

Adaptive Simulation Using Perfect Control Variates

Shane Henderson, Cornell University

Burt Simon, University of Colorado at Denver

Abstract

We introduce adaptive-simulation schemes for estimating performance measures for stochastic systems based on the method of control variates. We consider several possible methods for adaptively tuning the control-variate estimators, and describe their asymptotic properties. Under certain assumptions, including the existence of a “perfect control variate”, all of the estimators considered converge faster than the canonical rate of $n^{-1/2}$, where n is the simulation runlength. Perfect control variates for a variety of stochastic processes can be constructed from “approximating martingales.” We prove a central limit theorem for an adaptive estimator that converges at rate $n^{-1}\sqrt{\ln n}$. A similar estimator converges at rate n^{-1} . An exponential rate of convergence is also possible under suitable conditions.

Keywords: Adaptive Monte Carlo, variance reduction, efficiency improvement, approximating martingales, control variates.

1 Introduction

Importance sampling is a stochastic simulation method used to reduce the variance of point estimators. If the importance sampling distribution is chosen appropriately, then one can often obtain large variance reductions. The primary difficulty in applying importance sampling is choosing the importance sampling distribution. Under certain conditions, there is a sampling distribution that yields a zero-variance estimator, and this ideal distribution can often be estimated as the simulation proceeds. Adaptive schemes for choosing the importance sampling distribution have been developed in a variety of application areas, e.g., [4, 5, 14, 1, 22, 21, 20, 16, 18, 9]. A survey of applications in rare-event simulation is given in [12]. It is possible to prove that certain adaptive schemes converge at an exponential rate [14, 2, 8, 7].

Another variance reduction method, known as control variates (see, e.g., [15]), can be viewed as an adaptive method. Suppose one wishes to estimate EX , where X is a real-valued random variable. If Y is a random variable with mean 0, then one can estimate EX by a sample mean of n i.i.d. replications of $X + uY$, where u is a real-valued parameter. The optimal (in the sense of minimizing variance) choice of u is $u^* = -\text{cov}(X, Y)/\text{var} Y$ which, in general, must be replaced by an estimator u_n based on sample covariances and variances. We can view this method as a parameterized variance reduction method, where u indexes the control variates $Y(u) = uY$. Under appropriate moment conditions, the estimator u_n converges almost surely to u^* . In general, the (minimal) variance of $X + Y(u^*)$ obtained from this method is strictly positive. Hence, while this method can reduce the simulation runlength required to reach a certain accuracy in the point estimate, it cannot reduce the *rate* of convergence, which remains at the canonical rate of $n^{-1/2}$. However, if $X + Y(u^*)$ has zero variance for some u^* (a “perfect control variate”) then, as with importance sampling, there may be ways to design adaptive schemes that converge faster than the canonical rate.

We define rates of convergence as follows. Throughout this paper “ \Rightarrow ” will denote weak convergence and “ \xrightarrow{P} ” will denote convergence in probability.

Definition 1 *Let $\{Z_n : n \geq 1\}$ be a sequence of real-valued random variables and $(h(n) : n \geq 1)$ be a sequence of deterministic, strictly-positive constants. We say that $\{Z_n\} = O_p(h(n))$ if, for all $\epsilon > 0$, there exists a deterministic constant $K > 0$ such that*

$$P\left(\frac{|Z_n|}{h(n)} > K\right) < \epsilon$$

for all $n \geq 1$.

This concept is also known as “tightness” or “boundedness in probability” [3, p. 37]. Observe that if $Z_n/h(n) \Rightarrow Z_\infty$ as $n \rightarrow \infty$, where Z_∞ is a real-valued random variable, then $Z_n = O_p(h(n))$.

Definition 2 *Suppose that $Z_n \xrightarrow{P} z$ as $n \rightarrow \infty$, where z is a deterministic constant. We say that Z_n converges to z at rate $h(n)$ if $\{Z_n - z\} = O_p(h(n))$.*

Note that rate of convergence is not uniquely defined. If Z_n converges at rate $h(n)$ and $g(n)/h(n)$ is bounded away from 0 as $n \rightarrow \infty$, then Z_n also converges at rate $g(n)$. In general, we will be interested in small values of $h(n)$. Also, note that if $h(n)^{-1}(Z_n - z) \Rightarrow Z_\infty$, where Z_∞ is a real-valued random variable, then Z_n converges to z at rate $h(n)$.

