Proceedings of the 2008 Winter Simulation Conference
S. J. Mason, R. R. Hill, L. Moench, O. Rose, eds.

APPROXIMATE ZERO-VARIANCE SIMULATION

Pierre L'Ecuyer

DIRO, Université de Montreal
C.P. 6128, Succ. Centre-Ville
Montréal (Québec), H3C 3J7, CANADA

Bruno Tuffin

INRIA Rennes Bretagne Atlantique
Campus de Beaulieu
35042 Rennes Cedex, FRANCE

ABSTRACT

Monte Carlo simulation applies to a wide range of estimation problems, but converges rather slowly in general. Variance reduction techniques can lower the estimation error, sometimes by a large factor, but rarely change the convergence rate of the estimation error. This error usually decreases as the inverse square root of the computational effort, as dictated by the central limit theorem. In theory, there exist simulation estimators with zero variance, i.e., that always provide the exact value. The catch is that these estimators are usually much too difficult (or virtually impossible) to implement. However, there are situations, especially in the context of rare-event simulation, where the zero-variance simulation can be approximated well enough to provide huge efficiency gains. Adaptive versions can even yield a faster convergence rate, including exponential convergence in some cases. This paper gives a brief overview of these methods and discuss their practicality.

1 INTRODUCTION

1.1 Monte Carlo Estimation

Monte Carlo (MC) is the method of choice to estimate a mathematical expectation defined as a high-dimensional integral (Fishman 1996, Asmussen and Glynn 2007). To recall how it works, suppose we want to estimate $\mu_0 = \mathbb{E}[X]$ for some random variable X. In its basic form, MC would generate n independent copies of X, say $X_1, \ldots, X_n$, and estimate μ_0 by the sample average

$$\bar{X}_n = \frac{1}{n}\sum_{i=1}^{n} X_i.$$

The strong law of large numbers guarantees the almost sure convergence $\bar{X}_n \to \mu_0$ when $n \to \infty$, and the accuracy of the MC method is described by the central limit theorem, which tells us that the interval

$$\left(\bar{X}_n - \frac{z_{1-\alpha/2}S_n}{\sqrt{n}},\ \bar{X}_n + \frac{z_{1-\alpha/2}S_n}{\sqrt{n}}\right)$$

is an asymptotically valid confidence interval for μ_0 with confidence level $1-\alpha$, where $z_{1-\alpha/2}$ is the $(1-\alpha/2)$-quantile of the standard normal distribution; that is, $\Phi(z_{1-\alpha/2}) = 1-\alpha/2$ where Φ is the standard normal distribution function, and

$$S_n^2 = \frac{1}{n-1}\sum_{i=1}^{n}(X_i - \bar{X}_n)^2$$

is the sample variance, which is an unbiased estimator of $\sigma^2 = \mathrm{Var}[X]$, assumed to be finite. The half-width of the confidence interval represents (in a probabilistic way) the estimation error. It converges to zero as $n^{-1/2}$ when $n \to \infty$.

1.2 Markov Chain Model

A finite-horizon discrete-event simulation model can in general be viewed as a *discrete-time Markov chain* (DTMC) $\{Y_j, j \geq 0\}$ that evolves in some general state space $\mathscr{Y}$ until reaching a stopping time τ. By putting enough information in the state Y_j, we can obtain that $\{Y_j, j \geq 0\}$ is indeed a Markov chain and also that τ is defined as the first time when the chain hits a given set of states, Δ. That is, $\tau = \inf\{j \geq 0 : Y_j \in \Delta\}$. The Markov chain is assumed to have a transition kernel $\{\mathbb{P}[\cdot \mid y], y \in \mathscr{Y}\}$ and initial distribution π_0, where $\mathbb{P}[B \mid y] = \mathbb{P}[Y_j \in B \mid Y_{j-1} = y]$ and $\pi_0(B) = \mathbb{P}[Y_0 \in B]$ for any measurable $B \subset \mathscr{Y}$. Suppose that each time we move from state $y \in \mathscr{Y}$ to state $z \in \mathscr{Y}$, we incur a non-negative cost $c(y,z)$. The total cost is then the random variable

$$X = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j).$$

Let $\mu(y) = \mathbb{E}_y[X]$ denote the expected value of X when $Y_0 = y$. We want to estimate $\mu_0 = \mathbb{E}[X] = \int_{y\in\mathscr{Y}} \mu(y)\pi_0(dy)$, the expected value of X for the initial distribution π_0.

We will exploit the fact that the function $\mu : \mathscr{Y} \to \mathbb{R}$ satisfies the recurrence equation (or Poisson equation)

$$\begin{aligned} \mu(y) &= \mathbb{E}_y[c(y,Y_1) + \mu(Y_1)] \\ &= \int_{\mathscr{Y}} [c(y,z) + \mu(z)]\, d\mathbb{P}[dz \mid y] \end{aligned} \tag{1}$$

for $y \notin \Delta$, and $\mu(y) = 0$ for $y \in \Delta$. This is a simplified form of the standard dynamic programming recurrence equation, without the minimization.

1.3 Variance Reduction

The usual way of improving accuracy is to change the estimator X to another estimator with the same expectation and smaller variance σ^2. Note that the bias (if any) and the computational cost are also important factors for the quality of an estimator. The computational cost is usually taken into account via a notion of efficiency, or work-normalized variance (Glynn and Whitt 1992, Glasserman 2004, Asmussen and Glynn 2007). In this paper, we concentrate on the variance only, and we assume that X has no bias.

We focus primarily on two families of variance reduction techniques, namely importance sampling and control variates. For each of them, we explain how zero-variance can be reached in principle, and we discuss situations where large gains can be obtained by approximating a zero-variance sampling scheme, in the DTMC setting discussed earlier. With adaptive schemes, it is possible to beat the $O(n^{-1/2})$ convergence rate implied by the central limit theorem. There are even cases where one can obtain convergence at an exponential rate. We also discuss how convergence at a faster rate than $O(n^{-1/2})$ can be achieved via generalized antithetic variates, or equivalently randomized quasi-Monte Carlo (**?**).

1.4 Rare Events

In the context of rare-event simulation, the system is often parameterized by a rarity parameter ε so that the important rare events become rarer (and accurate estimation becomes more difficult under naive simulation) when $\varepsilon \to 0$. In this type of setting, asymptotic approximations of μ_0, valid when $\varepsilon \to 0$, can often be used to approximate a zero-variance sampling scheme. The quality and robustness of the estimators thus obtained can be characterized by asymptotic properties such as logarithmic efficiency (also called asymptotic optimality), bounded relative error (or relative variance), and vanishing relative error (**?**). We will briefly survey these properties and give examples.

1.5 Outline

The remainder of the paper is organized as follows. Some of the main asymptotic robustness properties of estimators in a rare event setting are summarized in Section 2. In Section 3, we recall how to simulate a Markov chain with importance sampling, define a zero-variance estimator in that context, and give examples where this zero-variance estimator can be well approximated. In Section 4, we provide a similar discussion on how zero-variance can be approximated via control variates. Section 5 discusses adaptive learning algorithms to approach the zero-variance estimator, for the two families of variance reduction techniques discussed earlier. Difficulties and potential remedies are highlighted. Generalized antithetic variates are discussed in Section 6. Section 7 summarizes the important facts concerning zero-variance estimation and discusses some directions of research we believe are worthy to push forward.

