MATHEMATICS OF OPERATIONS RESEARCH
Vol. 27, No. 2, May 2002, pp. 253–271
Printed in U.S.A.

APPROXIMATING MARTINGALES FOR VARIANCE REDUCTION IN MARKOV PROCESS SIMULATION

SHANE G. HENDERSON AND PETER W. GLYNN

"Knowledge of either analytical or numerical approximations should enable more efficient simulation estimators to be constructed." This principle seems intuitively plausible and certainly attractive, yet no completely satisfactory general methodology has been developed to exploit it. The authors present a new approach for obtaining variance reduction in Markov process simulation that is applicable to a vast array of different performance measures. The approach relies on the construction of a martingale that is then used as an internal control variate.

Introduction. There are basically two different approaches that can be used to compute performance measures for complex stochastic processes. The first approach, which was the method of choice prior to the advent of computers, is to approximate the complex process by a simpler and more tractable process. The quality of the approximation is generally supported via a limit theorem of some kind. The second approach, more commonly used today, is to numerically compute the performance measure, often via simulation.

A natural question to ask is whether these two approaches can be sensibly combined. One way to attempt this is to use the tractable approximating process as an external control in a simulation of the process of interest (see Bratley et al. 1987, Ch. 2). In this method, both the original process X and its approximation Y are simulated simultaneously. It is hoped that correlation can be induced between the two processes and that the correlation can be used to reduce variance. The principal difficulty with this approach is *synchronization*, i.e., ensuring that the two processes are sufficiently "close" that significant correlation results. Indeed, for some approximations, it is entirely unclear how to perform synchronization, e.g., how does one synchronize a queue with a diffusion approximation? A second difficulty is that the approximating process must be simulated in concert with the process of primary interest. This additional overhead may prove very costly and may reduce the effectiveness of external control variates as an efficiency improvement technique.

In this paper, we propose a completely different means of exploiting the availability of a tractable approximation. Our methodology applies to a vast array of different performance measures related to Markov processes, including steady-state costs, finite-horizon cumulative costs, transient distributions, infinite-horizon discounted costs, and accrued costs up to a hitting time. The basic idea is to use the simpler approximating process to construct an appropriate martingale for the more complex process. The approach used to construct the martingale depends largely on the choice of performance measure. The martingale is then used as an internal control in a simulation of the original (less tractable) process. In contrast to the method of external controls, our approach does not involve a simultaneous simulation of the more tractable process. We call the method the "approximating martingale-process method" (AMPM).

It turns out that AMPM really demands only the availability of an approximation to the solution of an appropriately chosen system of linear equations. The particular linear system

Received October 13, 2000; revised July 11, 2001.
MSC 2000 subject classification. Primary: 65C40.
OR/MS subject classification. Primary: Probability/Markov processes.
Key words. Variance reduction, Markov process, simulation, martingale.

0364-765X/02/2702/0253/$05.00
1526-5471 electronic ISSN,

to be approximated depends on the performance measure to be computed. The appropriate martingale is then constructed from the approximating solution. One way to calculate an approximating solution is to use the exact solution to an approximating (more tractable) model. However, another approach to obtaining an approximate solution is to use a numerical algorithm to develop a crude approximation to the solution of the appropriate linear system. This algorithm may be simulation based, or it may involve some nonsampling-based methodology.

Henderson and Meyn (1997, 2000) apply AMPM to simulations of multiclass queueing networks. They use a combination of analytical and numerical techniques to obtain and compute the approximate solution.

The methodology presented in Andradóttir et al. (1993) is closely related to the theory presented in §6. The authors obtained variance reduction in steady-state simulations of discrete-time Markov chains on a finite state space.

Other authors have developed methodology for incorporating prior information into simulations. Schmeiser and Taaffe (1994) show how to use an approximating process as an external control variate and as a direct approximation simultaneously. Nelson et al. (1997) discuss three methods for combining deterministic approximations with simulation. Schmeiser et al. (2001) examine the small-sample properties of biased control variates. Emsermann and Simon (2000) develop "quasicontrol variates," which are control variates with a mean that is unknown and estimated in an auxiliary simulation.

This paper is organized as follows. In §1, we present a simple example to motivate the overall methodology. Then, in the next five sections, we show how to define approximating martingales for a variety of performance measures for discrete-time Markov chains (DTMCs) on a general state space and for continuous-time Markov chains (CTMCs) on a finite state space.

In §7, we show how our theory generalizes to processes satisfying stochastic differential equations. In that setting, Itô calculus provides the basic mechanism for defining an appropriate class of martingales. Indeed, the reader familiar with stochastic calculus may find this section the ideal starting point for reading this paper.

In §8, we provide a concrete example of the application of the theory in this paper. We define an estimator for the mean steady-state waiting time of customers in the single-server queue that has provably better performance in heavy traffic than a more standard estimator.

For general processes, one must typically adjoin supplementary variables to the state space to ensure that the resulting process is Markov. For such processes, the identification and application of approximating martingales is not as easy as it is for the processes we consider here, but it can still be done. This issue, and others, are addressed in Henderson (1997) and Henderson and Glynn (2001). Further comments on this issue are also given in §9. Section 9 also demonstrates that our approach may be applied in ways that at first may not appear obvious. In particular, we discuss how the approximating martingale may be "turned on and off," depending on the behaviour of the simulated process.

1. A first example. We begin with a simple example that demonstrates the key steps in applying our methodology.

EXAMPLE 1. Stephen G. Henry (any resemblance to a real person is unintentional) is hounded by the immigration departments of two countries and therefore divides his time between country C_0 and country C_1. He spends an exponentially distributed period with mean λ_i^{-1} in C_i before immigration catches up with him and sends him to country C_{1-i}. We wish to find α, the long-run fraction of time he spends in C_0.

Clearly, the answer is $\lambda_1/(\lambda_0+\lambda_1)$, but the following development should be instructive. Define $Y = (Y(t)\colon t \geq 0)$ to be a continuous-time Markov chain (CTMC) on state space

$S = \{C_0, C_1\}$. Define $f\colon S \to \mathbb{R}$ by $f(C_0) = 1 = 1 - f(C_1)$. A natural estimator of α is

$$\alpha(t) \triangleq \frac{1}{t}\int_0^t f(Y(s))\,ds, \tag{1}$$

the fraction of time spent in C_0 in $[0, t]$.

DEFINITION 1. Define $P_x(\cdot) = P(\cdot|Y(0) = x)$ and $P_\mu(\cdot) = \int_S P_x(\cdot)\mu(dx)$. We say that M is a P_μ martingale if it is a martingale under the probability law P_μ with respect to the natural filtration $(\mathcal{F}_t\colon t \geq 0)$, where $\mathcal{F}_t = \sigma(Y(s)\colon 0 \leq s \leq t)$.

It can be shown (see §2) that for any $u\colon S \to \mathbb{R}$, $M = (M(t)\colon t \geq 0)$ is a mean zero P_μ martingale, for any μ, where

$$M(t) = u(Y(t)) - u(Y(0)) - \int_0^t Au(Y(s))\,ds,$$

and A is the rate matrix of Y, given by

$$A = \begin{bmatrix} -\lambda_0 & \lambda_0 \\ \lambda_1 & -\lambda_1 \end{bmatrix}.$$

Since $EM(t) = 0$ for all t, we could estimate α by

$$\alpha'(t) = \alpha(t) - M(t)/t$$

without adding any bias to the estimator $\alpha(t)$.

Let us define a particular u by $u(0) = 0$ and $u(1) = -(\lambda_0 + \lambda_1)^{-1}$. (This choice may appear somewhat arbitrary, but we will see in §6 why it is appropriate.) Then,

$$Au(x) = -f(x) + \lambda_1/(\lambda_0 + \lambda_1), \tag{2}$$

and we see that

$$\begin{aligned} \alpha'(t) &= \frac{1}{t}\int_0^t f(Y(s))\,ds - \frac{u(Y(t)) - u(Y(0))}{t} + \frac{1}{t}\int_0^t Au(Y(s))\,ds \\ &= \frac{\lambda_1}{\lambda_0 + \lambda_1} + \frac{u(Y(0)) - u(Y(t))}{t}. \end{aligned}$$

The variance of $\alpha(t)$ is of the order t^{-1}, whereas that of $\alpha'(t)$ is of the order t^{-2}. Moreover, the estimators have the same expectation, so that $\alpha'(t)$ is the preferred estimator.