Perfect control variates can only occur in the parameterization $Y(u) = uY$ discussed above if X and Y are related through a direct linear relationship (i.e., $\text{Cor}(X, Y) = \pm 1$). This opportunity will rarely present itself in practice. However, it has recently been shown [11] that for a wide range of performance measures of Markov models, one can obtain a set of parameterized control variables $Y(u)$, based on approximating martingales, with $E(Y(u)) = 0$ for all u , and $\text{var}(X + Y(u^*)) = 0$ for some u^* . This suggests that one may be able to develop a class of adaptive control variate schemes that converge at rates faster than the canonical $n^{-1/2}$ rate. An *adaptive* simulation procedure is necessary to fully exploit a perfect control variate, since u^* is typically unknown. Adaptive in this context means using better and better estimates of u^* as the simulation proceeds. A nonadaptive scheme can do no better than the canonical convergence rate, even when a perfect control variate $Y(u^*)$ is available, if u^* is not known exactly. But adaptive schemes can provably obtain exponential rates of convergence under certain conditions; see Section 4.

An improvement in the convergence rate of a Monte Carlo scheme is very attractive, but a cautionary note is perhaps appropriate. A scheme may improve the convergence *rate* of the simulation method, but may require a large amount of work in each iteration. This additional work may mean that the scheme only outperforms standard Monte Carlo methods when seeking a high level of accuracy. With this in mind, we analyze several adaptive estimators with convergence rates of $n^{-1}\sqrt{\ln n}$, n^{-1} and $e^{-\gamma n}$ with $\gamma > 0$. The best estimator in practice may depend on factors in addition to the convergence rate.

A related issue is discussed in [6], where exponential convergence in the number of replications is empirically demonstrated, but the work per iteration is increasing. Nevertheless, a linear rate of convergence in terms of computational effort is empirically demonstrated. Our focus in this paper is on the case where a perfect control variate exists, but one

might also adopt adaptive methods with nonlinearly parameterized control variables when a perfect control variate does not exist. In this setting, one searches for the best of the parameterized control variates. Some results in this direction in the setting of steady-state simulation are stated in [10].

This paper is organized as follows. In Section 2 we describe a nontrivial example where perfect control variates exist. In Section 3 we provide a mathematical structure for the general estimation problem, and establish some useful properties of perfect control variate schemes. Then, in Section 4, we show how to exploit the structure that arises in the example in Section 2 to develop two adaptive estimators that, under certain conditions, converge at an exponential rate. This same structure is typical when one uses approximating martingales to develop perfect control variates, and so these estimators are quite generally applicable.

Some user intervention may be required to ensure that the estimators derived in Section 4 do, indeed, converge. Furthermore, the exponential rate of convergence relies on the presence of certain structure. In Section 5 we explore the performance of an adaptive procedure that requires no user intervention and considerably less structure. We prove a central limit theorem, thereby showing that a particular natural adaptive estimator, \mathcal{A}_n , converges at rate $n^{-1}\sqrt{\ln n}$. We also give an improved estimator, $\tilde{\mathcal{A}}_n$, that converges at rate n^{-1} . This is the best rate possible in the very general setting described in Section 5. Some conclusions and directions for future research are offered in Section 6.

Unless otherwise stated, all vectors are column vectors. The transpose of a vector v is denoted v' .

2 An Example

Let $Z = \{Z_n : n \geq 0\}$ be a discrete-time Markov chain on the finite state space $S = \{0, 1, \dots, d\}$ and suppose that Z reaches state 0 (the “absorbing state”) almost surely (and therefore in finite expected time) starting from any $Z_0 > 0$. Let $g : \{1, 2, \dots, d\} \rightarrow \mathbb{R}$

be a given cost function, and define

$$\mu(x) = E \left(\sum_{k=0}^{T-1} g(Z_k) \mid Z_0 = x \right), \quad 1 \leq x \leq d, \quad (1)$$

where $T = \inf\{k \geq 0 : Z_k = 0\}$ is the time till absorption. Then μ satisfies

$$\mu = g + P\mu, \quad (2)$$

where P is the sub-transition matrix of Z corresponding to the transient states $\{1, \dots, d\}$.

