
Derandomizing Variance Estimators

Author(s): Shane G. Henderson and Peter W. Glynn

Source: *Operations Research*, Vol. 47, No. 6 (Nov. - Dec., 1999), pp. 907-916

Published by: INFORMS

Stable URL: <http://www.jstor.org/stable/222980>

Accessed: 17/07/2010 01:42

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=informs>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



INFORMS is collaborating with JSTOR to digitize, preserve and extend access to *Operations Research*.

DERANDOMIZING VARIANCE ESTIMATORS

SHANE G. HENDERSON

University of Auckland, Auckland, New Zealand, sg.henderson@auckland.ac.nz

PETER W. GLYNN

Department of Engineering—Economic Systems and Operations Research, Stanford University, Stanford, California, 94305, glynn@leland.stanford.edu

(Received December 1995; revision received December 1997; accepted August 1998)

One may consider a discrete-event simulation as a Markov chain evolving on a suitably rich state space. One way that regenerative cycles may be constructed for general state-space Markov chains is to generate auxiliary coin-flip random variables at each transition, with a regeneration occurring if the coin-flip results in a success. The regenerative cycles are therefore *randomized* with respect to the sequence of states visited by the Markov chain. The point estimator for a steady-state performance measure does not depend on the cycle structure of the chain, but the variance estimator (that defines the width of a confidence interval for the performance measure) does. This implies that the variance estimator is randomized with respect to the visited states. We show how to “derandomize” the variance estimator through the use of conditioning. A new variance estimator is obtained that is consistent and has lower variance than the standard estimator.

INTRODUCTION

Operations researchers are often interested in computing steady-state performance measures of a stochastic system. For example, one might wish to compute the long-run mean number of jobs queued at a work station in a job shop, or the long-run probability of buffer overflow in a telecommunications switch. The chief difficulty in steady-state simulation typically is not in obtaining a point estimator for the quantity of interest, but instead is in obtaining an estimate of its accuracy. One may gauge the accuracy of a point estimator through estimates of the *time average variance constant* (TAVC) of the stochastic system (see (1) below). In this paper, we derive a new estimator of the TAVC that may be used when the stochastic process of interest has a special form of regenerative structure. We show both analytically and through numerical examples that the new estimator has better statistical properties than a more standard regenerative TAVC estimator.

We give our results in the case where the underlying stochastic process is described as a Markov chain $X = (X_n; n \geq 0)$ evolving on a general state-space S . This framework encompasses virtually any discrete-event simulation. Indeed, it has long been recognized that generalized semi-Markov processes (GSMPs) are very general models of discrete-event stochastic systems (Ho 1991). But GSMPs may be described mathematically (Glynn 1983, 1989) as discrete-time general state-space Markov chains (GSSMCs), by observing the state of the process when events occur. Of course, sufficient information must be augmented to the state-space to ensure that the resulting discrete-time process is Markov.

EXAMPLE 1. The number of customers in a single-server queue is a discrete-event system with events triggered by

the arrival of a customer, or the completion of a service of a customer. To represent this process as a GSSMC, one would let Q_n be the number of customers in the system immediately after the n th customer arrival or departure, and $C_n(a), C_n(d)$ be the time until the next customer arrival and departure, respectively. The process $X = (X_n; n \geq 0)$, where $X_n = (Q_n, C_n(a), C_n(d))$, is a Markov chain on state-space $S \subseteq \{0, 1, 2, \dots\} \times \mathbb{R}^+ \times \mathbb{R}^+$.

Another field where simulation of GSSMCs has become important is that of Markov chain Monte-Carlo simulation (MCMC). MCMC is used in global optimization algorithms (Boender and Romeijn 1995) and in statistical modeling situations where Bayesian methods are applied (Gilks et al. 1996). This application is explored in more depth in §4, where we describe how the GSSMC arises and give a numerical example where the techniques described in this paper prove useful.

Let us return to our more general setup. Let $f: S \rightarrow \mathbb{R}$ be a cost function defined on the state-space S of X . Under suitable conditions, the strong law of large numbers for Markov chains (Meyn and Tweedie 1993, p. 411) yields

$$\alpha_n \triangleq \frac{1}{n} \sum_{k=0}^{n-1} f(X_k) \rightarrow \alpha$$

almost surely (a.s.) as $n \rightarrow \infty$, where α is the steady-state mean of $(f(X_n); n \geq 0)$. Therefore, α_n is usually taken as a point estimator of the performance measure α . Again, under certain conditions (Meyn and Tweedie 1993, p. 411) a central limit theorem (CLT)

$$\sqrt{n}(\alpha_n - \alpha) \Rightarrow \sigma N(0, 1) \quad (1)$$

holds where \Rightarrow denotes convergence in distribution, and $N(0, 1)$ is a standard normal random variable (r.v.).

Subject classifications: Simulation: statistical analysis. Probability: regenerative processes; Markov processes.
Area of review: SIMULATION.

The constant σ^2 is the TAVC, and it is easily seen to provide some measure of the accuracy of the estimator α_n because an approximate $100(1 - p)\%$ confidence interval for α can be easily derived from (1) as $\alpha_n \pm z_{1-p/2} \sigma / \sqrt{n}$, where $P(N(0, 1) > z_{1-p/2}) = p/2$. The TAVC is also important if we wish to obtain a distributional approximation for $C_n = \sum_{k=0}^{n-1} f(X_k)$, the cumulative cost up to time n . The CLT (1) suggests that

$$C_n \stackrel{\mathcal{D}}{\approx} n\alpha + \sigma\sqrt{n}N(0, 1),$$

where $\stackrel{\mathcal{D}}{\approx}$ should be understood to mean “has approximately the same distribution.” Therefore, with estimators of α and σ^2 we may obtain approximations for quantities such as tail probabilities for C_n . Pich (1992) gives another context where reliable estimates of σ^2 are desirable. In that work, the TAVC appears as an input to a diffusion approximation for queueing systems. Clearly, reliable estimators of σ^2 are desirable.

One method for estimating σ^2 is regenerative simulation. The popularity of this method stems from the ease with which a consistent estimator of σ^2 may be constructed and the rate of convergence of the estimator to σ^2 . Indeed, Henderson and Glynn (1997b) generalize a result by Glynn and Iglehart (1987) to show that regenerative TAVC estimators have the fastest rate of convergence of all known TAVC estimators. For a simulation run of length n , the error in the TAVC estimator is of the order $n^{-1/2}$, which is the same as the order of the error in the point estimator. This compares favorably with, for instance, spectral estimators that typically converge at rate $n^{-\beta}$ for some $\beta < 1/2$ (Grenander and Rosenblatt 1984, p. 129). This fast rate of convergence is a strong argument for the use of the regenerative method whenever possible.

In the regenerative method, a sample path of X is decomposed into cycles by identifying regeneration times when the chain probabilistically “starts over.” The cycle structure is then used to identify an estimator for σ^2 . For example, if X is a discrete-time Markov chain on a finite state-space, then the chain regenerates every time it returns to a fixed state (we will call this the *returns* method for determining regenerations). In general state-space, however, it is possible that no single point is recurrent, so this method of identifying regeneration time fails. Fortunately, Athreya and Ney (1978) and Nummelin (1978) provide another method for determining regenerations.