There are three central themes that are exemplified above. The first is that we add a martingale (or a suitable multiple of a martingale) to a more standard estimator. The second is the identification of a suitable linear system of Equations (2), which dictates the optimal choice of martingale. The final theme is that with the "right" choice of martingale, virtually unlimited variance reduction is possible.

Let us now see how these same themes unfold in the context of a variety of performance measures.

2. Accrued costs prior to absorption. In this and the next four sections, $X = (X_n\colon n \geq 0)$ will denote a Markov chain evolving on a general (not necessarily finite or countably infinite) state space Σ, and $Y = (Y(t)\colon t \geq 0)$ will denote a CTMC evolving on a finite state space S. The function f will represent a real-valued cost function on either Σ or S, depending on the context.

Let $C \subset \Sigma$ be a set of absorbing states, and let C^c denote the complement of C in Σ. Define

$$T = \inf\{n \geq 1\colon X_n \in C\}$$

to be the time of absorption in C, and for $x \in C^c$, define

$$u^*(x) = E_x \sum_{k=0}^{T-1} f(X_k), \tag{3}$$

where $E_x(\cdot) \triangleq E(\cdot|X_0 = x)$. To ensure that $u^*(x)$ is well-defined and finite for all $x \in C^c$, we impose the following conditions:

CONDITION 1. $\|f\| \triangleq \sup_{x\in C^c} |f(x)| < \infty$.

CONDITION 2. There is a finite-valued nonnegative function g: $\Sigma \to \mathbb{R}$ such that for some $\epsilon > 0$ and for all $x \in C^c$, $Pg(x) \le g(x) - \epsilon$, where

$$Pg(x) \triangleq E_x g(X_1) = \int_\Sigma g(y)P(x, dy).$$

Condition 1 ensures that for $x \in C^c$, $u^*(x) \le \|f\| E_x T$, and Condition 2 ensures that $E_x T < \infty$ for all $x \in C^c$ (see, for example, Meyn and Tweedie 1993, Theorem 11.3.4).

Suppose that we wish to compute $\alpha \triangleq u^*(x)$. Let U_k be the observed cumulative cost for the kth realization of $(X_i: 0 \le i \le T)$, where $X_0 = x$. Then α may be estimated by α_n, where α_n is the sample average of the U_ks.

Define B to be the restriction of P to C^c, so that $B(x, dy) = P(x, dy)$ for $x, y \in C^c$. It is well known that u^* then satisfies the linear system

$$u = f + Bu,$$

where $Bu(x) \triangleq \int_{C^c} u(y)B(x, dy)$ for $x \in C^c$. It will prove convenient to have u^* and f defined for all $x \in \Sigma$. For $x \in C$, define $u^*(x) = f(x) = 0$, so that $u^* = f + Pu^*$.

Now, suppose that $M = (M_n: n \ge 0)$ is a P_x martingale with respect to the natural filtration $(\mathcal{F}_n: n \ge 0)$, where $\mathcal{F}_n = \sigma(X_0, \ldots, X_n)$, such that $E_x M_T = 0$. Let $M_T(k)$ be the value of M_T for the kth realization of $(X_i: 0 \le i \le T)$. Then the martingale estimator is given by

$$\alpha'_n = \frac{1}{n} \sum_{k=1}^{n} (U_k - M_T(k))$$

and is unbiased. The following proposition identifies a class of approximating martingales and identifies the "optimal" choice of martingale.

PROPOSITION 1. *Suppose that Conditions* 1 *and* 2 *hold. Let* u: $\Sigma \to \mathbb{R}$ *be such that* $u(x) = 0$ *for* $x \in C$, $0 \le u(x) < \infty$ *for* $x \in C^c$, *and for some constant* $b < \infty$, $|(P - I)u(x)| \le b$ *for* $x \in C^c$. *Then* $M = (M_n: n \ge 0)$ *is a* P_x *martingale for all* $x \in \Sigma$, *where*

$$M_n = u(X_n) - u(X_0) - \sum_{k=0}^{n-1} (P - I)u(X_k).$$

If, in addition, $E_x \sum_{k=0}^{T-1} u(X_k) < \infty$, *then* $E_x M_T = 0$ *for all* $x \in C^c$.

If $u = u^*$, *then under* P_x, $M_T = -u^*(x) + U$, *so that* α'_n *has zero variance.*

Proposition 1 shows that $u = u^*$ should be used to define the martingale M. Of course, we are trying to compute u^*, so it is unknown, but an approximate solution of u^* should be a good choice for u. In §8, we will see how to use approximations to the process X to obtain u. Other methods are of course possible. For instance, u could be an approximate numerical solution to $u = f + Bu$.

It is very reasonable to use an estimator of the form $n^{-1} \sum_{k=1}^{n} [U_k + \beta_n M_T(k)]$, where β_n is selected to attempt to minimize the variance. We chose $\beta_n = -1$ in the above analysis to simplify the exposition but one could also select β_n using more typical control variate techniques, e.g., as an estimate of $-\mathrm{Cov}(U, M_T)/\mathrm{Var}\, M_T$.

The analysis for the CTMC is similar. In this case, $T = \inf\{t \geq 0: Y(t) \in C\}$, and $u^*(x) = E_x \int_0^T f(Y(s))\, ds$. Note that with this definition, $u^*(x) = 0$ for $x \in C$. Suppose that for each $x \in C^c$, there is some $t = t(x)$ such that $P(Y(t) \in C | Y(0) = x) > 0$. Because of the finite state space of Y, a simple geometric trials argument shows that $E_\mu T < \infty$ for all μ. Hence, u^* is well defined and finite, and it is easy to show that u^* solves the linear system $Au = -f$.

Again, we define U_k as the cumulative cost accrued until absorption in C for the kth realization of $(Y(t): 0 \leq t \leq T)$ and let α_n be the sample average of the U_ks. The estimator α'_n is given by $\alpha_n - M(T, k)$, where $M(T, k)$ is the kth realization of a certain martingale observed at the stopping time T. The following proposition identifies an appropriate class of martingales.

PROPOSITION 2. *Let u: $S \to \mathbb{R}$, and define $M = (M(t)\colon t \geq 0)$ by*

$$M(t) = u(Y(t)) - u(Y(0)) - \int_0^t Au(Y(s))\, ds.$$

Then, for any μ, M is a P_μ martingale, and if $E_\mu T < \infty$, then $E_\mu M(T) = 0$.

If we choose $u = u^*$, then α'_n has zero variance.

3. Infinite horizon discounted costs. Let f be bounded, and suppose that $0 \leq h(x) \leq \delta < 1$ for all $x \in \Sigma$. Then the infinite-horizon discounted cost $u^*(x) = E_x U$ is well defined and bounded (by $\|f\|/(1-\delta)$), where

$$U = \sum_{n=0}^{\infty} f(X_k) \prod_{k=0}^{n-1} h(X_k). \tag{4}$$

Because of the infinite horizon involved, it is not possible to simulate replicates of U in finite time. Nevertheless, Fox and Glynn (1989) show how to construct unbiased estimators of $u^*(x)$ that may be simulated in finite time. They suggest randomizing over the infinite horizon in (4), thereby replacing (4) by a finite sum up to a random limit. A similar approach may be applied to the estimators constructed in this section, and so henceforth we consider this issue settled.

It is straightforward to show that u^* satisfies the linear system,

$$u(x) = f(x) + h(x)Pu(x) \quad \forall\, x \in \Sigma.$$

Suppose we were to simulate replicates of $(X_n: n \geq 0)$ under P_x, obtaining independent replicates $U_1, \ldots, U_n$ of U. We could then estimate $\alpha = u^*(x)$ by $\alpha_n = n^{-1} \sum_{k=1}^n U_k$. We will instead estimate α by

$$\alpha'_n = \frac{1}{n} \sum_{k=1}^{n} U_k - M_\infty(k),$$

where $M_\infty(k)$ is the kth replicate of the limiting value of an approximating martingale. The following proposition provides a wide class of approximating martingales and identifies the optimal choice.