2 ASYMPTOTICS FOR RARE-EVENT SIMULATION

In rare-event simulation, asymptotic analysis of estimators is usually done under the following type of framework. The quantity $\mu_0 = \mu_0(\varepsilon) > 0$ to be estimated is parameterized by a rarity parameter $\varepsilon > 0$, in a way that $\lim_{\varepsilon\to 0^+} \mu_0(\varepsilon) = 0$. We have a family of estimators $X = X(\varepsilon)$ taking their values in $[0,\infty)$, such that $\mathbb{E}[X(\varepsilon)] = \mu_0(\varepsilon)$ and $\mathrm{Var}[X(\varepsilon)] = \sigma_0^2(\varepsilon)$ for each $\varepsilon > 0$. For a given model, different parameterizations may specify different asymptotic regimes (Heidelberger 1995, Juneja and Shahabuddin 2006). For example, in a queueing system with buffer size B and s servers, we can take $\varepsilon = 1/B$ if we are interested in very large values of B, and $\varepsilon = 1/s$ if we are interested in what happens when there is a large number of servers. In reliability models, some of the failure rates are often taken as polynomial functions of ε (Shahabuddin 1994, Nakayama 1996).

In this type of setting, we would like to be able to estimate $\mu_0(\varepsilon)$ to a given relative accuracy with a computing budget that does not increase with ε. This corresponds to the notion of an estimator with *bounded relative error* (BRE), defined below. There are situations, however, where the best available estimators do not have BRE, but require a computing budget that increases only very slowly with ε for a given relative accuracy. These estimators often enjoy the slightly weaker property of *logarithmic efficiency* (LE), also known as *asymptotic optimality*. Importance sampling estimators constructed via the theory of large deviations often have this property (Heidelberger 1995, Juneja and Shahabuddin 2006, Asmussen and Glynn 2007). A third property, also defined below, is that of *vanishing relative error* (VRE). It means that the relative error converges to zero when $\varepsilon \to 0$, for a given computing budget.

Definition 1 (a) The estimator $X(\varepsilon)$ has *bounded relative error* (BRE) (Heidelberger 1995), or equivalently *bounded relative variance*, if

$$\limsup_{\varepsilon\to 0} \frac{\sigma_0(\varepsilon)}{\mu_0(\varepsilon)} < \infty. \tag{2}$$

(b) It is *logarithmically efficient* (LE) (Heidelberger 1995, Asmussen and Glynn 2007) if

$$\lim_{\varepsilon\to 0} \frac{\ln \mathbb{E}[X^2(\varepsilon)]}{2\ln \mu_0(\varepsilon)} = 1. \tag{3}$$

(c) It has *vanishing relative error* (VRE) (**?**) if

$$\limsup_{\varepsilon\to 0} \frac{\sigma_0(\varepsilon)}{\mu_0(\varepsilon)} = 0, \tag{4}$$

or equivalently if

$$\limsup_{\varepsilon\to 0} \mathbb{E}[X^2(\varepsilon)]/\mu_0^2(\varepsilon) = 1. \tag{5}$$

The VRE property is obviously stronger than BRE, which is itself stronger than LE. VRE is strongly related to zero-variance simulation in the sense that VRE implies that $X(\varepsilon)$ is generated from a probability law that converges in the L_∞ norm to a zero-variance measure (**?**). We will see examples where this property holds.

Dupuis and Wang (2004) give a Lyapunov (sufficient) condition for LE, interpreting the Lyapunov function as a *subsolution* to the recurrence equation of a stochastic game in which we select a change of measure for importance sampling, and then an adversary selects a set of sample paths with the worst possible variance contribution. Other sets of Lyapunov conditions for the BRE, LE, and VRE properties can be found in Blanchet and Glynn (2007) and L'Ecuyer et al. (2008). These conditions are often convenient to verify the BRE, LE, and VRE properties, and can also be used to design IS schemes that satisfy these properties.

3 ZERO VARIANCE VIA IMPORTANCE SAMPLING

Suppose that the estimator X of interest can be written as $X = h(\omega)$ for some (measurable) function $h : \Omega \to \mathbb{R}$, where ω obeys a probability measure $\mathbb{P}$ defined over a measurable space $(\Omega, \mathscr{F})$, where Ω is the sample space and $\mathscr{F}$ is a σ-field. In *importance sampling* (IS), we replace $\mathbb{P}$ by another probability measure $\tilde{\mathbb{P}}$ over the same measurable space and for which $d\tilde{\mathbb{P}}(\omega) \neq 0$ whenever $h(\omega)d\mathbb{P}(\omega) \neq 0$. Under this condition, we have

$$\begin{aligned}\mathbb{E}[X] &= \int h(\omega) d\mathbb{P}(\omega) \\ &= \int h(\omega) \frac{d\mathbb{P}(\omega)}{d\tilde{\mathbb{P}}(\omega)} d\tilde{\mathbb{P}}(\omega) \\ &= \tilde{\mathbb{E}}\left[h(\omega)L(\omega)\right],\end{aligned}$$

where $L = d\mathbb{P}/d\tilde{\mathbb{P}}$ is the likelihood ratio (or Radon-Nikodým derivative) between the two measures and $\tilde{\mathbb{E}}$ is the expectation associated with $\tilde{\mathbb{P}}$. Then we can generate n independent realizations of ω from $\tilde{\mathbb{P}}$, say $\omega_1, \ldots, \omega_n$, and use the average

$$\bar{X}_{\mathrm{is},n} = \frac{1}{n}\sum_{i=1}^{n} h(\omega_i)L(\omega_i)$$

as an unbiased estimator of μ_0. The aim is to select $\tilde{\mathbb{P}}$ so that

$$\tilde{\mathrm{Var}}[h(\omega)L(\omega)] = \tilde{\mathbb{E}}[(h(\omega)L(\omega))^2] - \mu_0^2$$

is as small as possible, and preferably much smaller than $\mathrm{Var}[X] = \mathbb{E}[h(\omega)]$. If h is non-negative, one can check that $d\tilde{\mathbb{P}}(\omega) = \frac{h(\omega)}{\mathbb{E}[h(\omega)]} d\mathbb{P}(\omega)$ yields variance zero. In the case where $\mathbb{P}$ and $\tilde{\mathbb{P}}$ are (possibly multivariate) continuous probability distributions with densities with respect to the Lebesgue measure, L is the ratio of these densities. If they are discrete distributions, then L is the ratio of probabilities.

We shall concentrate on the case where the system is modeled by a DTMC $\{Y_j, j \geq 0\}$ as in Section 1.2. Here, ω represents the sample path $Y_0, Y_1, \ldots, Y_\tau$ and we have $X = h(Y_0, \ldots, Y_\tau) = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j)$. To estimate $\mu_0 = \mathbb{E}[X]$ by IS, we restrict ourselves to changes of measure that replace $\mathbb{P}$ and π_0 by a different kernel and (possibly) a modified initial distribution, so the model remains a Markov chain under IS. It turns out that a zero-variance change of measure can be found within this class.