We assume that μ is unknown. Our goal is to estimate μ . Let $X \in \mathbb{R}^d$ be a column vector where the x th component, X^x , is a realization of $\sum_{k=0}^{T-1} g(Z_k)$ with $Z_0 = x$, so that $E(X) = \mu$. (The joint distribution of the components of X does not matter at this stage.)

One way to generate such an X is to simulate Z until absorption in 0 starting from each initial state $x = 1, 2, \dots, d$, but there are other possibilities as well.

We now describe our parameterized control variates, $Y(u)$. For any $u \in \mathbb{R}^d$ let $\tilde{u} = (0, u(1), \dots, u(d))'$ be the vector u augmented with an additional 0. Let $(\mathcal{G}_n : n \geq 0)$ be the natural filtration defined by $\mathcal{G}_n = \sigma(Z_0, \dots, Z_n)$ for $n \geq 0$. Then $M(u) = (M_n(u) : n \geq 0)$ is a martingale with respect to $(\mathcal{G}_n : n \geq 0)$, for any distribution on Z_0 , where

$$M_n(u) = \tilde{u}(Z_n) - \tilde{u}(Z_0) - \sum_{k=0}^{n-1} [(\tilde{P} - I)\tilde{u}](Z_k), \quad n \geq 0,$$

\tilde{P} is the transition matrix of Z , and I is the identity matrix. This is just the Dynkin martingale; see e.g., [13, p. 310]. The optional-sampling theorem allows us to conclude that $E(M_T(u) \mid Z_0 = x) = 0$ for any initial state $x \geq 1$. If

$$Y(u) = [-M_{T_1}^1(u), -M_{T_2}^2(u), \dots, -M_{T_d}^d(u)]'$$

where $M_{T_x}^x(u)$ is a realization of $M_T(u)$ with $Z_0 = x$, then

$$\hat{X}(u) \equiv X + Y(u) = [X^1 - M_{T_1}^1(u), \dots, X^d - M_{T_d}^d(u)]'$$

satisfies

$$E(\hat{X}(u)) = E(X) = \mu.$$

Thus, $Y(u)$ is a parameterized control variate for X . Furthermore, from (2) we see that if we choose $u(x) = \mu(x)$ for $x \geq 1$ then

$$Y(\mu) = \mu - X,$$

so

$$\hat{X}(\mu) = \mu$$

is not random. In other words, if we take $u = \mu$, the resulting estimator has zero variance, and so $Y(\mu)$ is a perfect control variate.

Notice that in this example, the parameter $u^* = \mu$ that gives zero variance is also the parameter that we wish to compute. This observation proves key in Section 4 where we establish an exponential rate of convergence for certain adaptive schemes that take advantage of this structure. First we generalize the setting and structure the general problem.

3 Structuring the Problem

Let (Ω, \mathcal{F}, P) be an underlying probability space for our simulation, and let $(\mathcal{F}_n : n \geq 0)$ be a filtration on (Ω, \mathcal{F}) , i.e., an increasing sequence of sigma fields. Intuitively, \mathcal{F}_n represents the information we have available after the n th simulation replication. Let $X \in \mathfrak{R}^{d'}$ be a random vector with $E(X) = \mu$ (the quantity of interest). We assume $E(\|X - \mu\|_{d'}^2) < \infty$, where $\|\cdot\|_{d'}$ is a metric on $\mathfrak{R}^{d'}$. Let $\{X_n : n \geq 1\}$ be i.i.d. replicates of X , where X_n is measurable with respect to \mathcal{F}_n ($X_n \in m\mathcal{F}_n$) for all $n \geq 1$. Let $u \in U$ index our control variates, where U is a subset of \mathbb{R}^d equipped with the metric $\|\cdot\|_d$. We will simply use $\|\cdot\|$ to denote our metrics since the context is always clear. (In many cases we can let $d' = 1$ without loss of generality since the components μ_i , $i = 1, \dots, d'$ of μ can be estimated separately. However, in Section 4 it is crucial that $d' = d$ since the elements of X cannot be estimated separately.) For $u \in U$, let $Y(u) \in \mathfrak{R}^{d'}$ be a random vector with $E(Y(u)) = 0$ and $E(\|Y(u)\|^2) < \infty$. For each fixed $u \in U$, $\{Y_n(u) : n \geq 1\}$ is a sequence of i.i.d. replicates of $Y(u)$, and furthermore, $Y_i(u)$ is independent of $Y_j(v)$ for $i \neq j$ and any $u, v \in U$. Also, $Y_n(u) \in m\mathcal{F}_n$ for all $n \geq 0$ and all $u \in U$.