PROPOSITION 3. *Let h be defined as above, and suppose that u: $\Sigma \to \mathbb{R}$ is bounded. Then $M = (M_n\colon n \geq 0)$ is a P_x martingale for all $x \in \Sigma$, where*

$$M_n = \sum_{k=1}^{n} [u(X_k) - Pu(X_{k-1})] \prod_{j=0}^{k-1} h(X_j).$$

Furthermore, $M_n \to M_\infty$ a.s., where $E_x M_\infty = 0$.

If $u = u^*$, then $M_\infty = U - u^*(X_0)$, so that under P_x, α'_n is a zero variance estimator for $u^*(x)$.

Note that again, u^* is the "optimal" choice, suggesting that an approximation to u^* might be an effective choice.

In the CTMC case, assume that $h(x) \geq \delta > 0$ for all $x \in S$, so that the expected infinite-horizon cost $u^*(x) = E_x U$ is well defined and bounded (by $\|f\|/\delta$), where

$$U = \int_0^\infty f(Y(t))e^{-V(t)}\,dt,$$

and $V(t) = \int_0^t h(Y(s))\,ds$.

Again, we define α_n to be the sample average of independent replicates $U_1, \dots, U_n$ of U, and $\alpha_n' = \alpha_n - n^{-1}\sum_{k=1}^n M(\infty, k)$, where $M(\infty, k)$ is the limiting value of a certain martingale from the kth replication.

It is straightforward to show that u^* satisfies the linear system,

$$Au(x) - h(x)u(x) = -f(x) \quad \forall\, x \in S,$$

and the appropriate class of martingales is now given by the following result.

PROPOSITION 4. *Let u: $S \to \mathbb{R}$ and let $M = (M(t)\colon t \geq 0)$ be given by*

$$M(t) = e^{-V(t)}u(Y(t)) - u(Y(0)) - \int_0^t e^{-V(s)}[Au(Y(s)) - h(Y(s))u(Y(s))]\,ds.$$

If $h(x) > 0$ for all x, then M is a P_μ martingale for all μ. Furthermore, $M(t) \to M(\infty)$ a.s. as $t \to \infty$, and $E_\mu M(\infty) = 0$ for all μ.

If $u = u^*$, then under P_x, $M(\infty) = -u^*(x) + U$, and then α_n' is a zero-variance estimator.

We will establish an analogue of this result for processes satisfying stochastic differential equations in §7, and so the proof is omitted.

4. Transient distributions. We now extend the approach outlined in earlier sections to the problem of computing $u_n^*(x) = E_x f(X_n)$. To ensure that this quantity exists, we will assume that f is bounded. A reasonable estimator of $u_n^*(x)$ is given by

$$\alpha_m \triangleq \frac{1}{m}\sum_{k=1}^m f(X_n(k)),$$

where $X_n(k)$ is the observed value of X_n on the kth replication of $(X_0, \dots, X_n)$ under P_x. We define the alternative estimator α_m' by

$$\alpha_m' \triangleq \frac{1}{m}\sum_{k=1}^m f(X_n(k)) - M_n(k),$$

where $M_n(k)$ is the kth realization of an appropriate martingale at time n.

Before defining a class of approximating martingales, observe that $(u_i^*(x)\colon i = 0, \dots, n)$ is the exact solution to the linear system,

$$\begin{aligned} &u_0 = f \quad \text{and} \\ &u_j = Pu_{j-1} \quad \text{for } j = 1, \dots, n. \end{aligned} \tag{5}$$

PROPOSITION 5. *Let $(u_j\colon 0 \leq j \leq n)$ be a sequence of $n+1$ bounded real-valued functions on Σ. For fixed n, define $M = (M_j\colon 0 \leq j \leq n)$ by $M_0 = 0$ and*

$$M_j = \sum_{k=1}^j u_{n-k}(X_k) - Pu_{n-k}(X_{k-1}).$$

Then M is a P_μ martingale for all μ. Furthermore, if $u_i = u_i^*$ for $0 \leq i \leq n$, then under P_x, $M_n = f(X_n) - u_n^*(x)$, and then α_n' is a zero variance estimator of $u_n^*(x)$.

In continuous time, our goal is to estimate $u^*(t,x) = E_x f(Y(t))$, where $t > 0$. The estimator α_n is again a sample average of replicates of $f(Y(t))$, and the alternative estimator is a sample average of replicates of $f(Y(t)) - M(t)$, where $M(t)$ is the value at time t of a certain martingale.

For a differentiable (in t) function $u(t,x)$, define

$$u_1(s,x) = \left.\frac{\partial u(t,x)}{\partial t}\right|_{t=s}.$$

Before defining a class of approximating martingales, observe that $u^*(t,x)$ is the exact solution to the linear system,

$$\begin{aligned} u_1(s,\cdot) &= Au(s,\cdot), \quad \text{where} \\ u(0,\cdot) &= f(\cdot). \end{aligned}$$

Note that this is just the Kolmogorov backwards equation for Y (see, e.g., Karlin and Taylor 1981, Chapter 14).

PROPOSITION 6. *Suppose that $(u(s,\cdot): 0 \le s \le t)$ is a set of real-valued functions defined on S with the property that $u(s,x)$ is continuously differentiable in s for all $x \in S$, i.e., $u_1(s,x)$ is continuous in s for all x. Then $M = (M(s): 0 \le s \le t)$ is a P_μ martingale for all μ, where*

$$M(s) = u(t-s, Y(s)) - u(t, Y(0)) - \int_0^s [Au(t-r, Y(r)) - u_1(t-r, Y(r))]\,dr.$$

If $u(s,\cdot) = u^*(s,\cdot)$ for $0 \le s \le t$, then under P_x, $M(t) = f(Y(t)) - u^*(t,x)$, so that α'_t is a zero variance estimator of $E_x f(Y(t))$.

5. Finite horizon cumulative costs. Suppose we are interested in computing $u_n^*(x)$, where

$$u_n^*(x) \stackrel{\triangle}{=} E_x \sum_{k=0}^{n} f(X_k).$$

To ensure that u_n^* exists, we require that $\|f\| < \infty$. We define the estimator α_m as the sample mean of m realizations of $\sum_{k=0}^n f(X_k)$. The alternative estimator α'_m will again be defined as the sample mean of m realizations of $\sum_{k=0}^n f(X_k) - M_n$, where M_n is a martingale of the form defined in Proposition 5. To get some idea of a good choice of functions u used to define the martingale, note that $u_n^*(x)$ satisfies the following linear system:

$$\begin{aligned} u_0 &= f, \\ u_j &= Pu_{j-1} + f \quad \text{for } j = 1, \ldots, n. \end{aligned}$$

It is easy to show that, if M_n is defined as in Proposition 5 and $u_j = u_j^*$ for $j = 0, \ldots, n$, then

$$M_n = \sum_{j=0}^{n} f(X_j) - u_n(X_0),$$

so that under P_x, α'_n is a zero-variance estimator for $u_n^*(x)$.

In continuous time, our goal is to estimate $u^*(t,x) = E_x \int_0^t f(Y(s))\,ds$, where $t > 0$. The estimator α_n is again a sample average of replicates of $\int_0^t f(Y(s))\,ds$, and the alternative estimator is a sample average of replicates of $\int_0^t f(Y(s))\,ds - M(t)$, where $M(t)$ is the value at time t of a martingale of the form given in Proposition 6.

To get some idea of a good choice of functions $u(t,\cdot)$ to define the martingale M, observe that $u^*(t,x)$ satisfies the linear system,

$$\begin{aligned} u_1(t,\cdot) &= Au(t,\cdot) + f(\cdot), \quad \text{where} \\ u(0,\cdot) &= 0. \end{aligned}$$

Again, we expect that if u approximately satisfies this linear system, then the estimator α_t' will have lower variance than α_t.