In general, for a DTMC, IS consists in replacing the transition kernel $\{\mathbb{P}[\cdot \mid y], y \in \mathscr{Y}\}$ by a new kernel $\{\tilde{\mathbb{P}}[\cdot \mid y], y \in \mathscr{Y}\}$, and the initial distribution π_0 by $\tilde{\pi}_0$, chosen so that for all $B \in \mathscr{F}$, $\tilde{\mathbb{P}}[B \mid y] > 0$ whenever $\int_B \mu(y')\mathbb{P}[dy' \mid y] > 0$, $\tilde{\pi}[B] > 0$ whenever $\int_B \mu(y')\pi[dy'] > 0$, and $\tilde{\mathbb{E}}[\tau] < \infty$. We then have $\mu_0 = \tilde{\mathbb{E}}[X_{\mathrm{is}}]$ where

$$X_{\mathrm{is}} = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j) L(Y_0, \ldots, Y_j) \tag{6}$$

and

$$L(Y_0, \ldots, Y_j) = \frac{\partial \pi_0}{\partial \tilde{\pi}_0}(Y_0) \prod_{i=1}^{j} \frac{\partial \mathbb{P}}{\partial \tilde{\mathbb{P}}}(Y_i \mid Y_{i-1})$$

is the likelihood ratio associated with the choice of initial state and the first j transitions.

For more concreteness and simplicity, suppose that the state space $\mathscr{Y}$ is discrete, so the transition kernel is determined by a matrix of transition probabilities $p(y,z) = \mathbb{P}[Y_j = z \mid Y_{j-1} = y]$ for $x,z \in \mathscr{Y}$, and initial probabilities $\pi_0(y) = \mathbb{P}[Y_0 = y]$ for $y \in \mathscr{Y}$. We replace the probabilities $p(y,z)$ and $\pi_0(y)$ by new probabilities $\tilde{p}(y,z)$ and $\tilde{\pi}_0(y)$. Then,

$$\mathbb{P}[\tau = n,\ (Y_0,\ldots,Y_\tau) = (y_0,\ldots,y_n)] = \pi_0(y_0)\prod_{j=1}^{n} p(y_{j-1},y_j)$$

and the likelihood ratio for the initial state and the first j transitions becomes

$$L(Y_0,\ldots,Y_j) = \frac{\pi_0(Y_0)}{\tilde{\pi}_0(Y_0)}\prod_{i=1}^{j}\frac{p(Y_{i-1},Y_i)}{\tilde{p}(Y_{i-1},Y_i)}.$$

In the case of a continuous state space, the transition probabilities are replaced by densities and the likelihood ratio has the same form.

Suppose now that we change the probabilities so that

$$\tilde{p}(y,z) = p(y,z)[c(y,z)+\mu(z)]/\mu(y), \tag{7}$$
$$\tilde{\pi}(y) = \pi(y)\mu(y)/\mu_0, \tag{8}$$

where p represents a probability in the discrete case and a density in the continuous case. Then, one can verify that the IS estimator X_{is} in (6) equals μ_0 with probability one under $\tilde{\mathbb{P}}$, so it has zero variance (Booth 1987, Kuruganti and Strickland 1996, Kollman et al. 1999, L'Ecuyer and Tuffin 2007). Indeed, we have

$$\begin{aligned} X_{\text{is}} &= \sum_{j=1}^{\tau} c(Y_{j-1},Y_j)L(Y_0,\ldots,Y_j) \\ &= \sum_{j=1}^{\tau} c(Y_{j-1},Y_j)\frac{\pi_0(Y_0)}{\tilde{\pi}_0(Y_0)}\prod_{i=1}^{j}\frac{p(Y_{i-1},Y_i)}{\tilde{p}(Y_{i-1},Y_i)} \\ &= \sum_{j=1}^{\tau} c(Y_{j-1},Y_j)\frac{\mu_0}{\mu(Y_0)}\prod_{i=1}^{j}\frac{\mu(Y_{i-1})}{c(Y_{i-1},Y_i)+\mu(Y_i)} \\ &= \mu_0, \end{aligned}$$

where the last equality is obtained by induction on the value taken by τ, using the fact that $\mu(Y_\tau) = 0$. These zero-variance modified probabilities are also unique in the sense that no other probabilities $\tilde{\pi}(y)$ and $\tilde{p}(y,z)$ give a Markov chain whose IS estimator has zero variance.

Of course, we cannot implement (7) and (8) exactly, because the function $\mu : \mathscr{Y} \to \mathbb{R}$ is unknown, but we can approximate this zero-variance change of measure if we can replace μ in these equations by a good proxy (or good approximation) chosen so that it is not too difficult to generate variates from the corresponding IS distributions.

Example 1 We give a simple academic example where zero variance can be reached exactly. Of course, simulation is not needed in this case; this is just an illustrative example. Consider a DTMC with state-space $\mathscr{Y} = \{0,1,\ldots,B\}$, for which $p(y,y+1) = p_y$ and $p(y,y-1) = 1-p_y$, for $y = 1,\ldots,B-1$, and we take $\Delta = \{0,B\}$ (so the probabilities from 0 and from B are irrelevant). We want to compute the probability of reaching B before reaching 0, starting from a given y, i.e., $\mu(y) = \mathbb{P}[Y_\tau = B \mid Y_0 = y]$. This is a stylized version of a problem that arises in telecommunications, where one is interested in how frequently a router becomes congested. Here, the recurrence equation (1) becomes

$$\mu(y) = p_y\mu(y+1) + (1-p_y)\mu(y-1)$$

for $y = 1,\ldots,B-1$, with the boundary conditions $\mu(0) = 0$ and $\mu(B) = 1$. Assuming that we start from a fixed state $Y_0 = y_0$, the zero-variance change of measure is defined by

$$\tilde{p}_y = \frac{p_y\mu(y+1)}{\mu(y)}.$$

Because $\mu(0) = 0$, we see that $\tilde{p}_1 = 1$ and that no sample path will ever return to 0 under zero-variance IS.

When $p_y = p < 1$ for $y = 1,\ldots,B-1$, this problem is known as the gambler's ruin problem. Its solution is $\mu(y) = (1-\rho^{-y})/(1-\rho^{-B})$ if $\rho = p/(1-p) \neq 1/2$, and $\mu(y) = y/B$ if $\rho = 1/2$. This gives

$$\tilde{p}_y = \frac{1-\rho^{-y-1}}{1-\rho^{-y}}\ p.$$

Note that the new probabilities $\tilde{p}_y$ do not depend on B, but they depend on the distance from 0. Moreover, one can verify that $(p_{y-1}/\tilde{p}_{y-1})(1-p_y)(1-\tilde{p}_y) = 1$, which means that any cycle (subpath going from one state to the same state) in a sample path does not contribute to the likelihood ratio. In other words, all paths leading to B are essentially the same path and have the same likelihood ratio. If $p < 1/2$, we have $\rho < 1$ and $\tilde{p}_y > 1/2$ for all y, so the system is attracted toward B under IS. If $p > 1/2$, we also have $\tilde{p}_y > p > 1/2$.

For large B, we can use the approximation $\mu(y) = (\rho^{B-y}-\rho^B)/(1-\rho^B) \approx \rho^{B-y}$, which corresponds to the probability of the sample path that goes directly from y to B. This path is the *dominating* path when ρ is small, in the sense that it has a much larger probability than the other paths leading to B, and it is easy to see that using this approximation in (7) leads to an estimator with VRE for $p \to 0$, or $B \to \infty$ with $p < 1/2$, because then $\rho^{B-y}/\mu(y) \to 1$. □

In the previous example, there is essentially a single path leading to B. For more complex Markov chains, the number

of distinct paths to the target set can be much larger, and this can greatly complicate the computation of the zero-variance IS. Yet, good approximations are sometimes possible, based on asymptotic analysis, as shown by the next example, taken from **?**.

