k))(1 - I_{12}(k)) \prod_{j<l} (1 - I_{22}(\delta_j(k))) \right] \\
&= \sum_{k \in H(1;2)} (1 - I_{12}(k))\, E_* \left[\sum_{l=1}^{R(k)} D_{22}(\delta_l(k)) \prod_{j<l} (1 - I_{22}(\delta_j(k))) \right] \\
&= \sum_{k \in H(1;2)} (1 - I_{12}(k))\, E_* \left[\sum_{l=1}^{R'(k)} D_{22}(\delta_l(k)) \right],
\end{aligned} \tag{36}$$

where $R'(k) = \min(R(k), \inf\{j \geq 1 : I_{22}(\delta_j(k)) = 1\})$.

LEMMA 8.

$$E_* \left[\sum_{l=1}^{R'(k)} D_{22}(\delta_l(k)) \right] = d_0 \frac{p}{r + h_{12}} + d_1 \frac{r}{r + h_{12} - 1}.$$

PROOF. Suppose m balls are placed in $n \geq m$ boxes in a line. Let Z be the number of empty boxes on the left end. Then

$$P(Z \geq k) = \frac{\binom{n-k}{m}}{\binom{n}{m}}$$

for $0 \leq k \leq n - m$. Therefore,

$$P(Z = k) = \frac{\binom{n-k}{m} - \binom{n-k-1}{m}}{\binom{n}{m}} = \frac{\binom{n-k-1}{m-1}}{\binom{n}{m}},$$

and

$$\begin{aligned}
E(Z) &= \sum_{k=0}^{n-m} k \frac{\binom{n-k-1}{m-1}}{\binom{n}{m}} \\
&= \frac{1}{\binom{n}{m}} \sum_{k=0}^{n-m} k \binom{n-k-1}{m-1} \\
&\quad \text{(substitute } j = n - m - k\text{)} \\
&= \frac{1}{\binom{n}{m}} \sum_{j=0}^{n-m} (n - m - j) \binom{m-1+j}{m-1} \\
&= (n - m) - \frac{1}{\binom{n}{m}} \sum_{j=1}^{n-m} j \binom{m-1+j}{m-1} \\
&= (n - m) - \frac{1}{\binom{n}{m}} m \sum_{j=1}^{n-m} \binom{m-1+j}{m} \\
&\quad \text{(substitute } i = j - 1\text{)} \\
&= (n - m) - \frac{1}{\binom{n}{m}} m \sum_{i=0}^{n-m-1} \binom{m+i}{m} \\
&\quad \left(\text{use the identity } \sum_{i=0}^{p} \binom{q+i}{q} = \binom{q+p+1}{q+1} \right) \\
&= (n - m) - \frac{1}{\binom{n}{m}} m \binom{n}{m+1} \\
&= \frac{n-m}{m+1}.
\end{aligned}$$

To get the mean number of D_{22}'s that do not hit F, we use the above formula with $n = p + r + (h_{12} - 1)$ and $m = r + (h_{12} - 1)$. □

Finally, putting together Lemma 8 and (34), (35), and (36), we get our new estimator for ξ. The unbiasedness and variance reduction of our two estimators follow directly from Theorem 1.