6. Average steady-state cost. In the previous sections, we considered performance measures that may be computed using terminating simulations. In this section, we show how to define approximating martingales for steady-state simulation. Suppose that X is a positive Harris recurrent discrete-time Markov chain with stationary probability measure π. Suppose that $\pi|f| \triangleq \int_\Sigma |f(x)|\pi(dx) < \infty$ and that we wish to compute $\alpha = \pi f$.

By the strong law (Asmussen 1987, Proposition 3.7),

$$\alpha_n \triangleq \frac{1}{n}\sum_{k=0}^{n-1} f(X_k) \to \alpha \text{ a.s.}$$

as $n \to \infty$, so that α_n is a consistent estimator of α.

It is not immediately clear what the appropriate linear system for this performance measure is. However, consider *Poisson's equation*,

$$(P-I)u(x) \triangleq Pu(x) - u(x) = -(f(x) - \alpha) \quad \forall x \in \Sigma. \tag{6}$$

Suppose u^* is a solution to (6). Then $(P-I)u^*(x) + f(x)$ is a zero variance estimator of α. Accordingly, we define the estimator,

$$\begin{aligned} \alpha_n' &= \frac{1}{n}\sum_{k=0}^{n-1}[f(X_k) + (P-I)u(X_k)] \\ &\approx \alpha_n - n^{-1}M_n, \end{aligned} \tag{7}$$

where $M = (M_n\colon n \geq 0)$ is an appropriate martingale (see Proposition 7). Observe that if u is π integrable, then $(P-I)u$ is π integrable, and $\pi[(P-I)u] = [\pi(P-I)]u = [\pi - \pi]u = 0$. So by the strong law, $\alpha_n' \to \alpha$ a.s. as $n \to \infty$.

PROPOSITION 7. *Suppose that $\pi|f| < \infty$ and u is π integrable. Then $M = (M_n\colon n \geq 0)$ is a P_π martingale with $E_\pi M_n = 0$ for all n, where*

$$M_n = u(X_n) - u(X_0) - \sum_{k=0}^{n-1}(P-I)u(X_k).$$

REMARK 1. One might prefer to use an estimator of the form $\alpha_n - M_n/n$ over α_n', as it introduces no further bias. We prefer the use of α_n' over such an estimator for three reasons. First, if $u = u^*$, then α_n' has zero variance. This is not the case for $\alpha_n - M_n/n$. Second, the bias of α_n' is the sum of the bias of α_n and the expected value of the second term in (7). It is not clear then that the bias present in α_n' is greater than that of α_n. Third, the bias of these estimators is typically of the order n^{-1}, as is the variance. Since the mean-squared error is the sum of *squared* bias and variance, bias is asymptotically negligible, and so computation of the "peripheral term," $n^{-1}(u(X_0) - u(X_n))$, does not appear especially beneficial.

In §8, we will apply this theory to estimating the mean steady-state waiting time in the single-server queue.

Now suppose that the CTMC Y is irreducible and positive recurrent with stationary distribution π. Again, our goal is to estimate $\alpha = \pi f$. It is known that the estimator

$$\alpha(t) \triangleq \frac{1}{t}\int_0^t f(Y(s))\,ds \to \alpha \text{ a.s.}$$

as $t \to \infty$, as follows from regenerative process theory.

The continuous-time version of Poisson's equation (see, e.g., Glynn and Meyn 1996) is given by

$$Au(x) = -(f(x) - \alpha) \quad \forall\, x \in S.$$

If u^* solves this equation, and we take $u = u^*$, then $\alpha'(t)$ will have zero variance, where

$$\begin{aligned}\alpha'(t) &\triangleq \frac{1}{t}\int_0^t f(Y(s)) + Au(Y(s))\,ds \\ &\approx \alpha(t) - t^{-1}M(t),\end{aligned}$$

and the P_μ martingale (for all μ) $M = (M(t)\colon t \geq 0)$ is defined by

$$M(t) = u(Y(t)) - u(Y(0)) - \int_0^t Au(Y(s))\,ds.$$

Hence, the theory carries over to CTMCs on finite state space without difficulty. Henderson (1997) also discusses CTMCs on a countably infinite state space.

7. Stochastic differential equations. In the previous sections, we have shown that, for a wide variety of performance measures, one can define approximating martingales for both discrete-time Markov processes on a general state space and CTMCs on a finite state space. The application of the approximating martingale method is not, however, limited to such processes. In this section, we demonstrate how to obtain approximating martingales for processes satisfying stochastic differential equations.

Let $X = (X_t\colon t \geq 0)$ be a real-valued continuous-time process satisfying

$$dX_t = \mu(X_t)\,dt + \sigma(X_t)\,dB_t, \tag{8}$$

where $\mu, \sigma\colon \mathbb{R} \to \mathbb{R}$ are continuous, $B = (B_t\colon t \geq 0)$ is standard Brownian motion, and (8) is interpreted in the Itô sense. We will assume that μ and σ are bounded to simplify the exposition. (Note that under this assumption, the process X is nonexplosive.)

The first performance measure we will consider is infinite-horizon discounted costs. Suppose that $f, h\colon \mathbb{R} \to \mathbb{R}$ are bounded, and $h(x) \geq \delta > 0$ for all $x \in \mathbb{R}$. Set

$$U = \int_0^\infty f(X_t)e^{-V_t}\,dt,$$

where $V_t = \int_0^t h(X_s)\,ds$. Define $u^*(x) = E_x U \leq \delta^{-1}\|f\|$.

As pointed out in §3, it is impossible to obtain replicates of U from a simulation, as this would require an infinite amount of computation. However, Fox and Glynn (1989) demonstrate that it is possible to obtain unbiased estimators of u_x^* that may be computed in finite time by "randomizing" over the infinite horizon. This point is somewhat peripheral to our discussion, and we henceforth assume it settled.

The estimator α_n of $u^*(x)$ is defined to be the sample mean of n realizations of U under P_x. The estimator α_n' is defined to be the sample mean of n realizations of $U - M_\infty$, where

M_∞ is the almost sure limit of an appropriate martingale. Before defining a useful class of martingales, observe that u^* satisfies the linear system

$$Au(x) - h(x)u(x) = -f(x) \quad \forall x \in \mathbb{R},$$

where

$$Au(x) \stackrel{\triangle}{=} \mu(x)u'(x) + \frac{1}{2}\sigma^2(x)u''(x).$$

On page 191 of Karlin and Taylor (1981), similar systems are derived for related performance measures.

PROPOSITION 8. *If h is bounded away from 0 and ∞, and if u: $\mathbb{R} \to \mathbb{R}$ is twice continuously differentiable with u' bounded, then $M = (M_t\colon t \ge 0)$ is a P_μ martingale for all μ, where*

$$M_t = e^{-V_t}u(X_t) - u(X_0) - \int_0^t e^{-V_s}[Au(X_s) - h(X_s)u(X_s)]\,ds.$$

Furthermore, $M_t \to M_\infty$ a.s. as $t \to \infty$, where $E_\mu M_\infty = 0$.

If $u = u^*$, then α_n' is a zero variance estimator of $u^*(x)$.

We believe that the proof of this result is instructive, and so it will be included in this section. Before presenting this proof, it is worth reiterating the point from earlier sections that this proposition shows that infinite variance reduction results if u is chosen as u^*. This result is somewhat akin to importance sampling where infinite variance reduction is possible in some contexts if the right change of measure is employed. Unfortunately, the right change of measure depends on knowledge of the quantity that is to be estimated, and so it is unobtainable in practice. However, as in importance sampling, Proposition 8 suggests that if u is a good approximation to u^*, then considerable variance reduction may result. The approximation u to u^* may be obtained in any fashion. In the next section, we will see how to use an approximation to the process X that yields an analytical approximation to u^*. Another possible approach is to use an approximation to u^* that is obtained from some numerical scheme.