The “splitting method” (as we shall call it) generates auxiliary coin-flip r.v.s at each transition, and a regeneration occurs if the coin-flip results in a success. So, in contrast to regeneration times determined by hitting times of a single state, the regeneration times are now *randomized* with respect to the sample path of X .

The splitting method was originally introduced as a theoretical tool in the analysis of GSSMCs. The idea was adapted to a simulation setting in Glynn (1982), and a discussion of the method can be found in Glynn and L’Ecuyer (1993). In a recent paper, Andradóttir et al.

(1995) apply the splitting method to discrete state space Markov chains, achieving variance reduction in estimating σ^2 over the more traditional returns method. It should be noted that the variance reduction technique described in this paper is easily applied to the setting in Andradóttir et al. (1995) and is *in addition* to the gains obtained there.

The standard point estimator α_n clearly does not depend on regeneration times, while the variance estimator \tilde{V}_n does (see (8)). Because the regeneration times are randomized with respect to a sample path of X , the variance estimator is also randomized (but the point estimator is not). This suggests that after obtaining a sample path (X_0, X_1, \dots, X_n) of X , we could independently resample the regeneration times k times to obtain conditionally independent variance estimators $\hat{s}_1^2, \hat{s}_2^2, \dots, \hat{s}_k^2$, say. We could then average the results to obtain a single estimator s^2 . It is evident that s^2 is, in fact, an estimator of $E(\tilde{V}_n | X_0, \dots, X_n)$, which has a lower variance (and the same mean) than \tilde{V}_n because it is a conditional expectation (see Theorem 2). So we might now ask whether we can find an explicit formula for this expression. The answer is yes, and we do so in this paper.

In §1, we review pertinent results from Henderson and Glynn (1997b) in which the structural relation between GSMPs that model discrete-event stochastic systems and their associated GSSMCs is investigated. In particular, we introduce the notion of Harris recurrence (which generalizes that of recurrence for discrete state space Markov chains) and describe the structure of GSSMCs that arise from GSMP models. Such chains typically possess an m -minorization (see §1 for a definition), where m is strictly greater than 1, so it is important to consider such chains in the subsequent analysis.

The regenerative method is reviewed in §2. In §3 we introduce the derandomized estimator, show that it is weakly consistent, and prove that it achieves variance reduction over the standard regenerative TAVC estimator. Finally, in §4 we present two examples. The first is a discrete state-space Markov chain that demonstrates the application of derandomization to the class of Markov chains perhaps most familiar to the reader. The second is a real MCMC application where a useful variance reduction is achieved. This second example establishes both the feasibility and the effectiveness of the derandomized TAVC estimator. The less informative proofs appear in §5.

1. DISCRETE-EVENT DYNAMIC SYSTEMS AND MARKOV CHAINS

In this section we briefly review certain results from Henderson and Glynn (1997b). Our goal is to explain how a discrete-event stochastic simulation may be viewed as the simulation of a related GSSMC and, further, to give some idea of the structure of the resulting GSSMC. See Whitt (1980) and Glynn (1989) for more precise discussions.

An enormous class of discrete-event stochastic systems may be formulated as GSMPs. A GSMP consists of a countable set of states S representing the configuration of a

system plus a finite set of events E . Associated with each state is an active set of events $E(s)$ that compete to cause a state change. Each active event has a clock associated with it, indicating the time remaining until the event “occurs.” When the clock reaches zero, the event triggers a state change. Event clocks may be discarded at time of a state change if the event is no longer active in the new state. For each newly activated event, a clock reading is sampled from a probability distribution.

The GSSMC is constructed by observing the GSMP only at the instants at which a state transition occurs and providing a rich enough state-space that the resulting discrete-time process is Markov. For $n \geq 0$, let S_n and C_n be the state and vector of clock readings immediately after the n th transition (where we assume that the 0th transition occurs at time 0). Define $X = (X_n: n \geq 0)$, where $X_n = (S_n, C_n)$. Then X is a Markov chain on a general state-space because the future construction of the GSMP requires only the current state and clock readings.

Just as in discrete state space Markov chains, X must satisfy some notion of recurrence in order to exhibit stable long-term behavior. The appropriate notion in general state-space is that of Harris recurrence. In discrete state-space, a single state may be hit infinitely often, and this forms the basis for the formal definition of recurrence. In general state-spaces, it is possible that no point is hit infinitely often, even though the process seems recurrent from an intuitive standpoint.

EXAMPLE 2. Consider the chain $X = (X_n: n \geq 0)$, where $X_n = U_n - [U_n]$, and $U = (U_n: n \geq 0)$ is a random walk on the real line. The value X_n is simply the fractional part of U_n , and thus it evolves on state-space $\Omega = [0, 1)$. If the increments of U have a continuous distribution, then all nonempty intervals of Ω are eventually visited from every starting point. Intuitively, then X should be recurrent, even though no point in Ω will be hit infinitely often.

Observe that for this example the distribution of X_{n+1} should not vary greatly if $X_n = x$ is known to lie in some “small” set, so the distribution of X_{n+1} should contain some “common component” φ , say, the magnitude of which may vary with x . Although no single point is recurrent, the “common component” is. This is the essence of the following definition of Harris recurrence.

DEFINITION 1. Let $X = \{X_n: n \geq 0\}$ be a Markov chain on a complete separable metric space Σ . We say that X is Harris recurrent if there exists an $m \geq 1$, an $\varepsilon > 0$, a probability distribution φ , and a nonnegative function λ such that:

1. $P^m(x, \cdot) \triangleq P(X_m \in \cdot | X_0 = x) \geq \lambda(x)\varphi(\cdot) \forall x \in \Sigma$; and
2. $P(\lambda(X_n) \geq \varepsilon \text{ infinitely often} | X_0 = x) = 1 \forall x \in \Sigma$.

Harris chains automatically possess a unique (up to a multiplicative constant) nontrivial σ -finite invariant measure π . If π has finite total mass, then we say that X is positive Harris

recurrent, and we may take π to have total mass 1, so that it is a probability.

EXAMPLE 3. If $X = (X_n: n \geq 0)$ is an irreducible, recurrent, discrete-time Markov chain on a countable state-space Σ , then it is Harris recurrent. To see this, note that every state x is hit infinitely often. Let x^* be an arbitrary state, and set $\lambda(x^*) = 1$, and $\lambda(x) = 0$ for $x \neq x^*$. Next, define $m = 1$, $\varphi(\cdot) = P(x^*, \cdot)$ and let $\varepsilon \in (0, 1)$. It is then easy to see that with these definitions, X is Harris recurrent. If, in addition, X is positive recurrent, then X possesses a unique invariant distribution π (Chung 1967), and we may take $\pi(\Sigma) = 1$. Therefore, we may conclude that X is positive Harris recurrent. We will return to this example shortly.

DEFINITION 2. We say that (λ, φ) is an m -minorization of X if conditions 1 and 2 hold.

For conditions that ensure Harris recurrence of the GSSMC associated with a GSMP, see Glynn (1989). In Henderson and Glynn (1997b), it is shown that the GSSMC associated with a GSMP will typically have an m -minorization where $m > 1$, but will not have a 1-minorization. This result has implications for the application of the regenerative method, as we shall see in the next section.