Example 2 Suppose that our Markov chain has finite state space $\mathscr{Y} = \{0,\dots,M-1\}$, a fixed initial state y_0, and that we want to estimate the probability μ_0 that the chain hits a given set $B \subset \mathscr{Y}$ before returning to y_0. We have $\Delta = B \cup \{y_0\}$. Suppose that the transition probabilities are parameterized as

$$p(y,z) = p(y,z,\varepsilon) = \mathbb{P}[Y_j = z \mid Y_{j-1} = y] = a(y,z)\varepsilon^{b(y,z)},$$

where $a(y,z)$ and $b(y,z)$ are non-negative constants for each pair (y,z). We can then write $\mu_0 = \mu_0(\varepsilon)$.

Models of highly reliable Markovian systems considered in Shahabuddin (1994), Nakayama (1996), L'Ecuyer and Tuffin (2007), for example, fit this framework. In these models, B represents the set of state where the system is failed and y_0 is the state of a new system. Each transition corresponds either to a failure of one or more components (it brings the system closer to B) or to a repair (it brings the system farther from B). Typically, $b(y,z) > 0$ for the failure transitions and $b(y,z) = 0$ for the repair transitions, so $\mu_0(\varepsilon) \to 0$ when $\varepsilon \to 0$.

For $y \notin \Delta$, let $\Pi_B(y)$ be the set of all sample paths $\pi = (y, y_1, \dots, y_\tau)$ where $y_\tau \in B$ and $y_j \notin \Delta$ for $j < \tau$. Let $\Pi_B = \Pi_B(y_0)$. Suppose that there is a set $\Pi_1 \subset \Pi_B$ such that

$$p(\pi,\varepsilon) = \prod_{j=1}^{\tau} p(y_{j-1}, y_j, \varepsilon) = a(\pi)\varepsilon^b + o(\varepsilon^b)$$

for all $\pi \in \Pi_1$, where $a(\pi) > 0$ and $b > 0$, and $p(\pi,\varepsilon) = o(\varepsilon^b)$ for all $\pi \notin \Pi_1$. We also assume that all cycles that belong to some path $\pi \in \Pi_B$ have probability $O(\varepsilon^\delta)$, for some constant $\delta > 0$. This implies that Π_1 must be finite, because it cannot contain paths having a cycle. When $\varepsilon \to 0$, Π_1 is the set of dominant paths of Π_B, in the sense that

$$\lim_{\varepsilon\to 0} \frac{1}{\mu_0(\varepsilon)} \sum_{\pi\in\Pi_1} p(\pi,\varepsilon) = 1.$$

We now apply IS with modified probabilities for which

$$\tilde{p}(\pi,\varepsilon) = \prod_{j=1}^{\tau} \tilde{p}(y_{j-1}, y_j, \varepsilon) = \frac{a(\pi)}{a} + o(1) \tag{9}$$

when $\varepsilon \to 0$, for all $\pi \in \Pi_1$, where $a = \sum_{\pi\in\Pi_1} a(\pi)$. We assume that for each path $\pi \in \Pi_B \setminus \Pi_1$ for which $p(\pi,\varepsilon) = \Theta(\varepsilon^{b(\pi)})$ for $b(\pi) > b$, we have $\tilde{p}(\pi,\varepsilon) = \Theta(\varepsilon^{c(\pi)})$ for some constant $c(\pi)$ that satisfies

$$2[b(\pi) - b] > c(\pi) > 0. \tag{10}$$

We also assume that for any state $y \notin \Delta$ that belongs to a path $\pi \in \Pi_B$, the probability of returning to y before hitting Δ under IS is strictly less than 1, and the likelihood ratio associated with any such cycle does not exceed 1 if ε is small enough.

Under these conditions, the IS estimator is $X_{\text{is}}(\varepsilon) = p(\pi,\varepsilon)/q(\pi,\varepsilon) = a\varepsilon^b + o(\varepsilon^b)$ if we reach B via a path $\pi \in \Pi_1$, which happens with probability $1 + o(1)$, and all other paths have $o(1)$ contribution (**?**).

L'Ecuyer and Tuffin (2007) propose heuristic methods to approximate zero-variance IS in this setting and illustrate the effectiveness of their proposals by numerical examples of highly reliable Markovian systems. After observing that

$$\mu(y) = \sum_{\pi\in\Pi_B(y)} p(\pi,\varepsilon),$$

they start with the simple idea of replacing $\mu(y)$ in (7) by its crude lower bound

$$v_0(y) = \max_{\pi\in\Pi_B(y)} p(\pi,\varepsilon),$$

which is usually easy to compute from any state y by solving a shortest path problem. This estimate is fine when a single path dominates the sum, but otherwise it may underestimate $\mu(y)$ significantly. Typically, this underestimation becomes worse when we are farther away from B. L'Ecuyer and Tuffin (2007) suggest refinements to this crude lower bound. A first improvement is to sum the probabilities of a set of disjoint dominant paths leading to B instead of just a single one. This gives a tighter lower bound, $\tilde{v}_0(y)$. In another refinement, they estimate $\mu(y_0)$ in preliminary runs with a crude IS strategy, then approximate $\mu(y)$ by a function $v_1(\cdot)$ that interpolates this estimate of μ_0 at y_0 and on the set B (where $\mu(y) = 1$) by raising $v_1(\cdot)$ to some power. They also propose further refinements.

In their numerical experiments, IS with these types of approximations performs much better than the popular heuristics such as simple failure biasing (SFB), balanced failure biasing (BFB), and balanced likelihood ratios, proposed earlier for these types of models (L'Ecuyer and Tuffin 2007). **?** show that if $\mu(y)$ in (7) is replaced by $v(y)$ where $v(y)$ is the sum of probabilities of all the dominant paths from y (those with the smallest power of ε), then we have VRE. In this sense, the method performs increasingly better when ε gets smaller. If we just take $v(y)$ as the probability of the most probable path to failure from y, then we have BRE but not VRE in general. With SFB or BFB, one can achieve BRE (under some conditions), but never VRE. □

In the previous examples, we have a DTMC model with a stopping time naturally defined by a subset of states. In other situations, the model might have to be reformulated appropriately so it fits our DTMC framework. This can always be done by including enough information in the state. For example, if we consider a continuous time Markov chain (CTMC) in which the cost depends on the holding time at each step, or for which the time horizon is fixed (instead of being the time when the CTMC hits a given set of states), then it suffices to include the current time (or the time that remains before reaching the horizon) in the state of the DTMC. Likewise, if we have a DTMC model with a fixed number of steps, then we must include the current step number (or the number of steps that remain) in the state.

Example 3 This example is adapted from De Boer et al. (2007). In Example 2, suppose that we have a CTMC $\{X(t), t \geq 0\}$ with state space $\mathscr{X}$ and that the transition probabilities are transition rates instead. Thus, the time until the next transition is always an exponential random variable with state-dependent rate. There is a fixed time horizon t_* and we want to estimate the probability μ_0 of reaching the set of states $\tilde{B}$ before time t_*, from a given initial state y_0.