PROOF OF PROPOSITION 8. Consider the two-dimensional process $W = (W_t\colon t \ge 0)$, where $W_t = (X_t, V_t)$. Observe that

$$dW_t = \begin{pmatrix} \mu(X_t)\,dt + \sigma(X_t)\,dB_t \\ h(X_t)\,dt \end{pmatrix}.$$

Define $g(x, v) = e^{-v}u(x)$. By Itô's formula (Chung and Williams 1990, Theorem 5.10, p. 109),

$$g(W_t) - g(W_0) = \int_0^t \frac{\partial g}{\partial x}(W_s)\,dW_s + \int_0^t \frac{\partial g}{\partial v}(W_s)\,dV_s + \frac{1}{2}\int_0^t \frac{\partial^2 g}{\partial x^2}(W_s)\sigma^2(X_s)\,ds.$$

We therefore obtain

$$\text{(9)} \quad e^{-V_t}u(X_t) - u(X_0) - \int_0^t e^{-V_s}(Au(X_s) - h(X_s)u(X_s))\,ds = \int_0^t e^{-V_t}u'(X_s)\sigma(X_s)\,dB_s.$$

Observe that the right-hand side of (9) is a stochastic integral with bounded integrand and, hence, a martingale. Therefore, M is a martingale.

Observe that M_t is bounded by $\|u'\|\|\sigma\|/\delta$ for all t, and so $M_t \to M_\infty$ a.s. by the martingale convergence theorem. Dominated convergence then implies that $EM_\infty = 0$.

Finally, if $u = u^*$, then note that

$$M_\infty = -u^*(X_0) + \int_0^\infty e^{-V_t} f(X_t)\, dt = -u^*(X_0) + U,$$

and therefore, under P_x, $\alpha_n' = u^*(x)$. □

As a second example, we show how to obtain approximating martingales for estimating transient distributions. Let $u^*(t, x) \triangleq E_x f(X_t)$, where f is assumed to be bounded. A reasonable estimator of $u^*(t, x)$ is the sample mean α_n of n replicates of $f(X_t)$ under P_x. Define the estimator α_n' as the sample mean of n replicates of $f(X_t) - M_t$ under P_x, where M_t is the value of a certain martingale at time t.

For a given function $u(t, x)$, define

$$u_1(s, x) = \left.\frac{\partial u(t, x)}{\partial t}\right|_{t=s} \quad \text{and} \quad u_2(t, y) = \left.\frac{\partial u(t, x)}{\partial x}\right|_{x=y}.$$

Before defining a suitable class of martingales, observe that $u^*(t, x)$ satisfies the linear system

$$\begin{aligned} u_1(s, \cdot) &= Au(s, \cdot), \quad \text{where} \\ u(0, \cdot) &= f(\cdot). \end{aligned}$$

For a proof, See Karlin and Taylor (1981).

PROPOSITION 9. *Suppose that f is bounded and let $u(t, x)$ be twice continuously differentiable in x and continuously differentiable in t. Suppose that for some constant K, $|u_2(s, x)| \le K$ for all $s \in [0, t]$ and all $x \in \mathbb{R}$. Then $M = (M_s\colon 0 \le s \le t)$ is a P_μ martingale for all μ, where*

$$M_s = u(t-s, X_s) - u(t, X_0) - \int_0^t [Au(t-s, X_s) - u_1(t-s, X_s)]\, ds.$$

If $u = u^*$, then $M_t = f(X_t) - u^*(t, X_0)$, so that under P_x, α_n' is a zero variance estimator of $u^*(t, x)$.

8. Waiting times in the $GI/G/1$ queue. In this section, we will apply our methodology to estimating the mean steady-state waiting time in the single-server (GI/G/1) queue. We will show that in heavy traffic, the martingale estimator outperforms the standard estimator. To do so, we need to deal with sequences of GI/G/1 queues. We begin by defining these systems.

Let $(\overline{V}_n\colon n \ge 0)$ and $(\overline{U}_n\colon n \ge 1)$ be independent sequences of i.i.d. r.v.s, with $E\overline{V}_0 = E\overline{U}_1 = \mu^{-1}$. Consider now a family of queues, defined in terms of these building blocks and parameterized by $\rho < 1$.

The ρth system consists of the sequences $(V_n(\rho)\colon n \ge 0)$ and $(U_n(\rho)\colon n \ge 1)$, where $V_n(\rho) = \overline{V}_n$ and $U_n(\rho) = \overline{U}_n/\rho$. The ρth system then has an arrival rate of $\lambda \triangleq \mu\rho$ and traffic intensity ρ. Let $X_n(\rho) \triangleq V_{n-1}(\rho) - U_n(\rho)$ and let $W(\rho) = (W_n(\rho) : n \ge 0)$ be the customer waiting time sequence (in the queue) for the ρth system. (When there is no danger of confusion, we will drop the index ρ, as we have already done with the arrival rate λ.) Then $W_0 = 0$ (assuming the first customer arrives at time 0), and for $n \ge 1$, $W_n = [W_{n-1} + X_n]_+$, where $[x]_+$ denotes the positive part of x; see, e.g., Asmussen (1987). Observe that W is a Markov chain on $[0, \infty)$.

If $\rho < 1$, then $W_n \Rightarrow W_\infty$ (Asmussen 1987) and Kiefer and Wolfowitz (1956) showed that if $\rho < 1$ and $EV_0^{k+1} < \infty$, then $EW_\infty^k < \infty$ for $k \geq 1$. We wish to estimate EW_∞. The estimator α_n, as defined in §6, is given by

$$\alpha_n = \frac{1}{n}\sum_{k=0}^{n-1} f(W_k),$$

where $f(w) = w$. Other estimators have been suggested by Asmussen (1990) and Minh and Sorli (1983). The martingale estimator proposed here is compared with α_n and the Minh-Sorli estimator in Henderson (1997); see the remarks following Theorem 10.

According to §6, to define the martingale estimator, we need to obtain an approximation to the solution to Poisson's equation. It is known (see, e.g., Asmussen 1992), that the process $W(\rho)$ may be approximated by a reflected Brownian motion (RBM) $\widetilde{W} = (\widetilde{W}(t)\colon t \geq 0)$, with drift $-(1-\rho)/\lambda$ and diffusion coefficient $\sigma^2 = \mathsf{Var}\,\overline{U}_1 + \mathsf{Var}\,\overline{V}_0$. We will solve Poisson's equation for $\widetilde{W}$ to obtain the necessary approximation u to the solution to Poisson's equation for $W(\rho)$.

If $\rho < 1$, then the stationary distribution of $\widetilde{W}$ is exponential with mean $\eta = \lambda\sigma^2/(2(1-\rho))$ (see, e.g., Harrison 1990). Thus, Poisson's equation for $\widetilde{W}$ is

$$\frac{\sigma^2}{2}u''(x) - \frac{1-\rho}{\lambda}u'(x) = -(x-\eta), \tag{10}$$

subject to the initial conditions $u(0) = u'(0) = 0$. Note that if u solves (10), then so does $u + c$ for any constant c, so that the first initial condition picks out one solution. The second condition is required to ensure that u lies in the domain of the generator of the RBM; see Henderson (1997) for details. The solution to this ordinary differential equation is $u(x) = \lambda x^2/(2(1-\rho))$.

Now,

$$\begin{aligned}
(P-I)u(x) &= E_x u(W_1) - u(x) \\
&= E_x u([x+X_1]_+) - u(x) \\
&= E_x u(x+X_1) - u(x) + u(0)P(x+X_1 < 0) - E(u(x+X_1); x+X_1 < 0) \\
&= \frac{\lambda}{2(1-\rho)}\big(EX_1^2 + 2xEX_1 - E((x+X_1)^2; x+X_1 < 0)\big).
\end{aligned}$$

Thus, we find that the estimator α_n' is given by

$$\begin{aligned}
\alpha_n' &= \frac{1}{n}\sum_{k=0}^{n-1} W_k + (P-I)u(W_1) \\
&= \frac{\lambda EX_1^2}{2(1-\rho)} - \frac{\lambda}{2n(1-\rho)}\sum_{k=0}^{n-1} h_2(W_k) \\
&= \frac{EX_1^2}{-2EX_1} + \frac{1}{2nEX_1}\sum_{k=0}^{n-1} h_2(W_k),
\end{aligned}$$

where $h_2(x) = h_2(x;\rho) = E((x+X_1(\rho))^2; x+X_1(\rho) < 0)$.