To see why discrete-event systems do not typically have a 1-minorization, consider a system with no event cancellation, i.e., a system in which all clocks remain active until they trigger an event. Suppose that at least k clocks are active in every state. Then a minimum of k state transitions must occur before all the current clock readings are replaced or discarded. If all the clock-setting distributions have densities (with respect to Lebesgue measure), then the distribution of the state (S_n, C_n) will be concentrated on a set of Lebesgue measure 0 for at least k transitions. This situation cannot arise in a system with an m -minorization where $m < k$.

2. REGENERATION AND HARRIS CHAINS

In this section we show how to define regeneration times for positive Harris recurrent Markov chains and review the relevant regenerative process theory. Let $X = (X_n: n \geq 0)$ be a Harris recurrent Markov chain with an m -minorization.

2.1. Determining Regeneration Times

We only review the idea here. For more details the reader is referred to Glynn and L'Ecuyer (1993). We first assume that $m = 1$, but we will consider the case $m > 1$ shortly. Let us write

$$P(x, \cdot) = \lambda(x)\varphi(\cdot) + (1 - \lambda(x))Q(x, \cdot), \tag{2}$$

where $Q(x, \cdot)$ is given by

$$\frac{P(x, \cdot) - \lambda(x)\varphi(\cdot)}{1 - \lambda(x)},$$

when $\lambda(x) < 1$ and (arbitrarily) a point mass at x when $\lambda(x) = 1$. This decomposition lies at the heart of the splitting variable approach to regenerative simulation.

The decomposition (2) suggests generating a transition of the Markov chain from $X_n = x$ in the following way. First, generate a Bernoulli r.v. Z_n with success probability $\lambda(x)$. If $Z_n = 1$, generate X_{n+1} from φ , otherwise generate X_{n+1} from Q . The important idea here is that if $Z_n = 1$, X_{n+1} is distributed according to φ *independently* of X_n . We can simulate X using this two-step process and when $Z_n = 1$, $n + 1$ is a regeneration time.

Simulating r.v.s with distribution φ or Q may be difficult, so a related approach is usually taken. Suppose that $X_0 = x$. First, generate X_1 from the starting point x , and then generate Z_0 with success probability $w(x, X_1)$, where $w(x, y)$ is the conditional probability that a regeneration occurred, given that x and then y were observed. Intuitively (from acceptance/rejection ideas), the value $w(x, y)$ should satisfy

$$w(x, y) = \frac{\lambda(x)\varphi(dy)}{P(x, dy)}.$$

That is, $w(x, \cdot)$ is a density of $\lambda(x)\varphi(\cdot)$ with respect to $P(x, \cdot)$. As before, if $Z_0 = 1$, then the chain regenerates at time $t = 1$ (Glynn and L'Ecuyer 1993). The simulation can proceed by generating X_{n+1} from X_n and then generating Z_n as a Bernoulli($w(X_n, X_{n+1})$) r.v.

EXAMPLE 3 (*continued*). The decomposition (2) is exactly

$$P(x, \cdot) = I(x = x^*)P(x^*, \cdot) + I(x \neq x^*)Q(x, \cdot),$$

where $I(\cdot)$ is the indicator function that is 1 if its argument is true and 0 otherwise. The two-step process referred to above is as follows. If $X_n = x^*$, then $Z_n = 1$, X_{n+1} is generated according to $P(x^*, \cdot)$, and $n + 1$ is a regeneration time. If $X_n = x \neq x^*$, then $Z_n = 0$, X_{n+1} is generated according to $P(x, \cdot)$, and $n + 1$ is not a regeneration time. This is precisely the returns method for determining regeneration times. The second method of determining regenerations also coincides with the returns method. In this case, $w(x^*, y) = 1$ for all y , and $w(x, y) = 0$ if $x \neq x^*$.

Now consider the case where $m > 1$. We may write

$$P^m(x, \cdot) = \lambda(x)\varphi(\cdot) + (1 - \lambda(x))Q(x, \cdot), \tag{3}$$

where $Q(x, \cdot)$ is given by

$$\frac{P^m(x, \cdot) - \lambda(x)\varphi(\cdot)}{1 - \lambda(x)}, \tag{4}$$

when $\lambda(x) < 1$ and (arbitrarily) a point mass at x when $\lambda(x) = 1$. This decomposition suggests that starting from $X_0 = x$, we should generate Z_0 as a Bernoulli($\lambda(x)$) r.v. If $Z_0 = 1$, then X_m should be generated according to the distribution φ , otherwise from Q . If $Z_0 = 1$, then X_m will be distributed according to φ *independently* of X_0 , so that in fact, (X_m, X_{m+1}, \dots) and X_0 are independent, and m is (in some sense) a regeneration time. Notice however, that after generating X_0 and X_m , we must generate X_1, \dots, X_{m-1} conditional

on those two values, which may be difficult. However, a similar trick to the case $m = 1$ may also be performed here. First generate X_0, \dots, X_m . Then compute a Bernoulli r.v. Z_0 with success probability $w(X_0, X_m)$, where $w(x, \cdot)$ is a density of $\lambda(x)\varphi(\cdot)$ with respect to $P^m(x, \cdot)$. If $Z_0 = 1$, then X_m is distributed according to φ independently of X_0 , and a regeneration is recorded. No matter which method is used to determine whether m is a regeneration time, X_m may be correlated with X_1, \dots, X_{m-1} . This leads to 1-dependent cycles (adjacent cycles may be dependent, but nonadjacent cycles are independent), as we will see shortly.

When $m > 1$, how are regeneration times constructed? One method is to generate Z_0, Z_m, Z_{2m}, \dots , and declare $X_{(k+1)m}$ to be a regeneration time if $Z_{km} = 1$. A second method is to test for regenerations more selectively by generating the Bernoulli r.v.s only when X visits some set A . Of course, one must wait until at least m transitions have gone by before generating the next Bernoulli r.v.

In either case, coin-flips are calculated at a possibly random subsequence of transitions $\{\Gamma_n: n \geq 0\}$ separated by at least m time units, with success probabilities $w(X_{\Gamma_n - m}, X_{\Gamma_n})$. If the coin-flip is successful, then X_{Γ_n} is distributed according to φ *independently* of the distribution of $X_{\Gamma_n - m}$, and we record a regeneration.

As noted earlier, the resulting cycles are no longer independent but are 1-dependent (nonadjacent cycles are independent). To see this, note that if X_n is a regeneration time, then X_n is independent of X_0, X_1, \dots, X_{n-m} , but no conclusions can be drawn about dependence between X_n and $X_{n-m+1}, \dots, X_{n-1}$. Nonadjacent cycles are separated by at least m time units and are therefore independent.

As discussed in the previous section, if X is a GSSMC associated with a GSMP model, then it typically will not have a 1-minorization but will have an m -minorization where $m > 1$. Therefore, we must be prepared to deal with the more complicated $m > 1$ case in our subsequent analysis.

REMARK. In order to use the splitting approach to determine regenerations, we need to know the densities $w(x, \cdot)$ for all $x \in S$. As noted above, w is a density of the minorization with respect to the m -step transition probabilities. Because calculation of the m -step transition probabilities can be extremely difficult, the densities $w(x, \cdot)$ may also be difficult to compute for many problems. This is the essence of the difficulty in implementing the regenerative method for chains with m -minorizations when $m > 1$.