Let $\{X_j, j \geq 0\}$ be the embedded DTMC (the sequence of visited states). We can easily reformulate this process as a DTMC by redefining the state at step j (just after the jth transition) as $Y_j = (X_j, R_j)$, where R_j is the remaining clock time before reaching t_*, $B = \{(X_j, R_j) : X_j \in \tilde{B} \text{ and } R_j > 0\}$ and $\Delta = B \cup \{(X_j, R_j) : R_j \leq 0\}$. The zero-variance change of measure is defined as usual, based on Y_j. Let $\mu(y) = \mu(x,t)$ be the probability of reaching $\tilde{B}$ before the time limit if the DTMC is in state $y = (x,t)$. We have

$$\mu(x,t) = \begin{cases} 1 & \text{if } x \in \tilde{B} \text{ and } t > 0, \\ 0 & \text{if } t \leq 0, \\ \mathbb{E}_{(x,t)}[\mu(X_1, R_1)] & \text{otherwise.} \end{cases}$$

In the latter case, we have

$$\mathbb{E}_{(x,t)}[\mu(X_1, R_1)] = \int_0^t \sum_{x' \in \mathscr{X}} \mu(x', t-\delta) p(x,x') \lambda_x \exp[-\lambda_x \delta] d\delta,$$

where λ_x is the total jump rate when the CTMC is in state x, and $p(x,x')$ is the probability that this jump is to state x'. Here, $p(x', t-\delta \mid x,t) = p(x,x')\lambda_x \exp[-\lambda_x \delta]$, for $\delta > 0$, is the density of the next state at $(x', t-\delta)$ when the current state is (x,t) for the DTMC.

To obtain a zero-variance IS estimator, we change the density p to a density $\tilde{p}$ defined by

$$\tilde{p}(x', t-\delta \mid x,t) = \frac{\mu(x', t-\delta)}{\mu(x,t)} \pi(x', t-\delta \mid x,t)$$

for $x' \in \mathscr{X}$ and $0 < \delta < t$, and 0 elsewhere, if $\mu(x,t) > 0$ (De Boer et al. 2007). Under this new density, the time until the next jump is no longer exponential, so we no longer have a CTMC. This time is nonzero only over the interval $(0,t)$, and it is not a truncated exponential either; in fact, since $\mu(x', t-\delta)$ is decreasing in δ, the right "tail" of the new distribution of the time to the next jump (before truncation) decreases faster than for the exponential distribution. When $\mu(x,t) = 0$, the density is unchanged: $\tilde{p}(x', t-\delta \mid x,t) = p(x', t-\delta \mid x,t)$.

De Boer et al. (2007) observe that sampling from this zero-variance IS is unpractical and propose a simple approximation, similar to those of the previous example, valid for the case where the expected number of transitions within the time horizon t_* is small. The idea is to approximate $\mu(y)$ by the sum of probabilities of a few dominant paths (those having the largest probability) among those leading to the rare set B. They give numerical examples where this approach performs well. In these examples, they consider a single dominant path, namely the "shortest path" leading to B. They also perform an asymptotic analysis in which the jump rates are parameterized by ε in the same way as the transition rates in Example 2. They show that their approach provides an IS estimator with BRE in general, and VRE under additional conditions similar to the VRE conditions in Example 2. □

4 ZERO VARIANCE VIA CONTROL VARIATES

A zero-variance estimator can also be defined via *control variates* (CV), as follows (Henderson and Glynn 2002, Henderson and Simon 2004, Kim and Henderson 2006). We again assume the same DTMC model as in Section 1.2. The idea is to replace the estimator $X = \sum_{j=1}^{\tau} c(Y_{j-1}, Y_j)$ by

$$X_{\text{cv}} = X - M_\tau$$

where

$$\begin{aligned} M_\tau &= \sum_{j=1}^{\tau} \big(c(Y_{j-1}, Y_j) + \mu(Y_j) - \mu(Y_{j-1})\big) \\ &= \sum_{j=1}^{\tau} \big(c(Y_{j-1}, Y_j) + \mu(Y_j) \\ &\qquad - \mathbb{E}[c(Y_{j-1}, Y_j) + \mu(Y_j) \mid Y_{j-1}]\big). \end{aligned}$$

One can easily verify that $\mathbb{E}[M_\tau] = 0$, so $\mathbb{E}[X_{\text{cv}}] = \mathbb{E}[X]$ (unbiasedness), and that $X_{\text{cv}} = \mu(y_0)$, a constant, because $\mu(Y_\tau) = 0$. Thus, we have a zero-variance estimator.

Of course, the function μ is unknown, but we still have $\mathbb{E}[M_\tau] = 0$ if we define

$$M_\tau = \sum_{j=1}^{\tau} \left(c(Y_{j-1}, Y_j) + v(Y_j) - \mathbb{E}[c(Y_{j-1}, Y_j) + v(Y_j) \mid Y_{j-1}] \right), \tag{11}$$

where v is an approximation of the function μ such that $v(y) = 0$ for $y \in \Delta$. This M_τ is a martingale, regardless of the choice of v. With a reasonably good approximation, the variance can sometimes be reduced significantly. Adaptive learning methods to find a good v are discussed in Section 5. Note that this methodology is usually not appropriate for rare-event simulation, because it does not make the rare events occur more frequently.

Kim and Henderson (2007) extends the method to regenerative simulation for the estimation of steady-state performance measure. Henderson and Glynn (2002) consider applications to infinite-horizon models with discounting and to stochastic differential equations.

5 APPROXIMATIONS AND ADAPTIVE LEARNING

For both IS and CV, the zero-variance schemes described earlier require the knowledge or a good approximation of the function μ. In previous examples, such approximations were obtained *a priori* by the sum of probabilities of a few dominating sample paths. More generally, in the rare event setting of Section 2, one can do well in some cases by replacing μ by an approximation obtained from a priori knowledge of the asymptotic behavior of the system, e.g., from large deviation's theory (Dupuis and Wang 2004, Juneja and Shahabuddin 2006, Asmussen and Glynn 2007, Blanchet and Glynn 2007).

Most practical approaches to approximate μ work by first selecting a parametric class of functions $\mathscr{V} = \{v(\cdot;\theta) : \mathscr{Y} \to \mathbb{R}, \theta \in \Theta\}$, where $\Theta \subseteq \mathbb{R}^m$, and $\theta = (\theta_1, \ldots, \theta_m)$ is a vector of parameters that we try to optimize so that $v = v(\cdot;\theta)$ is close to μ in some sense.

One natural and convenient way of doing this is to select a fixed set of functions $v_1, \ldots, v_m$, independent over $\mathbb{R}$, and define $\mathscr{V}$ as the space of all linear combinations of these functions, i.e.,

$$\mathscr{V} = \left\{ v = v(\cdot;\theta) = \sum_{i=1}^{m} \theta_i v_i(\cdot) \right\} \tag{12}$$

where $\theta = (\theta_1, \ldots, \theta_m) \in \mathbb{R}^m$. That is, $\mathscr{V}$ is the linear space spanned by the basis functions $v_1, \ldots, v_m$. In this setting, the vector θ is usually found by least-squares approximation: an estimate $\hat{\mu}(y)$ of $\mu(y)$ is obtained at a finite number of states y, say $y_1, \ldots, y_r \in \mathscr{Y}$, and then we compute

$$\theta \quad = \quad \arg\min_{\theta \in \Theta} \sum_{\ell=1}^{r} |v(y_\ell;\theta) - \hat{\mu}(y_\ell)|^2$$

via standard least-squares linear regression.

More generally, the parameterization can be nonlinear, in which case the optimal θ can be much more difficult to find. It can be approximated by statistical adaptive learning (Juneja and Shahabuddin 2006), using techniques such as stochastic approximation and (iterative) sample average approximation, for example (Kim and Henderson 2006). This can work even if no explicit characterization of μ is available a priori.