B. PSEUDO-CODE

B.1 Estimator in Theorem 2

Below is the pseudo-code for the estimator in Theorem 2. (The pseudo-code for estimator in Theorem 5 is similar.) Note that it is specifically for a discrete-event simulation of a continuous-time process. Discrete-event processes $\{X_n : n = 0, 1, 2, \ldots\}$ can be handled by letting the inter-event time Δ always be 1. The estimator is denoted by `alpha`.

```
m ← number of regenerative 1-cycles to simulate;
k ← 0;              // counter for number of regenerative 1-cycles
t ← 0;              // simulation time
h12 ← 0;            // cardinality of H(1;2)
sumy12 ← 0;         // sum of the Y_12(k) over H(1;2)
sumy21 ← 0;         // sum of the Y_21(k) over H(1;2)
sumy22 ← 0;         // sum of the Y_22(k) over J(2;1)
sumysq ← 0;         // sum of the Y(k)^2 over J(1;2)
sumy12sq ← 0;       // sum of the Y_12(k)^2 over H(1;2)
sumy21sq ← 0;       // sum of the Y_21(k)^2 over H(1;2)
sumy22sq ← 0;       // sum of the Y_22(k)^2 over J(2;1)
sumy12y21 ← 0;      // sum of the Y_12(k)Y_21(ψ(k)) over H(1;2)
accum ← 0;          // accumulator for Y over the current segment
laststoptime ← 1;   // type (i.e., 1 or 2) of last stopping time to occur

generate initial state x;
do while (k < m)
   generate inter-event time Δ;
   t ← t + Δ;
   accum ← accum + g(x)*Δ;
   generate next state x;

   if (T_1 occurs at current time t) then
      if (laststoptime == 1) then
         sumysq ← sumysq + accum*accum;
      else
         sumy21 ← sumy21 + accum;
         sumy21sq ← sumy21sq + accum*accum;
         lasty21 ← accum;
      endif
      laststoptime = 1;
      k ← k + 1;
      accum ← 0;
   endif

   if (T_2 occurs at current time t) then
```

```
      if (laststoptime == 1) then
         sumy12 ← sumy12 + accum;
         sumy12sq ← sumy12sq + accum*accum;
         h12 ← h12 + 1;
         if (h12 == 1) then
            firsty12 = term;
         else
            sumy12y21 ← sumy12y21 + term*lasty21;
         endif
      else
         sumy22 ← sumy22 + accum;
         sumy22sq ← sumy22sq + accum*accum;
      endif
      laststoptime = 2;
      accum ← 0;
   endif
enddo

if (h12 < 3) then
   alpha ← sumysq/m;
else
   sumy12y21 ← sumy12y21 + firsty12*lasty21;
   alpha ← (1/m) * ( sumysq + sumy12sq + sumy21sq
      + (2/(h12-1))*(sumy12*sumy21 - sumy12y21) + sumy22sq
      + (2/h12)*(sumy12 + sumy21)*sumy22
      + (2/(1+h12))*(sumy22*sumy22 - sumy22sq) );
endif
end
```

REFERENCES

Bratley, P., Fox, B., and Schrage, L. 1987. *A Guide to Simulation, Second Edition*. Springer-Verlag, New York.

Calvin, J. and Nakayama, M. 1997a. Central limit theorems for permuted regenerative estimators. Working draft.

Calvin, J. and Nakayama, M. 1997b. A new variance-reduction technique for regenerative simulations of Markov chains. In *Proceedings of the 1997 Winter Simulation Conference* (December 1997). IEEE Press. to appear.

Conover, W. 1980. *Practical Nonparametric Statistics, Second Edition*. John Wiley and Sons, New York.

Crane, M. and Iglehart, D. 1975. Simulating stable stochastic systems, III: Regenerative processes and discrete-event simulations. *Operations Research 23*, 1 (January-February), 33–45.

Efron, B. 1979. Bootstrap methods: Another look at the jackknife. *Annals of Statistics 7*, 1 (February), 1–26.

Fox, B. and Glynn, P. 1990. Discrete-time conversion for simulating finite-horizon Markov processes. *SIAM Journal of Applied Mathematics 50*, 1457–1473.

Glynn, P. and Iglehart, D. 1987. A joint central limit theorem for the sample mean and regenerative variance estimator. *Annals of Operations Research 8*, 41–55.

Glynn, P. and Whitt, W. 1992. The asymptotic efficiency of simulation estimators. *Operations Research 40*, 3 (May-June), 505–520.

Goyal, A., Shahabuddin, P., Heidelberger, P., Nicola, V., and Glynn, P. 1992. A unified framework for simulating Markovian models of highly dependable systems. *IEEE Transactions on Computers C-41*, 1 (January), 36–51.

Hammersley, J. and Handscomb, D. 1964. *Monte Carlo Methods*. Methune, London.

Hordijk, A., Iglehart, D., and Schassberger, R. 1976. Discrete-time methods of simulating continuous-time Markov chains. *Advances in Applied Probability 8*, 772–788.

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

Kingman, J. 1972. *Regenerative Phenomena*. John Wiley and Sons, New York.

Mood, A., Graybill, F., and Boes, D. 1974. *Introduction to the Theory of Statistics, Third Edition*. McGraw-Hill, Inc., New York.

Sen, P. 1977. Some invariance principles relating to jackknifing and their role in sequential analysis. *Annals of Statistics 5*, 316–329.

Serfling, R. 1980. *Approximation Theorems of Mathematical Statistics*. John Wiley and Sons, New York.

Shahabuddin, P. 1994. Importance sampling for highly reliable Markovian systems. *Management Science 40*, 3, 333–352.

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