2.2. The Regenerative Method

Let $f: \Sigma \rightarrow \mathbb{R}$ be a real-valued cost function on Σ . If $\pi|f| \triangleq \int_{\Sigma} |f(x)|\pi(dx) < \infty$, then the strong law of large numbers for positive Harris recurrent Markov chains (Meyn and Tweedie 1993, p. 411) implies that

$$\alpha_n \triangleq n^{-1} \sum_{k=0}^{n-1} f(X_k) \rightarrow \pi f \triangleq \alpha, \tag{5}$$

as $n \rightarrow \infty$. Therefore, α may be interpreted as the long-run average cost of running the system X and can be estimated consistently by α_n .

Let $T_0 = 0$, let T_1, T_2, \dots be regeneration times for X and let $\tau_k = T_k - T_{k-1}$ (for $k = 1, 2, \dots$) be the associated cycle lengths. Let $g: \Sigma \rightarrow \mathbb{R}$, and define $Y_k(g)$ as the cost accumulated over the k th cycle using cost function g , i.e.,

$$Y_k(g) \triangleq \sum_{j=T_{k-1}}^{T_k-1} g(X_j).$$

At the regeneration times $T_k, X_{T_k} \sim \varphi$ where \sim should be read "has distribution."

Define $E_\varphi(\cdot)$ as $\int_\Sigma E(\cdot | X_0 = x) \varphi(dx)$, the conditional expectation given that $X_0 \sim \varphi$ and $P_\varphi(A) = E_\varphi I_A$, the expected value of the indicator function of A . If $E_\varphi(Y_1(|f|)^2 + \tau_1^2) < \infty$, then (Meyn and Tweedie 1993, p. 419)

$$\sqrt{n}(\alpha_n - \alpha) \Rightarrow \sigma N(0, 1),$$

where \Rightarrow denotes convergence in distribution.

REMARK. It is possible to weaken the hypothesis that $E_\varphi(Y_1(|f|)^2 + \tau_1^2) < \infty$ to $E_\varphi(\tau_1 + Y_1(f) - \alpha\tau_1)^2 < \infty$, although we will not require this greater level of generality. A proof of this result for the 1-minorization case may be found in Glynn and Iglehart (1993). The extension to the m -minorization case ($m > 1$) is similar.

Let $f_c(x) \triangleq f(x) - \alpha$ and $f_n(x) \triangleq f(x) - \alpha_n$ be the cost function centered by the steady-state mean and its estimate, respectively. Meyn and Tweedie (1993) show that if the regenerations are based on the 1-minorization, then the constant σ^2 may be written as

$$\sigma^2 = \frac{E_\varphi(Y_1(f_c)^2)}{E_\varphi \tau_1}; \tag{6}$$

whereas if the regenerations are based on an m -minorization with $m > 1$, so that the cycles are 1-dependent, Glynn (1982) shows that

$$\sigma^2 = \frac{E_\varphi(Y_1(f_c)^2) + 2E_\varphi(Y_1(f_c)Y_2(f_c))}{E_\varphi \tau_1}. \tag{7}$$

Because of the dependence structure of the cycles,

$$EY_1(f_c)Y_2(f_c) = ER_1(f_c)Y_2(f_c),$$

where, for a given function g , $R_j(g) \triangleq \sum_{k=T_j-m+1}^{T_j-1} g(X_k)$; i.e., $R_j(g)$ is the cost accumulated over the last portion of the cycle. Natural estimators for σ^2 for the two cases are given by

$$W_n \triangleq \frac{1}{n} \sum_{j=1}^{\ell(n)} Y_j(f_n)^2,$$

and

$$W_n^{(m)} \triangleq \frac{1}{n} \sum_{j=1}^{\ell(n)} Y_j(f_n)^2 + \frac{2}{n} \sum_{j=1}^{\ell(n)-1} R_j(f_n)Y_{j+1}(f_n),$$

where $\ell(n) \triangleq \sup\{k \geq 0 : T_k \leq n\}$ is the number of regenerations that occur by time n .

3. DERANDOMIZATION

The goal of this section is to describe the derandomized estimator, show that it consistently estimates σ^2 , and compare its statistical properties with the standard regenerative estimator. Let us first assume that the minorization constant $m = 1$. This is the case, for instance, in the application of splitting to countable state space Markov chains as in Andradóttir et al. (1995). We will consider the case $m > 1$ at the conclusion of this section.

By the standard estimator we mean \tilde{V}_n , where

$$\begin{aligned} \frac{1}{n} \sum_{j=1}^{\ell(n)} Y_j(f_n)^2 &= \frac{1}{n} \sum_{j=0}^{T_{\ell(n)}-1} f_n(X_j)^2 \\ &+ \frac{2}{n} \sum_{j=0}^{T_{\ell(n)}-1} f_n(X_j) \sum_{k=j+1}^{T_{\ell(n)+1}-1} f_n(X_k) \end{aligned} \tag{8}$$

$$\begin{aligned} &\approx \frac{1}{n} \sum_{j=0}^{n-1} f_n(X_j)^2 \\ &+ \frac{2}{n} \sum_{j=0}^{n-1} f_n(X_j) \sum_{k=j+1}^{(T_{\ell(n)+1}-1) \wedge n} f_n(X_k) \end{aligned} \tag{9}$$

$$\triangleq \tilde{V}_n,$$

and $a \wedge b = \min(a, b)$. Although it is more common to refer to (8), as the standard estimator, it is easy to show that if $E_\varphi \tau_1^2 < \infty$ and $E_\varphi Y_1(|f|)^2 < \infty$, then the difference between (8) and \tilde{V}_n converges to 0 a.s. as $n \rightarrow \infty$. Under these same conditions, (8) converges a.s. to σ^2 as $n \rightarrow \infty$ (Wolff 1989, p. 123). It follows that $\tilde{V}_n \rightarrow \sigma^2$ a.s. as $n \rightarrow \infty$, i.e., \tilde{V}_n is a strongly consistent estimator for σ^2 .

As in §2, let $w(x, \cdot)$ be a density of $\lambda(x)\varphi(\cdot)$ with respect to $P(x, \cdot)$, and let $\bar{w}(x, y) = 1 - w(x, y)$. The derandomized estimator V_n is given by (assuming $X_0 \sim \varphi$)

$$\begin{aligned} V_n \triangleq E_\varphi(\tilde{V}_n | X) &= E_\varphi \left\{ \frac{1}{n} \sum_{k=0}^{n-1} f_n(X_k)^2 + \frac{2}{n} \sum_{j=0}^{n-1} f_n(X_j) \right. \\ &\quad \cdot \left. \sum_{k=j+1}^{(T_{\ell(n)+1}-1) \wedge n} f_n(X_k) | X \right\} \\ &= \frac{1}{n} \sum_{k=0}^{n-1} f_n(X_k)^2 + \frac{2}{n} E_\varphi \left\{ \sum_{j=0}^{n-1} f_n(X_j) \right. \\ &\quad \cdot \left. \sum_{k=j+1}^n f_n(X_k) I(T_{\ell(n)+1} - 1 \geq k) | X \right\} \end{aligned}$$

$$\begin{aligned}
 &= \frac{1}{n} \sum_{k=0}^{n-1} f_n(X_k)^2 + \frac{2}{n} \sum_{j=0}^{n-1} f_n(X_j) \\
 &\quad \cdot \sum_{k=j+1}^n f_n(X_k) P_\varphi(T_{\ell(j)+1} - 1 \geq k | X) \\
 &= \frac{1}{n} \sum_{k=0}^{n-1} f_n(X_k)^2 + \frac{2}{n} \sum_{j=0}^{n-1} f_n(X_j) \\
 &\quad \cdot \sum_{k=j+1}^n f_n(X_k) \prod_{i=j}^{k-1} \bar{w}(X_i, X_{i+1}). \tag{10}
 \end{aligned}$$

The next result shows that V_n is a consistent estimator for σ^2 . The proof is given in the Appendix.