Observe that α_n' consists of two components, the first of which is the standard heavy traffic approximation for EW_∞. The second component represents a correction to this quantity. According to the discussion before Proposition 7, α_n' will be a consistent estimator of EW_∞ if u is π integrable. This follows if W_∞ has a finite second moment, so that α_n' is consistent if $\rho < 1$ and $EV_0^3 < \infty$. Our next result compares the behaviour of the martingale estimator with the standard estimator in heavy traffic.

THEOREM 10. *Suppose that $\rho < 1$ and that both $\overline{U}_1$ and $\overline{V}_0$ possess a moment generating function that is finite in a neighbourhood of the origin. Then, $\alpha_n \to EW_\infty$ a.s.,*

$$n^{1/2}(\alpha_n - EW_\infty) \Rightarrow \sigma_1 N(0,1),$$

as $n \to \infty$, and the variance constant $\sigma_1^2 = \sigma_1^2(\rho) = O((1-\rho)^{-4})$.

Under the same conditions, $\alpha'_n \to EW_\infty$ a.s.,

$$n^{1/2}(\alpha'_n - EW_\infty) \Rightarrow \sigma_2 N(0,1)$$

as $n \to \infty$, and the variance constant $\sigma_2^2 = \sigma_2^2(\rho) = O((1-\rho)^{-2})$ as $\rho \to 1$.

The constant σ^2 appearing in these CLTs is known as the *time average variance constant* (TAVC) and allows one to compare the performance of the estimators, since, for example, confidence intervals based on these CLTs will have widths that are proportional to the square root of the TAVC.

REMARK 2. Theorem 10 establishes that as $\rho \to 1$, the TAVC for the standard estimator grows at rate $(1-\rho)^{-4}$, whereas that of the martingale estimator grows at rate $(1-\rho)^{-2}$. Thus, we see that the martingale estimator outperforms the standard estimator in heavy traffic. This observation is also exhibited through numerical examples in Henderson (1997).

REMARK 3. Theorem 10 is generalized to higher order moments of the steady-state waiting time in Henderson (1997). In particular, it is shown there that when estimating EW_∞^k, the TAVC of the standard estimator is typically $O((1-\rho)^{-2k-2})$, while that of the martingale estimator is $O((1-\rho)^{-2k})$.

REMARK 4. It is also shown in Henderson (1997) that the Minh-Sorli estimator of EW_∞^k typically has a TAVC that is $O((1-\rho)^{-2k+1})$. Thus, in heavy traffic, we can expect the Minh-Sorli estimator to outperform both the standard and martingale estimators.

REMARK 5. The moment generating function assumption in Theorem 10 can be weakened to an assumption of finite moments of high-enough order. See Henderson (1997) for details.

9. Extensions/issues. In this paper, we have shown how to obtain variance reduction in simulations of a large class of performance measures for Markov processes. The basic approach was to add an appropriately defined zero mean martingale to a given estimator. Since the martingale has zero expectation, it does not introduce bias, and if the martingale is chosen carefully, one can expect large variance reductions.

The method can be applied in ways other than those we have presented here. For example, one may "switch on and off" the martingale. (Consider a single-server queue in which the server works at a load-dependent rate.) When the amount of work in the system is large (there are more than N customers in the system), the server works at maximum efficiency, say μ, but when the amount of work is below N, the server is less efficient. Suppose we wish to estimate the mean steady-state number of customers in the system α. (Of course, this can be computed using the theory of birth-death processes, but the problem serves as an illustrative example.) A reasonable estimator of α is given by

$$\frac{1}{t}\int_0^t X(s)\,ds,$$

where $X(s)$ is the number of customers in the system at time s.

One might approximate this system by an M/M/1 queue with constant service rate μ. We would solve Poisson's equation for the M/M/1 system to obtain u. We might then use the estimator α'_t as discussed in §6. Notice, however, that our approximation is only considered to be reasonable when $X(t) > N$. We could instead use an approximating martingale estimator that is only "on" when $X(t) > N$. Let T_i denote the ith time t at which $X(t) > N$, with $X(t-) = N$, and let $U_i > T_i$ be the first time after time T_i at which $X(t) = N$, with $X(t-) > N$. In this case, the estimator might be of the form,

$$\frac{1}{t}\int_0^t X(s)\,ds + \frac{1}{t}\sum_{k=1}^{N(t)}\left\{u(X(U_k)) - u(X(T_k)) - \int_{T_k}^{U_k} Au(X(s))\,ds\right\},$$

where A is the generator of the process $X = (X(s): s \geq 0)$ and $N(t)$ is the number of time intervals $(T_i, U_i]$ that are contained in the time interval $[0, t]$. Under moderate conditions, this estimator is consistent and can be expected to yield variance reduction.

One might also consider using different approximations at different times. In the above example, one might take the solution v to Poisson's equation for an M/M/1 system with a "slow" service rate and apply it when $X(t) \leq N$. The estimator will then have a similar form to that given above.

Our discussion throughout has focused on Markov processes. Of course, one must almost invariably attach "supplementary variables" to the state space of simulated processes to make them Markov. In Henderson and Glynn (2001), we consider this issue in more depth. We show that a direct application of the ideas in this paper is certainly possible, but that one must explicitly take into account the supplementary variables to obtain the greatest benefit.

Acknowledgments. The work of the first author was partially supported by New Zealand Public Good Science Fund Grant UOA803 and National Science Foundation Grant DMI-0085165. The work of the second author was partially supported by Army Research Office Contract No. DAAG55-97-1-0377-P0001 and National Science Foundation Grant DMS-9704732-001.

Appendix: Proofs.

PROOF OF PROPOSITION 1. It is well known that M is a local martingale, so that the only condition that needs to be verified is that $E_x|M_n| < \infty$ for all n. If $x \in C$, then $M_n = 0$ for all n, so suppose that $x \in C^c$. If $y \in C$, then $(P-I)u(y) = 0$, and so, defining $a \wedge b = \min\{a, b\}$, we have that

$$\begin{aligned}
E_x|M_n| &\leq E_x u(X_n) I(n < T) + u(x) + E_x \sum_{k=0}^{n\wedge T-1} |(P-I)u(X_k)| \\
&\leq B^n u(x) + u(x) + nb \\
&\leq b + B^{n-1}u(x) + u(x) + nb \\
&\leq 2nb + 2u(x) < \infty,
\end{aligned}$$

where the third and fourth inequalities follow from the fact that for $x \in C^c$, $|Bu(x)| = |Pu(x)| \leq u(x) + b$. To show that $E_x M_T = 0$ under the condition given, observe that

$$\begin{aligned}
E_x(|M_{k+1} - M_k| \mid \mathcal{F}_k) &= E_x(|u(X_{k+1}) - Pu(X_k)| \mid X_k) \\
&\leq E_x(u(X_{k+1})|X_k) + Pu(X_k) \\
&= 2Pu(X_k) \\
&\leq 2(u(X_k) + b).
\end{aligned}$$

Hence, $E_x\varphi < \infty$, where

$$\varphi = \sum_{k=1}^{T} |M_{k+1} - M_k|.$$

But $M_{T\wedge m} \to M_T$ a.s. as $m \to \infty$, $\sup_m |M_{T\wedge m}| \le \varphi$, and $E_x M_{T\wedge m} = 0$ for all m, and so by dominated convergence, $E_x M_T = 0$.

To prove that u^* is the optimal choice of u, note that $(P - I)u^* = -f$, and so under P_x,

$$M_T = -u^*(x) + \sum_{k=0}^{T-1} f(X_k).$$

Hence, $E_x M_T = 0$, and $\alpha_n' = u^*(x)$. □

Proof of Proposition 2. Proposition 1.2 in Henderson (1997, p. 15) shows that $Au(x)$ is the bounded pointwise limit

$$\lim_{t\to 0^+} \frac{E_x u(Y(t)) - u(x)}{t}.$$

Lemma 3.4 of Kurtz (1969), then establishes that M is a P_μ martingale for all μ.