We assume that there is an optimal choice u^* of u such that $E(\|X + Y(u^*)\|^2) = 0$, which implies that $X_n + Y_n(u^*) = \mu$ almost surely for all $n \geq 1$. For $n \geq 0$, let $u_n \in m\mathcal{F}_n$ denote our approximation of u^* at stage n .

For $n \geq 1$, let

$$\hat{X}_n = X_n + Y_n(u_{n-1})$$

denote an estimator of μ using the approximation u_{n-1} of u^* , and let

$$\nu(u) = E(\|X + Y(u) - \mu\|^2), \quad u \in U.$$

The \hat{X}_n 's are typically dependent since \hat{X}_n is a function of u_{n-1} . However, they are uncorrelated, as the following result shows.

Proposition 1 *The sequence $\{\hat{X}_n : n \geq 1\}$ consists of uncorrelated random vectors, and for $n \geq 1$,*

$$E(\|\hat{X}_n - \mu\|^2) = E(\nu(u_{n-1})).$$

Proof: Observe that for all $n \geq 1$, $E(\hat{X}_n | \mathcal{F}_{n-1}) = \mu$ almost surely. Now, let $i, j \in \{1, 2, \dots, d'\}$. For $n \geq 2$ and $1 \leq k \leq n-1$,

$$\begin{aligned} \text{cov}(\hat{X}_n(i), \hat{X}_{n-k}(j)) &= E(\hat{X}_n(i)\hat{X}_{n-k}(j)) - \mu_i\mu_j \\ &= E(E(\hat{X}_n(i)\hat{X}_{n-k}(j) | \mathcal{F}_{n-1})) - \mu_i\mu_j \\ &= E(E(\hat{X}_n(i) | \mathcal{F}_{n-1})\hat{X}_{n-k}(j)) - \mu_i\mu_j \\ &= E(\mu_i\hat{X}_{n-k}(j)) - \mu_i\mu_j = 0. \end{aligned}$$

To compute the expected squared error, observe that

$$\begin{aligned} E(\|\hat{X}_n - \mu\|^2) &= E(E(\|\hat{X}_n - \mu\|^2 | \mathcal{F}_{n-1})) \\ &= E(E(\|X_n + Y_n(u_{n-1}) - \mu\|^2 | \mathcal{F}_{n-1})) \\ &= E\nu(u_{n-1}). \end{aligned}$$

■

An adaptive estimator is specified by the method used to construct the sequence $\{u_n\}$, and the method of combining $\hat{X}_1, \dots, \hat{X}_n$ to obtain the n th estimate of μ .

4 Exponential Convergence

For certain problems it is possible to develop an adaptive estimator that converges at an exponential rate. In this section we describe the conditions that must prevail and establish that two estimators converge at an exponential rate. One of these conditions is that EX and u^* coincide, so in this section we have $d' = d$.

Let $X \in \mathfrak{R}^d$ be a random vector with $E(X) = \mu$. Suppose that we wish to compute $\alpha = f(\mu)$ for some known real-valued function f . Since f is known, this problem essentially reduces to estimating μ . Let $\{Y(u) : u \in U\}$ be a parameterized family of control variates with $U = \mathbb{R}^d$ containing a perfect control variate, $Y(u^*)$. The key requirement is that $\mu = u^*$, i.e., the parameter $u^* \in U$ that gives a perfect control variate corresponds with the vector μ that we wish to compute. The example presented in Section 2 is one such case. This condition also holds for all of the examples presented in [11] where approximating martingales were developed for a range of performance measures, so that we have strong motivation for considering this special case.

Let $\hat{X}(u) = X + Y(u)$. Under these conditions,

$$E(\hat{X}(u)) = u^*, \quad \forall u \in U,$$

and $\nu(u^*) = 0$, where

$$\nu(u) = E(\|\hat{X}(u) - u^*\|^2).$$

For $u \in U$, let $\{\hat{X}_{ni}(u), n, i = 1, 2, \dots\}$ be i.i.d. replicates of $\hat{X}(u)$. (Here we have extended our general framework to allow multiple replicates within each iteration, so that $\hat{X}_{ni}(u) \in m\mathcal{F}_n$ for all $n \geq 1, i \geq 1$ and $u \in U$.) Choose $u_0 \in U$ and $m < \infty$, and define, for $n \geq 1$,

$$u_n = \frac{1}{m} \sum_{i=1}^m \hat{X}_{ni}(u_{n-1}).$$

Let

$$\mathcal{E}_n = f(u_n)$$

be our estimator for α at stage $n \geq 1$.