THEOREM 1. *Let $X = \{X_n : n \geq 0\}$ be a positive Harris recurrent Markov chain on a complete, separable, metric state-space S . Suppose that X has a 1-minorization as described in §1. Let X_0 have distribution φ . If $E_\varphi(\tau_1^2 + Y_1(|f|)^2 \tau_1^2) < \infty$, then $P_\varphi(|V_n - \sigma^2| > \varepsilon) \rightarrow 0$ as $n \rightarrow \infty$ for each $\varepsilon > 0$, so V_n is consistent.*

REMARK. In this theorem we show that V_n is weakly consistent (V_n converges in P_φ probability to σ^2). In fact, if $\bar{w}(X_i, X_{i+1}) = 0$ infinitely often, then V_n enjoys regenerative structure, and thus under appropriate conditions is strongly consistent for σ^2 (converges a.s. to σ^2) and satisfies a central limit theorem (Henderson and Glynn 1997a). The observation that under certain conditions V_n is strongly consistent for σ^2 may be important in the context of sequential stopping methods. It is one of the sufficient conditions given by Glynn and Whitt (1992) that ensure that sequential stopping rules are asymptotically valid. Using a device called *rerandomization*, Henderson and Glynn (1997a) show how to ensure that $\bar{w}(X_i, X_{i+1}) = 0$ infinitely often at the expense of some statistical efficiency.

Clearly the assumption that X_0 has distribution φ is unrestrictive. One way to achieve this is to simulate X until a regeneration occurs and then discard the initial segment. To compare the statistical efficiency of the derandomized estimator with the standard estimator we present the following result.

THEOREM 2. *Let \tilde{V}_n be defined as in (9) and V_n be given by (10). Assume the conditions of Theorem 1.*

- Under P_φ , the L_1 error of V_n is less than that of \tilde{V}_n , i.e.,

$$E_\varphi |V_n - \sigma^2| \leq E_\varphi |\tilde{V}_n - \sigma^2|.$$

- Additionally, $\text{var}_\varphi V_n \leq \text{var}_\varphi \tilde{V}_n$. The inequality fails to be strict only when $\lambda(X_n) \in \{0, 1\} \forall n$ P_φ a.s.

- Under P_φ , then mean squared error (MSE) of V_n is less than that of \tilde{V}_n , i.e.,

$$E_\varphi (V_n - \sigma^2)^2 \leq E_\varphi (\tilde{V}_n - \sigma^2)^2.$$

PROOF. The first result is a simple consequence of the conditional version of Jensen’s inequality (see, e.g.,

Billingsley 1986, p. 470). For the second result note that

$$\begin{aligned}
 \text{var } \tilde{V}_n &= \text{var } E(\tilde{V}_n | X) + E(\text{var}(\tilde{V}_n | X)) \\
 &= \text{var } V_n + E(\text{var}(\tilde{V}_n | X)).
 \end{aligned}$$

But $E(\text{var}(\tilde{V}_n | X)) \geq 0$, with equality if and only if there is no randomization in the cycles, i.e., $\lambda(X_n)$ is either 0 or 1 for every n . Hence the second result follows. To prove the final result, note that for any r.v. W , $\text{MSE}(W) = \text{var}(W) + \text{bias}(W)^2$. Since $E V_n = E E(\tilde{V}_n | X) = E \tilde{V}_n$, $\text{bias}(V_n) = \text{bias}(\tilde{V}_n)$, and the final result follows from the second. \square

We immediately obtain the following corollary.

COROLLARY 1. *Suppose the conditions of Theorem 1 hold. Then*

$$\limsup_{n \rightarrow \infty} n \text{MSE}_\varphi V_n \leq \limsup_{n \rightarrow \infty} n \text{MSE}_\varphi \tilde{V}_n.$$

REMARK If, in addition to the conditions of Theorem 1, $E Y_1(f_c)^4 < \infty$, then the CLT proved by Glynn and Iglehart (1987) holds for \tilde{V}_n ; i.e., there is a constant η^2 such that $\sqrt{n}(\tilde{V}_n - \sigma^2) \Rightarrow \eta N(0, 1)$. Therefore, $n(\tilde{V}_n - \sigma^2)^2 \Rightarrow \eta^2 N(0, 1)^2$. If the sequence of r.v.s $(n(\tilde{V}_n - \sigma^2)^2 : n \geq 1)$ is uniformly integrable, then $\eta^2 = \limsup_{n \rightarrow \infty} n \text{MSE}_\varphi \tilde{V}_n$, so Corollary 1 proves that asymptotic variance reduction is achieved.

Given the fact that a large number of GSSMCs arising from discrete-event stochastic systems will have minorization constants $m > 1$, it is important to consider this case. The extension is, in fact, straightforward. When $m > 1$, we define the standard estimator of σ^2 to be

$$\begin{aligned}
 \tilde{V}_n &= \frac{1}{n} \sum_{j=0}^{n-1} f_n(X_j)^2 + \frac{2}{n} \sum_{j=0}^{n-1} f_n(X_j) \sum_{k=j+1}^{(T_{\ell(j)+1}-1) \wedge n} f_n(X_k) \tag{11} \\
 &\quad + \frac{2}{n} \sum_{j=0}^{n-1} f_n(X_j) I(T_{\ell(j)+1} - m < j) \sum_{k=T_{\ell(j)+1}}^{(T_{\ell(j)+2}-1) \wedge n} f_n(X_k). \tag{12}
 \end{aligned}$$

The first term in (7) is estimated by (11) and the second by (12). The derandomized estimator is given by

$$\begin{aligned}
 V_n &= E_\varphi(\tilde{V}_n | X) \\
 &= \frac{1}{n} \sum_{k=0}^{n-1} f_n(X_k)^2 + \frac{2}{n} \sum_{j=0}^{n-1} f_n(X_j) \sum_{k=j+1}^n f_n(X_k) \\
 &\quad \cdot \prod_{i:j < \Gamma_i \leq k} \bar{w}(X_{\Gamma_i-m}, X_{\Gamma_i}) \\
 &\quad + \frac{2}{n} \sum_i \sum_{j=\Gamma_i-m+1}^{\Gamma_i-1} f_n(X_j) w(X_{\Gamma_i-m}, X_{\Gamma_i}) \\
 &\quad \cdot \left(\sum_{k=\Gamma_i}^n f_n(X_k) \prod_{h:\Gamma_i < \Gamma_h \leq k} \bar{w}(X_{\Gamma_h-m}, X_{\Gamma_h}) \right).
 \end{aligned}$$

Table 1. Estimates of the variance of variance estimators.

Estimator	ε							
	0.5	0.4999	0.499	0.49	0.4	0.2	0.1	0.05
Returns	0.69	0.70	0.70	0.74	1.6	13	130	1200
Standard	7.5E-6	6.3E-5	5.9E-4	4.8E-3	0.13	2.8	54	580
Derandomized	7.5E-6	7.5E-6	8.9E-6	1.3E-4	1.9E-2	0.96	13	150
Ratio	1.0	8.4	66.3	36.9	6.8	2.9	4.2	3.9

Theorem 1 goes through as for the case $m = 1$ with a slight modification to the proof (discussed in the Appendix). Theorem 2 goes through without change. Hence the derandomized estimator is both consistent and statistically more efficient than the standard estimator for $m \geq 1$.