All these methods can converge to the zero-variance sampling only if the parameterized class $\mathscr{V}$ contains the function μ. Otherwise, the distance between μ and its closest approximation in $\mathscr{V}$ imposes a positive lower bound on the variance per simulation run. This means that if $\mathscr{V}$ is fixed and $\mu \notin \mathscr{V}$, the central-limit theorem will eventually kick in when the computing budget increases toward ∞, so the convergence rate is asymptotically $O(n^{-1/2})$ (in the probabilistic sense) in terms of the number n of simulation runs. But if $\mu \in \mathscr{V}$, we can do better. One simple case where $\mu \in \mathscr{V}$ is if the state space $\mathscr{Y}$ is finite, say $\mathscr{Y} = \{y_1, \ldots, y_m\}$, and we take the vector θ equal to the vector μ; that is, $\theta_i = \mu(y_i)$ for $i = 1, \ldots, m$. Kollman et al. (1999) consider a slightly more general case where $\mathscr{V}$ is a linear space of the form (12) that contains μ.

In general, the most difficult task is to select a good parameterized space $\mathscr{V}$, so that μ can be well approximated by at least one function from that space. In the linear case, this means a good set of basis functions $v_1, \ldots, v_m$. In the case of IS, another key issue is that an efficient method must be available to generate random variates from the transformed probabilities or densities. If we replace μ in (7) by an arbitrary approximation v, this would typically give a probability distribution for which no efficient random variate generation is available. This is an important limitation.

For IS, instead of parameterizing v, we can directly parameterize the IS distribution, and then minimize the variance of X_{is} with respect to these IS parameters. This type of approach has been the most popular in the past (Rubinstein and Shapiro 1993, Rubinstein 1999). It includes the cross-entropy method, for example. A major advantage is that we can restrict ourselves a priori to a parameterized class of distributions from which random variate generation is easy.

5.1 Adaptive Learning of Zero-Variance IS

Several algorithms have been proposed to approximate zero-variance IS by statistical learning of the function μ, or of

the IS distribution (or Markov chain transition kernel) that minimizes the variance within some class of distributions (Juneja and Shahabuddin 2006). Under certain conditions, the function μ and the zero-variance transition kernel can be learned adaptively at an exponential rate and this provides an estimator with exponential convergence, in the sense that the error converges as $O(e^{-\kappa n})$ for some constant $\kappa > 0$ instead of the usual $O(n^{-1/2})$ (**?**, Booth 1987, Kollman et al. 1999, **?**).

Kollman et al. (1999) have proposed an *adaptive importance sampling* method, also called *adaptive Monte Carlo* (AMC), which proceeds iteratively as follows. At step i, it uses an approximation $\mu^{(i)}$ of μ (or a guess of μ, at the first step), and performs r_i independent simulation replications using the probabilities (in the case of a discrete state space)

$$\tilde{P}^{(i)}(y,z) = \frac{P(y,z)(c(y,z)+\mu^{(i)}(z))}{\sum_{w\in\mathscr{Y}} P(y,w)(c(y,w)+\mu^{(i)}(w))}, \qquad (13)$$

obtained by replacing μ by $\mu^{(i)}$ in the zero-variance expression (7). For $\mu^{(i+1)}$, these authors use a linear model, which they fit to the simulation data at step i by least-squares regression. Under the assumptions that the linear model is exact (which is typically not quite realistic unless the model contains as many parameters as the number of states), and that the r_i's are large enough, they prove that $\mu^{(i)}$ converges to μ at an exponential rate when $i \to \infty$. This result also applies to variants of their algorithm, including one in which $\mu^{(i+1)}$ is defined as a convex combination of $\mu^{(i)}$ and the least-squares fit. **?** extend these results to general (continuous) state spaces. Note that this algorithm can be represented by a large Markov chain, whose state contains the current vector $\mu^{(i)}$. Desai and Glynn (2001) show that proving convergence to the zero-variance IS cannot be done in general with the standard techniques for analyzing geometric convergence of Harris-recurrent Markov chains, and requires specific methods of proof.

Ahamed, Borkar, and Juneja (2006) proposed another iterative learning algorithm that approximates μ by *adaptive stochastic approximation* (ASA), in the case where $\mathscr{Y}$ is finite. It starts with a given distribution of the initial state Y_0, an initial transition matrix $\tilde{P}^{(0)}$, and an initial guess $\mu^{(0)}$ of the function μ. The ASA algorithm runs a single sample path. At step j of the path, if the current state is $y_j \notin \Delta$, the current transition matrix $\tilde{P}^{(j)}$ is used to generate the next state y_{j+1}. From this new state, we update the estimate of $\mu(y_j)$ by the look-ahead rule:

$$\mu^{(j+1)}(y_j) = (1-a_j(y_j))\mu^{(j)}(y_j) + a_j(y_j)\left[c(y_j,y_{j+1}) + \mu^{(j)}(y_{j+1})\frac{P(y_j,y_{j+1})}{\tilde{P}^{(j)}(y_j,y_{j+1})}\right],$$

where $\{a_j(y),\ j \ge 0\}$, is a sequence of *step sizes* that satisfies $\sum_{j=1}^{\infty} a_j(y) = \infty$ and $\sum_{j=1}^{\infty} a_j^2(y) < \infty$ for each state y. The probability matrix is then updated as

$$\tilde{P}^{(j+1)}(y_j,y_{j+1}) = \max\left(\delta,\ P(y_j,y_{j+1})\frac{c(y_j,y_{j+1})+\mu^{(j+1)}(y_{j+1})}{\mu^{(j+1)}(y_j)}\right)$$

where $\delta > 0$ is a constant whose role is to ensure that the likelihood ratio remains bounded (to rule out the possibility that it takes huge values). For the other states, we leave $\mu^{(j+1)}(y) = \mu^{(j)}(y)$ and $\tilde{P}^{(j+1)}(y,z) = \tilde{P}^{(j)}(y,z)$. We then normalize via

$$\tilde{P}^{(j+1)}(y_j,y) = \frac{\tilde{P}^{(j+1)}(y_j,y)}{\sum_{z\in\mathscr{Y}} \tilde{P}^{(j+1)}(y_j,z)}$$

for all $y \in \mathscr{Y}$. When $y_j \in \Delta$, i.e., if the stopping time is reached at step j, y_{j+1} is generated again from the initial distribution, and the current probability matrix and estimate of μ are left unchanged. This method makes sense only if the state space is finite and not too large, because of the need to store the vector $\mu^{(j)}$ explicitly.

For larger state spaces, we can restrict ourselves a priori to a linear space as in (12), and use adaptive learning to optimize the parameter vector θ instead of learning $\mu(y)$ directly for each state y.