Theorem 2 Suppose that $E(\|u_0 - u^*\|^2) < \infty$, and for some $\sigma^2 < \infty$,

$$E\left(\left\|\frac{1}{k} \sum_{i=1}^k \hat{X}_{ni}(u) - u^*\right\|^2\right) \leq \frac{\sigma^2}{k} \|u - u^*\|^2, \quad (3)$$

for every $k \geq 1$. Then

$$E(\|u_n - u^*\|^2) \leq \left(\frac{\sigma^2}{m}\right)^n E(\|u_0 - u^*\|^2). \quad (4)$$

If, in addition, $f : U \rightarrow \mathbb{R}$ is a Lipschitz function with Lipschitz constant c satisfying $f(u^*) = \alpha$, then

$$E((\mathcal{E}_n - \alpha)^2) \leq \left(\frac{\sigma^2}{m}\right)^n c^2 E(\|u_0 - u^*\|^2). \quad (5)$$

Proof: From (3) we have for all $n \geq 1$ that

$$\begin{aligned} E(\|u_n - u^*\|^2) &= E\left(\left\|\frac{1}{m} \sum_{i=1}^m \hat{X}_{ni}(u_{n-1}) - u^*\right\|^2\right) \\ &\leq \frac{\sigma^2}{m} E(\|u_{n-1} - u^*\|^2), \end{aligned}$$

and so (4) follows immediately. Since f has Lipschitz constant c we have

$$|\mathcal{E}_n - \alpha| \equiv |f(u_n) - f(u^*)| \leq c \|u_n - u^*\|.$$

Thus,

$$E((\mathcal{E}_n - \alpha)^2) \leq c^2 E(\|u_n - u^*\|^2)$$

and (5) follows from (4). ■

It immediately follows that as long as $m > \sigma^2$ we get an exponential rate of convergence. We might ask how best to choose m . The following heuristic argument shows how one might make a choice and also describes the kind of convergence rate one can expect from \mathcal{E}_n .

Suppose that the expected computational time required to simulate $\hat{X}(u)$ is finite and doesn't depend on $u \in U$. By suitably choosing the time scale, we can assume that this time is 1. The expected time needed to simulate \mathcal{E}_n is therefore $t = nm$, so that $n = t/m$. If we define \mathcal{E}'_t to be our estimate of α after t units of computer time then Theorem 2 suggests that

$$E((\mathcal{E}'_t - \alpha)^2) = O\left(\left(\frac{\sigma^2}{m}\right)^{t/m}\right).$$

Differentiating the right-hand side with respect to m yields an optimal value, $m^* = \sigma^2 e$. If we use $m = m^*$ we find that

$$E((\mathcal{E}'_t - \alpha)^2) = O\left(\exp\left(\frac{-t}{\sigma^2 e}\right)\right).$$

Theorem 2 above shows that as long as m is large enough, i.e., enough averaging is performed at each step of the algorithm, one can expect an exponential rate of convergence in n . It is reasonable to ask whether one can also obtain exponential convergence when a single replication is obtained at each step, i.e., $m = 1$. Define

$$\hat{u}_n = w\hat{u}_{n-1} + (1-w)\hat{X}_n(\hat{u}_{n-1})$$

and

$$\hat{\mathcal{E}}_n = f(\hat{u}_n),$$

where $w \in (0, 1)$. Theorem 3 shows that as long as w is chosen sufficiently close to 1, an exponential rate of convergence is again assured. Notice that while this estimator takes $m = 1$, averaging is again playing an important role, since the exponential smoothing operation that defines \hat{u}_n is also a form of averaging.