4. NUMERICAL EXAMPLES

EXAMPLE 4. Let $X = (X_n; n \geq 0)$ be a Markov chain on state-space $S = \{0, 1\}$, with transition matrix

$$P_{xy} = \begin{cases} 1 - \varepsilon, & x = y, \\ \varepsilon, & x \neq y, \end{cases}$$

for some $\varepsilon \leq 1/2$. We wish to estimate the steady-state probability that X is in state 1, so that $f(x) = x$. Let φ be the uniform distribution on S . Set $\lambda(x) = 2\varepsilon$ for all x so that (λ, φ) is a 1-minorization. The regeneration density is given by

$$w(x, y) = \begin{cases} \frac{\varepsilon}{1-\varepsilon}, & x = y, \\ 1, & x \neq y. \end{cases}$$

We simulated 100 replications of 10,000 transitions, each time calculating certain variance estimators. By replicating the simulations 100 times we were able to obtain estimates of the TAVCs of the variance estimators. We chose several values of ε in $(0, 0.5]$. The results appear in Table 1.

The first row gives an estimate of the TAVC of the returns variance estimator using 0 as the return state. The next two rows give the corresponding results for the standard estimator (9) and the derandomized estimator (10). The final row reports the ratio of the variances of the standard estimator and the derandomized estimator. Observe the following.

- The standard and derandomized estimators achieve significant variance reduction over the returns estimator for all values of ε . To see why, note that the 1-minorization used is equally effective over both states of the Markov chain. However, the returns method has mostly short cycles (returning directly to state 0) and occasionally very long ones (the chain goes to state 1), inducing high variance in the variance estimator.

- As $\varepsilon \rightarrow 0$, the variances increase rapidly, but at approximately the same rate. As $\varepsilon \rightarrow 0.5$, the standard and derandomized estimators' variances converge to the same value. (This is to be expected because the standard and derandomized variance estimators coincide at $\varepsilon = 0.5$.)

- For small ε , the additional variance reduction achieved by the derandomized estimator over the standard estimator is moderate.

- The variance reduction achieved by the derandomized estimator over the standard estimator is maximized for some value of ε near 0.5. Indeed, the variance of the standard estimator is approximately 70 times the variance of the derandomized estimator when $\varepsilon = 0.499$.

The results of this experiment are somewhat representative of other finite state-space Markov chains we tried. The standard and derandomized estimators achieved quite large variance reductions over the returns estimator when the minorization used was effective over the entire range of the state space. Although the derandomized estimator occasionally achieved large reductions over the standard estimator, the reduction was usually moderate (less than a factor of 3, but still useful).

EXAMPLE 5. This example is intended to demonstrate that derandomization can be useful when applied to a difficult simulation problem. It further demonstrates that often no single regenerative state can be identified. This problem is an instance of the use of MCMC methods, which have recently received a great deal of attention in the literature (see Gilks et al. 1996, from which much of the material presented in this example is taken). They provide an effective method for sampling from a distribution π that is known up to a multiplicative normalizing constant. The basic idea is to construct a Markov chain that has π as its stationary distribution and then use the simulated sample path of the Markov chain as a dependent sample from π .

One particular area in which MCMC has seen extensive application is in the area of Bayesian inference. In Bayesian inference, both the observed data (D , say) and the parameters of interest (θ , say) are treated as random quantities. We then need to set up a joint distribution $P(D, \theta)$. This is done by specifying two parts: a *prior* distribution $P(\theta)$ on the parameters, and a *likelihood* $P(D|\theta)$ on the data conditional on the parameters. The joint distribution is then given by $P(D, \theta) = P(D|\theta)P(\theta)$. After observing the data D , we may then use Bayes theorem to determine the *posterior* distribution of θ conditional on D as

$$P(\theta|D) = \frac{P(\theta)P(D|\theta)}{\int P(\vartheta)P(D|\vartheta) d\vartheta}$$

Although all the elements of this expression for the posterior are known, the integral in the denominator will be known only in special cases. Hence, it is common that the posterior distribution is known only up to a normalizing constant.

Another field in which MCMC methods have seen extensive use is global optimization (Boender and Romeijn 1995).

Table 2. Observed pump failures.

<i>i</i>	<i>s_i</i>	<i>t_i</i>	<i>i</i>	<i>s_i</i>	<i>t_i</i>
1	5	94.320	6	19	31.440
2	1	15.720	7	1	1.048
3	5	62.880	8	1	1.048
4	14	125.760	9	4	2.096
5	3	5.240	10	22	10.480

For example, the method of pure adaptive search (Patel et al. 1988) for minimizing a function *f* proceeds iteratively. At each stage a current solution candidate is retained; call it *x_k*, say. The method then selects *x_{k+1}* to be uniformly distributed on the level set {*x* : *f*(*x*) ≤ *f*(*x_k*)}. (Obviously, the level set must have finite volume for this method to work.) Patel et al. (1988) show that under certain conditions, the number of iterations of this method required to solve the problem to a given precision increases only linearly with the dimension of the feasible set. The difficulty arises in being able to generate points uniformly distributed on the level set. This may be done by simulating a positive Harris recurrent Markov chain with limiting distribution that is uniform on the level set, i.e., through the use of MCMC methods.

The example that we will consider is examined in Mykland et al. (1994), and the data on which the example is based were originally analyzed in Gaver and O’Muircheartaigh (1987). Failures in 10 pumps at a nuclear power station are assumed to occur according to independent Poisson processes. Each pump is assumed to have its own failure rate—*λ*₁, . . . , *λ*₁₀, say. The pumps were observed for periods *t_i*, and the number of failures observed in that period *s_i* were recorded. The data appear in Table 2. Conditional on a parameter *β*, the individual failure rates are assumed to have independent Gamma *G*(*α*, *β*) distributions with density proportional to *x*^{*α*-1} *e*^{-*βx*}. The parameter *β* has a *G*(*γ*, *δ*) distribution, with *γ* = 0.01 and *δ* = 1. For *α*, Mykland et al. (1994) followed previous efforts in using the method of moments estimator 1.802, and so shall we.

For the purposes of our numerical experiment, we will estimate the posterior expected failure rate of pump 10.

The parameter *θ* in this problem is the vector (*λ*₁, . . . , *λ*₁₀, *β*), and we wish to construct a Markov chain that has the posterior distribution of *θ* as its stationary distribution. The posterior distribution of *θ* can be characterized by the fact that given *β*, the *λ_i*s are independent *G*(*α* + *s_i*, *β* + *t_i*) r.v.s, and given *λ*₁, . . . , *λ*₁₀, *β* has a *G*(*γ* + 10*α*, Σ*λ_i* + *δ*) distribution. This characterization allows one to use the Gibbs sampler to construct an appropriate Markov chain (see Gilks et al. 1996 for details on the Gibbs sampler). Each iteration of the Markov chain consists of sampling a new set of *λ_i*s from their conditional distribution given *β*, and then sampling *β* from its conditional distribution given the *λ_i*s. The stationary distribution of this Markov chain is the required posterior distribution.