To complete the proof, note that

$$|M(t \wedge T)| \le 2\|u\| + \|Au\|T.$$

However, $M(t\wedge T) \to M(T)$ P_μ a.s. and so, by dominated convergence, $E_\mu M(T) = 0$. □

Proof of Proposition 3. Let b be a bound on $|u(x)|$. Then, $M_n \le 2b/(1-\delta)$ for all n, and is thus integrable. Now, $u(X_k) - Pu(X_{k-1})$ are martingale differences, and $(\prod_{k=0}^{n-1} h(X_k) : n \ge 1)$ is previsible with respect to $(\mathscr{F}_n : n \ge 1)$, so that M_n is a discrete analogue of a stochastic integral, and therefore a martingale.

The martingale convergence theorem gives $M_n \to M_\infty$ a.s. and $EM_\infty = 0$ follows by bounded convergence.

Finally, observe that if $u(x) = u^*(x)$, then

$$\begin{aligned}
M_\infty &= \sum_{k=1}^{\infty} [u(X_k) - Pu(X_{k-1})] \prod_{j=0}^{k-1} h(X_j) \\
&= -h(X_0)Pu(X_0) + \sum_{k=1}^{\infty} [u(X_k) - h(X_k)Pu(X_k)] \prod_{j=0}^{k-1} h(X_j) \\
&= f(X_0) - u(X_0) + \sum_{k=1}^{\infty} f(X_k) \prod_{j=0}^{k-1} h(X_j) \\
&= U - u^*(X_0). \quad \square
\end{aligned}$$

Proof of Proposition 5. It is clear that M_j as defined is a sum of martingale differences. The boundedness of the functions u_j ensures that M_j is integrable. Finally, use the fact that the u_j^*s satisfy the linear system (5) to show that $M_n = f(X_n) - u_n^*(X_0)$. □

Proof of Proposition 6. For $0 \le s \le t$, define $v(s, x) = u(t - s, x)$. Theorem 12.6 of Karlin and Taylor (1981, p. 325) gives that $\widetilde{M} = (\widetilde{M}(s)\colon 0 \le s \le t)$ is a P_μ martingale for all μ, where

$$\widetilde{M}(s) = v(s, Y(s)) - v(0, Y(0)) - \int_0^s [Av(r, Y(r)) + v_1(r, Y(r))]\, dr.$$

However, note that $v_1(r, Y(r)) = -u_1(t - r, Y(r))$, so that

$$\widetilde{M}(s) = u(t-s, Y(s)) - u(t, Y(0)) - \int_0^s [Au(t-r, Y(r)) - u_1(t-r, Y(r))]\, dr = M(s).$$

Thus, we conclude that M is a P_μ martingale.

If $u = u^*$, then $M(t) = u^*(0, Y(t)) - u^*(t, Y(0)) = f(Y(t)) - u^*(t, Y(0))$, and the final statement in the proposition follows. □

PROOF OF PROPOSITION 7. It is well known that M is a local martingale. It is easy to show that M_n is P_π integrable, so that the result follows. □

PROOF OF PROPOSITION 9. Consider the process $W = (W_s\colon 0 \le s \le t)$, where $W_s = (t - s, X_s)$. Applying Itô's formula to W and the function u, we obtain (after some algebra)

$$u(t-s, X_s) - u(t, X_0) - \int_0^t [Au(t-s, X_s) - u_1(t-s, X_s)]\, ds$$
$$= \int_0^t u_2(t-s, X_s)\sigma(X_s)\, dB_s.$$

However, under our assumptions, the right-hand side of this expression is a martingale, and thus M is a martingale. □

PROOF OF THEOREM 10. The consistency of both estimators follows from the strong law for Harris chains and the result of Kiefer and Wolfowitz (1956).

The result for α_n follows from Theorem 5.1 of Asmussen (1992). We will prove the corresponding results for α_n' using regenerative process theory. Define $C = C(\rho) = \inf\{n \ge 1\colon W_n = 0\}$. If $\rho < 1$, then $C < \infty$ a.s. Furthermore, Lemma 4.1 of Asmussen (1992) establishes that $(1-\rho)EC(\rho)$ is bounded away from 0 and ∞ under our conditions. He also shows (Corollary 5.1) that $EC(\rho)^2 = O((1-\rho)^{-3})$ as $\rho \to 1$. Hence, since the function h_2 is bounded by $h_2(0)$, it follows from regenerative arguments that the CLT for α_n' holds with

$$\sigma_2^2 = \frac{E\big(\sum_{k=0}^{C-1}[a_1 + a_2 h_2(W_k) - E_\pi W_0]\big)^2}{EC}$$
$$= \frac{a_2^2}{EC} E\left(\sum_{k=0}^{C-1}[h_2(W_k) - \beta]\right)^2,$$

where $a_1 = EX_1^2/(-2EX_1)$, $a_2 = (2EX_1)^{-1}$, $\beta = \pi h_2$, and π is the stationary probability distribution. However, $a_2^2/EC = O((1-\rho)^{-1})$, so the proof will be complete if we show that $E(\sum_{k=0}^{C-1} h_2(W_k) - \beta)^2 = O((1-\rho)^{-1})$. Observe that

$$E\left(\sum_{k=0}^{C-1}[h_2(W_k) - \beta]\right)^2 \le 2E\left(\sum_{k=0}^{C-1} h_2(W_k)\right)^2 + 2\beta^2 EC^2.$$

From Lemma 12, $\beta = \beta(\rho) = O(1-\rho)$, so that the second term is $O((1-\rho)^{-1})$. Lemma 12 also provides the corresponding result for the first term, so that the proof is complete. □

LEMMA 11. *If $\overline{U}_1$ has a finite moment generating function in a neighbourhood of the origin, then there exists a $\gamma > 0$ such that the following results hold*:

1. *As* $\rho \to 1$, $Ee^{-\gamma W_\infty(\rho)} = O(1-\rho)$.
2. *For any* $0 < \rho_0 < 1$, *there exists* $d = d(\rho_0) < \infty$, *such that*

$$\sup_{\rho_0 \le \rho \le 1} h_2(x; \rho) \le de^{-\gamma x}.$$

PROOF. The first statement in the lemma asserts that the Laplace Stieltjes transform (LST) $Ee^{-\gamma W_\infty(\rho)}$ of the stationary waiting time $W_\infty(\rho)$ evaluated at γ is $O(1-\rho)$. Marshall (1968) showed that

$$(1 - \tilde{k}_\rho(s))\tilde{w}_\rho(s) = \frac{1 - \tilde{h}_\rho(-s)}{EC},$$

where $\tilde{k}_\rho$, $\tilde{w}_\rho$ and $\tilde{h}_\rho$ are the LSTs of the increment r.v. $X_1(\rho)$, the stationary waiting time, and the idle time in the ρth system, respectively. The assumption that the interarrival times $\overline{U}_n/\rho$ satisfy $P(\overline{U}_n/\rho > x) \le ce^{-\gamma'\rho x}$ for some $c, \gamma' > 0$ ensures that $X_1(\rho)$ has an exponentially bounded left tail also, and thus $\tilde{k}_\rho(s)$ is finite for $0 \le s < \gamma'$. Furthermore, a similar argument to that used in the proof of Lemma 4.1 in Asmussen (1992) shows that $\tilde{h}_\rho(-s)$ is bounded for $0 \le s < \gamma'$. Now, for small enough $\gamma_1 \in (0, \gamma')$, $\tilde{k}_\rho(\gamma_1)$ is bounded away from 1 as $\rho \to 1$. (This follows since $\tilde{k}_\rho(0) = 1$, $\tilde{k}'_\rho(0) = -EX_1(\rho) \ge 0$ for all ρ, and $\tilde{k}''_\rho(0) = EX_1(\rho)^2$, which is bounded away from 0 as $\rho \to 1$.) We then find that

$$\begin{aligned} Ee^{-\gamma_1 W_\infty(\rho)} &= \tilde{w}_\rho(\gamma_1) \\ &= \frac{1-\tilde{h}_\rho(-\gamma_1)}{EC(1-\tilde{k}_\rho(\gamma_1))}, \end{aligned}$$

and since $(1-\rho)EC$ is bounded away from 0 and ∞, $Ee^{-\gamma_1 W_\infty(\rho)} = O(1-\rho)$ for all γ_1 sufficiently small.