Example 4 Bolia, Juneja, and Glasserman (2004) use least-squares regression with a given set of basis functions to approximate the value function in the context of pricing an American option, as in the least-square Monte Carlo method (Glasserman 2004), and also approximate the zero-variance IS by plugging in (7) a least-squares approximation of the *continuation* value function (the value of the option conditional on not exercising it at the current step). In a first-stage, they learn good vectors of weights θ for each of these two value functions, and use them to approximate both the zero-variance IS and the optimal stopping rule. In a second stage, they use these approximations to estimate the value function more accurately. A key ingredient of their proposal is their choice of basis functions of the form $v_i(y) = \exp[a_i \log^2 y + b_i \log y]$, where the a_i and b_i are constants. With this particular choice, the IS densities at each step turn out to be mixtures of lognormal distributions, from which random variates are easy to generate. In their numerical illustrations, the variance is reduced by factors ranging from about 5 to 15, which is far from spectacular but still significant. The main factor limiting further variance reduction seems to be the lack of fit between the true value function and the best approximation with their selected basis functions. □

5.2 Adaptive Learning for Control Variates

For control variates, it is also natural to approximate μ via (12) and least-squares regression. If we denote the control variate (11) by $M_\tau(v)$, this gives

$$M_\tau(v(\cdot;\theta)) = \sum_{i=1}^{m} \theta_i M_\tau(v_i).$$

In other words, we have a linear combination of m control variates $M_\tau(v_i)$, so the optimal θ_i can be estimated via the standard least-squares regression procedure used for linear control variates (Glasserman 2004, Asmussen and Glynn 2007). Of course, the key issue again is the choice of basis functions $v_1,\ldots,v_m$. This is highly problem-dependent.

In the case where the parameterization of $v(y;\theta)$ is nonlinear, optimizing θ is more complicated. Kim and Henderson (2006) study and compare two stochastic optimization methods for that context, one based on stochastic approximation and the other based on sample average approximation. They demonstrate the feasibility of these methods by a numerical experiment with the pricing of barrier options. In their experiment, sample average approximation is the best performer.

6 ZERO VARIANCE VIA GENERALIZED ANTITHETIC VARIATES AND QUASI-MONTE CARLO

With *generalized antithetic variates* (GAV), we estimate $\mu_0 = \mathbb{E}[X]$ by the average of k dependent realizations of X, say $X^{(1)},\ldots,X^{(k)}$:

$$X_{\mathrm{a}} = \frac{1}{k}\sum_{i=1}^{k} X^{(i)}$$

(Wilson 1979, Ben-Ameur, L'Ecuyer, and Lemieux 2004, **?**). The variance of X_{a} is

$$\begin{aligned}\mathrm{Var}[X_{\mathrm{a}}] &= \frac{1}{k^2}\sum_{j=1}^{k}\sum_{\ell=1}^{k}\mathrm{Cov}[X^{(j)},X^{(\ell)}] \\ &= \frac{\mathrm{Var}[X]}{k} + \frac{2}{k^2}\sum_{j<\ell}\mathrm{Cov}[X^{(j)},X^{(\ell)}].\end{aligned}$$

We have $\mathrm{Var}[X_{\mathrm{a}}] < \mathrm{Var}[X]/k$ if and only if the sum of covariances is negative. It does not seem obvious a priori how zero variance can be achieved with the GAV scheme.

The classical *antithetic variates* (AV) scheme takes $k=2$, $X^{(1)} = f(U)$, and $X^{(2)} = f(1-U)$, assuming that $X = f(U)$ where $U \sim U(0,1)$. A very simple case is when $X = f(U) = aU + b$ where $U \sim U(0,1)$ (so f is an affine function). Then, AV gives $X_{\mathrm{a}} = a/2 + b$, a constant, so we have zero variance. This also works for a linear function of several variables. If f is monotone, AV always reduces the variance, but not to zero unless f is linear.

Zero variance (or faster convergence than $O(n^{-1/2})$) can also be obtained for certain spaces of nonlinear functions by more refined GAV schemes which are better known under the name of randomized quasi-Monte Carlo (RQMC). The usual RQMC framework is as follows. Suppose X can be written as $X = f(\mathbf{U})$ where $\mathbf{U}$ is a vector of s independent $U(0,1)$ random variables. Any expectation that admits an unbiased estimator by simulation can be written in this form, sometimes with $d = \infty$ (to cover the case where the required number of uniforms is random). An RQMC estimator has the form

$$\hat{\mu}_{n,\mathrm{rqmc}} = \frac{1}{n}\sum_{i=0}^{n-1} f(\mathbf{U}_i), \tag{14}$$

where $\{\mathbf{U}_0,\ldots,\mathbf{U}_{n-1}\}$ is a set of n random points that cover the unit hypercube $[0,1)^d$ very evenly (or uniformly, as a set) and such that each individual point is a random vector with the uniform distribution over $(0,1)^s$.

Two examples of such an RQMC point sets are a lattice rule with a random shift modulo 1 and a digital net with a random digital shift; see Niederreiter (1992), Sloan and Joe (1994), Owen (2003), L'Ecuyer and Lemieux (2002), **?** for their definitions. The uniformity of these point sets is often measured by a figure of merit that provides a bound on the worst-case variance over a given Hilbert space of functions. For example, for the Sobolev space of functions $f:[0,1)^s \to \mathbb{R}$ with square integrable mixed partial derivatives up to order α, for any $\delta > 0$, it is known that there exist RQMC rules giving $O(n^{-2\alpha-1+\delta})$ worst-case variance (**?**) and it is not difficult to find (by computer search) RQMC rules that achieve $O(n^{-2\alpha+\delta})$ worst-case variance (Dick et al. 2004, Dick et al. 2006).

In the case of a randomly-shifted lattice rule, the variance can be written exactly as

$$\mathrm{Var}[\hat{\mu}_{n,\mathrm{rqmc}}] = \sum_{\mathbf{0}\neq\mathbf{h}\in L_s^*} |\hat{f}(\mathbf{h})|^2, \tag{15}$$

where the $\hat{f}(\mathbf{h})$ are the coefficients of the Fourier expansion of the function f and L_s^* is the dual lattice, defined as the set of vectors $\mathbf{h} \in \mathbb{Z}^s$ whose scalar product modulo 1 with any point of the integration lattice is zero. If we can construct a lattice rule whose dual lattice contains none of the vectors $\mathbf{h}$ for which $\hat{f}(\mathbf{h}) \neq 0$, this gives a zero variance RQMC estimator. In particular, this is always possible (with a large enough n) if f has a Fourier expansion with a finite number of nonzero Fourier coefficients. For an arbitrary f, (15) indicates that we should try to select the lattice in order to minimize this expression. This can be achieved (in principle) by making sure that the dual of the selected lattice does not contain the vectors $\mathbf{h}$ associated with the largest square Fourier coefficients. This idea is discussed in L'Ecuyer

and Lemieux (2000) and **?**, for example. However, much research work is still needed for its practical implementation, which should involve adaptive algorithms.

Similar properties hold for digital nets with a random digital shift, with the coefficients of the Fourier expansion replaced by those of the Walsh coefficients (Dick 2008).

7 CONCLUSION

In theory, both IS and CV can achieve zero-variance for the simulation of a Markov chain. In practice, zero-variance sampling can only be approximated, usually by approximating the value function μ, and good approximations can be difficult to find in general. Least-squares approximation by a linear combination of a small (fixed) set of basis functions seems to be the natural solution in the case of continuous or very large state spaces. The main difficulty then is the choice of those basis functions. In the case of IS, another hurdle is that the approximation must be selected in a way that random variate generation from the modified distribution is not too hard. This problem does not appear with CV. Despite these difficulties, there are situations where approximate zero-variance sampling can provide very significant gains in efficiency, as exemplified by the numerical experiments reported in recent articles listed in our bibliography. More applications could benefit from this technology. This offers a fertile and exciting area for further research.

ACKNOWLEDGMENTS

This work has been supported by an NSERC-Canada Discovery Grant and a Canada Research Chair to the first author, INRIA's associated team MOCQUASIN to both authors, and the EuroNF Network of Excellence to the second author.