Mykland et al. (1994) define a regeneration density *w* based on three parameters—*Λ̃*, *d*₁, and *d*₂—and we used their suggestions for these parameters, namely 6.7, 1.591, and 3.109. The density is given by

$$w(x, y) = \begin{cases} \exp[(\tilde{\Lambda} - \Lambda(x)) \cdot (d(\Lambda(x)) - \beta(y))] & \text{if } d_1 \leq \beta(y) \leq d_2, \\ 0 & \text{otherwise.} \end{cases}$$

In this expression *x* and *y* are each vectors of the form (*λ*₁ . . . , *λ*₁₀, *β*), *Λ*(*x*) represents the sum of the *λ_i*s at *x*, and *d*(*z*) = *d*₁ if *z* < *Λ̃* and *d*₂ otherwise.

We performed 500 runs of the Gibbs sampler, each of length 1,000 transitions. The estimated TAVC of the standard estimator was 2.9, while that of the derandomized estimator was 1.7. This represents a variance reduction of 41%.

In summary, the derandomized estimator can be expected to achieve moderate but useful variance reductions over the standard variance estimator. Therefore, the derandomized estimator will probably be of most use in situations where the process is computationally expensive to simulate. In such a case, it is expensive to reduce estimator variance by increasing simulation run-length, and therefore the additional programming and computational effort required to use more sophisticated estimators is warranted.

To conclude this section, we provide an algorithm for computing the derandomized estimator *V_n*. Obviously, the first term in (10) is easy to compute, so we instead focus on computing

$$B_n \triangleq \sum_{j=0}^{n-1} f_n(X_j) \sum_{k=j+1}^n f_n(X_k) \prod_{i=j}^{k-1} \bar{w}(X_i, X_{i+1}).$$

Suppose (for now) that *w̄*(*X_i*, *X_{i+1}*) > 0 *∀i*. Define

$$\Lambda_j = \prod_{i=0}^{j-1} \bar{w}(X_i, X_{i+1})$$

for *j* ≥ 1 and *Λ*₀ = 1. Then

$$\begin{aligned} B_n &= \sum_{j=0}^{n-1} f_n(X_j) \sum_{k=j+1}^n f_n(X_k) \Lambda_k / \Lambda_j \\ &= \sum_{j=0}^{n-1} f_n(X_j) / \Lambda_j \sum_{k=j+1}^n f_n(X_k) \Lambda_k \\ &= \sum_{j=0}^{n-1} f_n(X_j) b_j / \Lambda_j, \quad \text{say.} \end{aligned}$$

Now consider the more general case, where we do not exclude the possibility that *w̄*(*X_i*, *X_{i+1}*) = 0 (i.e., given *X_i*, *i* + 1 is a regeneration time with probability 1). In this case, many of the terms that comprise *B_n* are 0. Let *U*(1), . . . , *U*(*r*^{*} - 1) be the indices such that *w̄*(*X_{U(i)-1}*, *X_{U(i)}*) = 0,

and define $U(0) = 0$, $U(r^*) = n + 1$. Then we can write

$$B_n = \sum_{r=1}^{r^*} \sum_{j=U(r-1)}^{U(r)-2} \sum_{k=j+1}^{U(r)-1} f_n(X_j) f_n(X_k) \prod_{i=j}^{k-1} \bar{w}(X_i, X_{i+1})$$

$$= \sum_{r=1}^{r^*} \sum_{j=U(r-1)}^{U(r)-2} f_n(X_j) b_j / \Lambda_j,$$

where

$$\Lambda_{U(r)} = 1, \quad \text{for } r = 0, 1, 2, \dots, r^* - 1,$$

$$\Lambda_k = \prod_{i=U(r-1)}^{k-1} \bar{w}(X_i, X_{i+1}) \quad \text{if } U(r-1) < k < U(r), \text{ and}$$

$$b_j = \sum_{k=j+1}^{U(r)-1} f_n(X_k) \Lambda_k \quad \text{if } U(r-1) \leq j < U(r) - 1.$$

We therefore obtain the following algorithm for computing B_n .

Algorithm 1. Calculation of B_n .

```

B = 0
for r = 1, ..., r*
  start = U(r - 1), stop = U(r) - 1
  if stop > start then (the sum is not vacuous)
    Λ_start = 1
    for j = (start + 1), ..., stop
      Λ_j = Λ_{j-1} w̄(X_{j-1}, X_j)
    end for
    b = 0
    for j = stop down to (start + 1)
      b = b + f_n(X_j) Λ_j
      B = B + b f_n(X_{j-1}) / Λ_{j-1}
    end for
  end if
end for
Return B_n = B
    
```

End: Calculation of B_n

This algorithm runs in $O(n)$ time and requires a maximum of $O(n)$ storage. This is comparable with a two-pass calculation of the standard regenerative TAVC estimator; see Shedler (1993).

Notice that in Algorithm 1 we divide certain quantities by $\Lambda_j \in (0, 1)$, where the Λ s are products of \bar{w} s. If Λ_j were represented on a finite-precision machine as 0, when in fact it was positive (but very small), the above algorithm could exhibit numerical instability. We did not encounter such difficulties on any example we considered, but we believe that unstable examples exist. Further details on this topic may be found in Henderson and Glynn (1997a), including an alternative approach that may be applied to unstable examples.

Henderson and Glynn (1997a) also consider implementing the derandomized estimator in sequential methods where the run-length is determined during the course of the simulation, and not beforehand. For such methods, an

updating scheme is needed to update previous estimates of the TAVC to reflect the new observations. Henderson and Glynn (1997a) give an updating scheme that, although computationally efficient, is not numerically stable. This is perhaps unsurprising because there is no known updating scheme for the standard regenerative method that is both computationally efficient and numerically stable (Shedler 1993, p. 122). Note, however, that Damerджи et al. (1997) introduce a nonstandard regenerative TAVC estimator that may be computed using a numerically stable, computationally efficient, one-pass algorithm.

APPENDIX

PROOF OF THEOREM 1. As stated in §3, it is easy to show that under our moment assumptions, $\tilde{V}_n \rightarrow \sigma^2 P_\mu$ a.s. $\forall \mu$. We will show that this convergence also occurs in L_1 — i.e., that $E_\phi |\tilde{V}_n - \sigma^2| \rightarrow 0$ as $n \rightarrow \infty$ — by showing that the sequence $\{\tilde{V}_n: n \geq 1\}$ is uniformly integrable. It will then follow that

$$E_\phi |E_\phi(\tilde{V}_n | X) - \sigma^2| \leq E_\phi |\tilde{V}_n - \sigma^2| \rightarrow 0,$$

as $n \rightarrow \infty$, so that L_1 convergence, and therefore convergence in probability, results.