The second result of the lemma follows by a direct calculation as follows. Let $x \ge 0$ and suppose that $\rho \ge \rho_0 > 0$. Using integration by parts,

$$\begin{aligned} |h_2(x)| &= \int_{-\infty}^{-x} (x+y)^2 P(X_1(\rho) \in dy) \\ &= \big[(x+y)^2 P(X_1(\rho) < y)\big]_{-\infty}^{-x} - \int_{-\infty}^{-x} 2(x+y) P(X_1(\rho) < y)\, dy \\ &\le -2\int_{-\infty}^{-x} (x+y) c e^{\gamma'\rho y}\, dy \\ &= \frac{2ce^{-\gamma'\rho x}}{(\gamma'\rho)^2} \\ &\le de^{-\gamma_2 x}, \end{aligned}$$

where $d = 2c/(\gamma'\rho_0)^2$ and $\gamma_2 = \gamma'\rho_0$. Taking $\gamma = \min\{\gamma_1, \gamma_2\}$ yields the result. □

LEMMA 12. *If $\overline{U_1}$ has a finite moment generating function in a neighbourhood of the origin, then as $\rho \to 1$, $Eh_2(W_\infty, \rho) = O(1-\rho)$, and*

$$E\left(\sum_{i=0}^{C-1} h_2(W_i)\right)^2 = O((1-\rho)^{-1}).$$

PROOF. From Lemma 11,

$$|Eh_2(W_\infty(\rho))| \le dEe^{-\gamma W_\infty(\rho)},$$

so that the first result follows. For the second result, note that

$$E\left(\sum_{i=0}^{C-1} h_2(W_i)\right)^2 = 2E\left(\sum_{n=0}^{C-1} h_2(W_n) \sum_{m=n}^{C-1} h_2(W_m)\right) - E\left(\sum_{i=0}^{C-1} h_2^2(W_i)\right). \tag{A1}$$

(The representation (A1) has been fruitfully exploited previously; see Asmussen 1992.) The second term on the right-hand side of (A1) is bounded by

$$E\sum_{i=0}^{C-1} d^2 e^{-2\gamma W_n} \le d^2 EC,$$

which is $O((1-\rho)^{-1})$ as $\rho \to 1$.

The first term in (A1) is bounded by

$$(A2) \qquad 2d^2 E\left(\sum_{n=0}^{C-1} e^{-\gamma W_n} E\left[\sum_{m=n}^{C-1} e^{-\gamma W_m} \,\middle|\, W_n, C > n\right]\right).$$

Turning to the conditional expectation in (A2), note that

$$E\left[\sum_{m=n}^{C-1} e^{-\gamma W_m} \,\middle|\, W_n, C > n\right] \le E_{W_n} C.$$

Suppose that $W_0 = x$ and let us determine $E_x C$. Note that $x + I = -\sum_{i=1}^{C} X_i$, where I is the length of the first idle period. Wald's identity then gives $x + E_x I = -EXE_x C$. Lemma 4.1 of Asmussen (1992) shows that there is a constant c such that $E_x I \le c$ for all $x \ge 0$ and for all ρ. Thus $E_{W_n} C \le (W_n + c)/(-EX)$, and so (A2) is bounded by

$$\frac{2d^2}{-EX} E \sum_{n=0}^{C-1} e^{-\gamma W_n}(W_n + c).$$

However, $(x+c)e^{-\gamma x} \le e^{-\gamma x/2}$ for x larger than some constant, say, b_1. Thus, (A2) is bounded by

$$(A3) \qquad \frac{b_2}{-EX} E\left(\sum_{n=0}^{C-1}\left[e^{-\gamma W_n/2} + b_3 I(W_n \le b_1)\right]\right)$$
$$= \frac{1}{1-\rho}\left(b_4 ECEe^{-\gamma W_\infty/2} + b_5 ECP(W_\infty \le b_1)\right),$$

for bounded deterministic constants b_i $(1 \le i \le 5)$.

The second result now follows by noting that $Ee^{-\gamma W_\infty/2} = O(1-\rho)$, $EC = O((1-\rho)^{-1})$, and

$$\begin{aligned} P(W_\infty \le b_1) &= P(e^{-\gamma W_\infty} \ge e^{-\gamma b_1}) \\ &\le e^{\gamma b_1} Ee^{-\gamma W_\infty} \\ &= O(1-\rho). \quad \square \end{aligned}$$

References

Andradóttir, S., D. P. Heyman, T. J. Ott. 1993. Variance reduction through smoothing and control variates for Markov chain simulations. *ACM Trans. Model. Comput. Simulation* **3** 167–189.

Asmussen, S. 1987. *Applied Probability and Queues.* Wiley, New York, 79, 154, 181.

———. 1990. Exponential families and regression in the Monte Carlo study of queues and random walks. *Ann. Statist.* **18** 1851–1867.

———. 1992. Queueing simulation in heavy traffic. *Math. Oper. Res.* **17** 84–111.

Bratley, P., B. L. Fox, L. E. Schrage. 1983. *A Guide to Simulation.* Springer-Verlag, New York.

Chung, K. L., R. J. Williams. 1990. *Introduction to Stochastic Integration*, 2nd ed. Birkhäuser, Boston, MA.

Emsermann, M., B. Simon. 2000. Improving simulation efficiency with quasi control variates. *Stochastic Models*. Forthcoming.

Fox, B. L., P. W. Glynn. 1989. Simulating discounted costs. *Management Sci.* **35** 1297–1315.

Glynn, P. W., S. P. Meyn. 1996. A Liapounov bound for solutions of the Poisson equation. *Ann. Probab.* **24** 916–931.

Harrison, J. M. 1990. *Brownian Motion and Stochastic Flow Systems*, 2nd ed. Krieger, Malabar, FL, 21, 94.

Henderson, S. G. 1997. Variance reduction via an approximating Markov process. Ph.D. thesis, Department of Operations Research, Stanford University, Stanford, CA.

———, P. W. Glynn. 2001. Approximating martingales for variance reduction in general discrete-event simulation. Working paper.

———, S. P. Meyn. 1997. Efficient simulation of multiclass queueing networks. S. Andradottir, K. J. Healy, D. H. Withers, B. L. Nelson, eds. *Proc. 1997 Winter Simulation Conf.*, 216–223.

———, ———. 2000. Variance reduction for simulation in multiclass queueing networks. Forthcoming.

Karlin, S., H. M. Taylor. 1981. *A Second Course in Stochastic Processes.* Academic Press, Boston, MA, 214.

Kiefer, J., J. Wolfowitz. (1956). On the characteristics of the general queueing process, with applications to random walk. *Ann. Math. Statist.* **27** 147–161.

Kurtz, T. G. 1969. Extensions of Trotter's operator semigroup approximation theorems. *J. Funct. Anal.* **3** 354–375.

Marshall, K. T. 1968. Some relationships between the distributions of waiting time, idle time, and interoutput time in the GI/G/1 queue. *Siam J. Appl. Math.* **16** 324–327.

Meyn, S. P., R. L. Tweedie. 1993. *Markov Chains and Stochastic Stability.* Springer-Verlag, Berlin, Germany, 265.

Minh, D. L., R. M. Sorli. (1983). Simulating the GI/G/1 queue in heavy traffic. *Oper. Res.* **31** 966–971.

Nelson, B. L., B. W. Schmeiser, M. R. Taaffe, J. Wang. 1997. Approximation-assisted point estimation. *Oper. Res. Lett.* **20** 109–118.

Schmeiser, B. W., M. R. Taaffe. 1994. Time-dependent queueing network approximations as simulation external control variates. *Oper. Res. Lett.* **16** 1–9.

———, ———, J. Wang. 2001. Biased control-variate estimation. *IIE Trans.* **33** 219–228.

S. G. Henderson: School of Operations Research and Industrial Engineering, Cornell University, Ithaca, New York 14853; e-mail: henderson@orie.cornell.edu

P. W. Glynn: Management Science and Engineering, Stanford University, Stanford, California 94305-4026; e-mail: GLYNN@stanford.edu