REFERENCES

Ahamed, I., V. S. Borkar, and S. Juneja. 2006. Adaptive importance sampling for Markov chains using stochastic approximation. *Operations Research* 54 (3): 489–504.

Asmussen, S., and P. W. Glynn. 2007. *Stochastic simulation*. New York: Springer-Verlag.

Ben-Ameur, H., P. L'Ecuyer, and C. Lemieux. 2004. Combination of general antithetic transformations and control variables. *Mathematics of Operations Research* 29 (4): 946–960.

Blanchet, J., and P. W. Glynn. 2007. Efficient rare-event simulation for the maximum of a random walk with heavy-tailed increments. Manuscript.

Bolia, N., S. Juneja, and P. Glasserman. 2004. Function-approximation-based importance sampling for pricing American options. In *Proceedings of the 2004 Winter Simulation Conference*, 604–611: IEEE Press.

Booth, T. E. 1987. Generalized zero-variance solutions and intelligent random numbers. In *Proceedings of the 1987 Winter Simulation Conference*, 445–451: IEEE Press.

De Boer, P. T., P. L'Ecuyer, G. Rubino, and B. Tuffin. 2007. Estimating the probability of a rare event over a finite horizon. In *Proceedings of the 2007 Winter Simulation Conference*: IEEE Press.

Desai, P. Y., and P. W. Glynn. 2001. A Markov chain perspective on adaptive Monte Carlo algorithms. In *Proceedings of the 2001 Winter Simulation Conference*, 379–384: IEEE Press.

Dick, J. 2008. Explicit constructions of quasi-Monte Carlo rules for the numerical integration of high-dimensional periodic functions. Manuscript.

Dick, J., I. H. Sloan, X. Wang, and H. Wozniakowski. 2004. Liberating the weights. *Journal of Complexity* 20 (5): 593–623.

Dick, J., I. H. Sloan, X. Wang, and H. Wozniakowski. 2006. Good lattice rules in weighted Korobov spaces with general weights. *Numerische Mathematik* 103:63–97.

Dupuis, P., and H. Wang. 2004. Importance sampling, large deviations, and differential games. *Stochastics and Stochastics Reports* 76:481–508.

Fishman, G. S. 1996. *Monte Carlo: Concepts, algorithms, and applications*. Springer Series in Operations Research. New York, NY: Springer-Verlag.

Glasserman, P. 2004. *Monte Carlo methods in financial engineering*. New York: Springer-Verlag.

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

Heidelberger, P. 1995. Fast simulation of rare events in queueing and reliability models. *ACM Transactions on Modeling and Computer Simulation* 5 (1): 43–85.

Henderson, S. G., and P. W. Glynn. 2002. Approximating martingales for variance reduction in Markov process simulation. *Mathematics of Operations Research* 27:253–271.

Henderson, S. G., and B. Simon. 2004. Adaptive simulation using perfect control variates. *Journal of Applied Probability* 41:859–876.

Juneja, S., and P. Shahabuddin. 2006. Rare event simulation techniques: An introduction and recent advances. In *Simulation*, ed. S. G. Henderson and B. L. Nelson, Handbooks in Operations Research and Management Science, 291–350. Amsterdam, The Netherlands: Elsevier. Chapter 11.

Kim, S., and S. G. Henderson. 2006. Adaptive control variates for finite-horizon simulation. *Mathematics of Operations Research* 32:508–527.

Kim, S., and S. G. Henderson. 2007. Non-linear control variates for regenerative steady-state simulation. In *Pro-*

ceedings of the 2007 Winter Simulation Conference, 430–438: IEEE Press.

Kollman, C., K. Baggerly, D. Cox, and R. Picard. 1999. Adaptive importance sampling on discrete Markov chains. *Annals of Applied Probability* 9 (2): 391–412.

Kuruganti, I., and S. Strickland. 1996. Importance sampling for Markov Chains: computing variance and determining optimal measures. In *Proceedings of the 1996 Winter Simulation Conference*, 273–280: IEEE Press.

L'Ecuyer, P., C. Lécot, and B. Tuffin. 2008. A randomized quasi-Monte Carlo simulation method for Markov chains. *Operations Research*. To appear.

L'Ecuyer, P., and C. Lemieux. 2000. Variance reduction via lattice rules. *Management Science* 46 (9): 1214–1235.

L'Ecuyer, P., and C. Lemieux. 2002. Recent advances in randomized quasi-Monte Carlo methods. In *Modeling Uncertainty: An Examination of Stochastic Theory, Methods, and Applications*, ed. M. Dror, P. L'Ecuyer, and F. Szidarovszky, 419–474. Boston: Kluwer Academic.

L'Ecuyer, P., and B. Tuffin. 2007. Effective approximation of zero-variance simulation in a reliability setting. In *Proceedings of the 2007 European Simulation and Modeling Conference*, 48–54. Ghent, Belgium: EUROSIS.

Nakayama, M. K. 1996. General conditions for bounded relative error in simulations of highly reliable Markovian systems. *Advances in Applied Probability* 28:687–727.

Niederreiter, H. 1992. *Random number generation and quasi-Monte Carlo methods*, Volume 63 of *SIAM CBMS-NSF Regional Conference Series in Applied Mathematics*. Philadelphia, PA: SIAM.

Owen, A. B. 2003. Variance with alternative scramblings of digital nets. *ACM Transactions on Modeling and Computer Simulation* 13 (4): 363–378.

Rubinstein, R. Y. 1999. The cross-entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability* 2:127–190.

Rubinstein, R. Y., and A. Shapiro. 1993. *Discrete event systems: Sensitivity analysis and stochastic optimization by the score function method*. New York: Wiley.

Shahabuddin, P. 1994. Importance sampling for the simulation of highly reliable Markovian systems. *Management Science* 40 (3): 333–352.

Sloan, I. H., and S. Joe. 1994. *Lattice methods for multiple integration*. Oxford: Clarendon Press.

Wilson, J. R. 1979. Proof of the antithetic-variates theorem for unbounded functions. *Math. Proc. Camb. Phil. Soc.* 86:477–479.

AUTHOR BIOGRAPHIES

PIERRE L'ECUYER is Professor in the Département d'Informatique et de Recherche Opérationnelle, at the Université de Montréal, Canada. He holds the Canada Research Chair in Stochastic Simulation and Optimization. His main research interests are random number generation, quasi-Monte Carlo methods, efficiency improvement via variance reduction, sensitivity analysis and optimization of discrete-event stochastic systems, and discrete-event simulation in general. He is currently Associate/Area Editor for *ACM Transactions on Modeling and Computer Simulation*, *ACM Transactions on Mathematical Software*, *Statistical Computing*, *International Transactions in Operational Research*, *The Open Applied Mathematics Journal*, and *Cryptography and Communications*. He obtained the *E. W. R. Steacie* fellowship in 1995-97, a *Killam* fellowship in 2001-03, and became an INFORMS Fellow in 2006. His recent research articles are available on-line from his web page: <http://www.iro.umontreal.ca/∼lecuyer>.

BRUNO TUFFIN received his PhD degree in applied mathematics from the University of Rennes 1 (France) in 1997. Since then, he has been with INRIA in Rennes. He spent 8 months as a postdoc at Duke University in 1999. His research interests include developing Monte Carlo and quasi-Monte Carlo simulation techniques for the performance evaluation of telecommunication systems, and developing new Internet-pricing schemes. He is currently Associate Editor for *INFORMS Journal on Computing*.