To prove uniform integrability, we begin by noting that

$$\begin{aligned} \tilde{V}_n &= \frac{1}{n} \sum_{j=1}^{\ell(n)} Y_j(f_n)^2 + \frac{1}{n} \left(\sum_{k=T(n)}^n f_n(X_k) \right)^2 - f_n(X_n)^2/n \\ &\leq \frac{1}{n} \sum_{j=1}^{\ell(n)+1} Y_j(f_n)^2 + Y_{\ell(n)+1}^2(|f_n|)/n \\ &\leq \frac{1}{n} \sum_{j=1}^{\ell(n)+1} (Y_j(f_c) + (\alpha - \alpha_n)\tau_j)^2 \\ &\quad + \frac{1}{n} (Y_{\ell(n)+1}(|f_c|) + |\alpha - \alpha_n|\tau_{\ell(n)+1})^2 \\ &\leq \frac{2}{n} \sum_{j=1}^{\ell(n)+1} (Y_j(|f_c|) + |\alpha - \alpha_n|\tau_j)^2 \\ &= \frac{2}{n} \sum_{j=1}^{\ell(n)+1} (Y_j(|f_c|))^2 + \frac{4|\alpha - \alpha_n|}{n} \\ &\quad \cdot \sum_{j=1}^{\ell(n)+1} Y_j(|f_c|)\tau_j + \frac{2(\alpha - \alpha_n)^2}{n} \sum_{j=1}^{\ell(n)+1} \tau_j^2. \end{aligned} \tag{13}$$

The first term converges both a.s. and in expectation to $2E_\phi Y_1(|f_c|)^2/E_\phi \tau_1$ by Wald's lemma. Therefore, it is uniformly integrable (Billingsley 1986, p. 223). We will now show how to deal with the final term only, because the treatment of the second term is entirely similar.

Note that $\alpha - \alpha_n = -\frac{1}{n} \sum_{j=0}^{n-1} f_c(X_j)$ so that

$$\begin{aligned} \frac{(\alpha - \alpha_n)^2}{n} \sum_{j=1}^{\ell(n)+1} \tau_j^2 &\leq n^{-3} \left(\sum_{j=1}^{\ell(n)+1} Y_j(|f_c|) \right)^2 \sum_{j=1}^{\ell(n)+1} \tau_j^2 \\ &\leq n^{-3} \left(\sum_{j=1}^{n+1} Y_j(|f_c|) \right)^2 \sum_{j=1}^{n+1} \tau_j^2 \\ &\rightarrow (E_\varphi Y_1(|f_c|))^2 E_\varphi \tau_1^2, \end{aligned}$$

as $n \rightarrow \infty$ a.s. But notice that this convergence also occurs in expectation (simply expand out the products and take expectations). Therefore, the last term in (13) is uniformly integrable. The second term is handled similarly, so the proof is complete. \square

EXTENSION OF THEOREM 1 TO THE $m > 1$ CASE. We may use the same approach used for the $m = 1$ independent case. A check of the proof reveals that the only possible problem occurs when Wald's lemma is applied to show that \tilde{V}_n is uniformly integrable. However, it is possible to avoid this difficulty (Janson 1983) by noticing that if $Y = \{Y_n; n \geq 0\}$ is an m -dependent, nonnegative sequence and N is a stopping time with respect to Y , then

$$E \sum_{j=0}^{N+m} Y_j = E \sum_{j=0}^{N+m} EY_j.$$

Hence, we bound \tilde{V}_n by adding extra terms to the summations. The proof then goes through as before. \square

ACKNOWLEDGEMENTS

The authors thank the associate editor and referees for suggestions that led to considerable improvements in both the presentation and content of this paper. One referee's suggestions led to simplifications in the proof of Theorem 1. This work was supported by the Army Research Office under contract no. DAAL03-91-G-0319.

REFERENCES

- Andradóttir, S., J.M. Calvin, P.W. Glynn. 1995. Accelerated regeneration for Markov chain simulations. *Probab. Engrg. Inform. Sci.* **9** 497–523.
- Athreya, K.B., P. Ney. 1978. A new approach to the limit theory of recurrent Markov chains. *Trans. Amer. Math. Soc.* **245** 493–501.
- Billingsley, P. 1986. *Probability and Measure*. 2nd ed. Wiley, New York.
- Boender, C.G.E., H.E. Romeijn. 1995. Stochastic methods. *Handbook of Global Optimization*. R. Horst, P. Pardalos, eds. Kluwer, New York, 829–869.

- Chung, K.L. 1967. *Markov Chains with Stationary Transition Probabilities*. 2nd Ed. Springer-Verlag, New York.
- Damerdjij, H., S.G. Henderson, P.W. Glynn. 1997. Computational efficiency evaluation in simulation output analysis. Working Paper. Department of Engineering Science, University of Auckland, New Zealand.
- Gaver, D.P., I.G. O'Muircheartaigh. 1987. Robust empirical Bayes analysis of event rates. *Technometrics* **29** 1–15.
- Gilks, W.R., S. Richardson, D.J. Spiegelhalter, eds. 1996. *Markov Chain Monte-Carlo in Practice*. Chapman & Hall.
- Glynn, P.W. 1982. Simulation output analysis for general state-space Markov chains. Ph.D. Thesis, Dept. of Operations Research, Stanford University, Stanford, CA.
- . 1983. On the role of generalized semi-Markov processes in simulation output analysis. *Proc. 1983 Winter Simulation Conf.*, 38–42.
- . 1989. A GSMP formalism for discrete event systems. *Proc. IEEE* **77** 14–23.
- , D.L. Iglehart. 1987. A joint central limit theorem for the sample mean and regenerative variance estimator. *Ann. Oper. Res.* **8** 41–55.
- , —. 1993. Conditions for the applicability of the regenerative method. *Management Sci.* **39** 1108–1111.
- , P. L'Ecuyer. 1993. Likelihood ratio gradient estimation for stochastic recursions. *Advances Appl. Probab.* **27** 1019–1053.
- , W. Whitt. 1992. The asymptotic validity of sequential stopping rules for stochastic simulations. *Ann. Appl. Probab.* **2** 180–198.
- Grenander, U., M. Rosenblatt. 1984. *Statistical Analysis of Stationary Time Series*. 2nd ed., Chelsea, New York.
- Henderson, S., P.W. Glynn. 1997a. Derandomizing and rerandomizing variance estimators. Technical Report. Department of Engineering-Economic Systems and Operations Research, Stanford University, Stanford, CA.
- , —. 1997b. Regenerative steady-state simulation of discrete-event systems. Working Paper. Department of Engineering Science, University of Auckland, New Zealand.
- Ho, Y.C., ed. 1991. *Discrete Event Dynamical Systems: Analyzing Complexity in the Modern World*. IEEE Press.
- Janson, S. 1983. Renewal theory for M -dependent variables. *Ann. Probab.* **11** 558–568.
- Meyn, S.P., R.L. Tweedie. 1993. *Markov Chains and Stochastic Stability*. Springer Verlag.
- Mykland, P., L. Tierney, B. Yu. 1994. Regeneration in Markov chain samplers. *J. Amer. Statist. Assoc.* **90** 233–241.
- Nummelin, E. 1978. A splitting technique for Harris recurrent Markov chains. *Z. Wahrsch. verw. Gebiete* **43** 309–318.
- Patel, N.R., R.L. Smith, Z.B. Zabinsky. 1989. Pure adaptive search in Monte Carlo optimization. *Math. Programming* **43** 317–328.
- Pich, M.T. 1992. The QNET method for two-moment analysis of open queueing networks with general workstation capabilities. Ph.D. Thesis, Dept. of Operations Research, Stanford University, Stanford, CA.
- Shedler, G.S. 1993. *Regenerative Stochastic Processes*. Academic Press, New York.
- Whitt, W. 1980. Continuity of generalized semi-Markov processes. *Math. Oper. Res.* **5** 494–501.
- Wolff, R.W. 1989. *Stochastic Modeling and the Theory of Queues*. Prentice-Hall, Englewood Cliffs, NJ.