Theorem 3 *Let $(U, \|\cdot\|)$ be \mathbb{R}^d ($d < \infty$) with Euclidean distance. Under the conditions of Theorem 2, if we select $w \in (0, 1)$ then, for all $n \geq 1$,*

$$E(\|\hat{u}_n - u^*\|^2) \leq \psi(w)^n E(\|\hat{u}_0 - u^*\|^2) \tag{6}$$

and

$$E((\hat{\mathcal{E}}_n - \alpha)^2) \leq \psi(w)^n c^2 E(\|\hat{u}_0 - u^*\|^2), \tag{7}$$

where $\psi(w) = w^2 + (1-w)^2\sigma^2$. Hence the convergence $\hat{u}_n \rightarrow u^*$ and $\hat{\mathcal{E}}_n \rightarrow \alpha$ is exponential for any $w > (\sigma^2 - 1)/(\sigma^2 + 1)$ with $w \in (0, 1)$. The function $\psi(w)$ is minimized at $w^* = \sigma^2/(\sigma^2 + 1)$ where $\psi(w^*) = w^*$.

Proof: Since \hat{u}_{n-1} and $\hat{X}_n(\hat{u}_{n-1})$ are uncorrelated,

$$\begin{aligned} E(\|\hat{u}_n - u^*\|^2) &= E(\|w(\hat{u}_{n-1} - u^*) + (1-w)(\hat{X}_n(\hat{u}_{n-1}) - u^*)\|^2) \\ &= w^2 E(\|\hat{u}_{n-1} - u^*\|^2) + (1-w)^2 E(\|\hat{X}_n(\hat{u}_{n-1}) - u^*\|^2) \\ &\leq (w^2 + (1-w)^2\sigma^2) E(\|\hat{u}_{n-1} - u^*\|^2). \end{aligned} \tag{8}$$

The final inequality in (8) follows from (3) with $k = 1$. Thus,

$$E(\|\hat{u}_n - u^*\|^2) \leq \psi(w)^n E(\|\hat{u}_0 - u^*\|^2).$$

The remainder of the proof is straightforward and omitted. \blacksquare

Remark 1 If we allow $d = \infty$ and/or we use a nonEuclidean norm, then we can make use of the inequality $\|x + y\|^2 \leq 2\|x\|^2 + 2\|y\|^2$ to obtain a similar result to that given above. In this case one may need $m > 1$ (multiple replications) at each iteration to ensure that $2w^2 + 2(1 - w)^2\sigma^2/m < 1$.

We now verify that an exponential rate of convergence is possible for the example presented in Section 2 by verifying the conditions of Theorem 2. For $\ell \in \{1, \dots, d\}$, the ℓ th component of $\hat{X}(u)$, denoted $\hat{X}^\ell(u)$, is given by

$$\begin{aligned} \hat{X}^\ell(u) &= \sum_{k=0}^{T-1} g(Z_k) + u(\ell) + \sum_{k=0}^{T-1} [(P - I)u](Z_k) \\ &= u(\ell) + \sum_{k=0}^{T-1} [(P - I)(u - \mu)](Z_k) \\ &= u(\ell) + \sum_{j=1}^d V_\ell(j) [(P - I)(u - \mu)](j) \\ &= u(\ell) + V_\ell'(P - I)(u - \mu), \end{aligned} \tag{9}$$

where

$$V_\ell(j) = \sum_{k=0}^{T-1} 1(Z_k = j)$$

is the number of visits to state j before absorption, starting from state ℓ , and V_ℓ is the corresponding column vector.

Using Theorem 2 or Theorem 3, we can construct an estimator with an exponential rate of convergence if we can verify condition (3). Take $\|\cdot\|$ to be the Euclidean norm and write,

$$\begin{aligned} E\left(\left\|\frac{1}{k} \sum_{i=1}^k \hat{X}_{ni}(u) - u^*\right\|^2\right) &= E \sum_{\ell=1}^d \left(\frac{1}{k} \sum_{i=1}^k (\hat{X}_{ni}^\ell(u) - u^*(\ell)) \right)^2 \\ &= \sum_{\ell=1}^d \frac{\text{var}(\hat{X}^\ell(u))}{k}. \end{aligned}$$

From (9) we write

$$\begin{aligned}\text{var}(\hat{X}^\ell(u)) &= \text{var}(V_\ell'(P - I)(u - u^*)) \\ &= (u - u^*)'(P - I)'\eta_\ell(P - I)(u - u^*),\end{aligned}$$

where η_ℓ is the covariance matrix of V_ℓ ,

$$\eta_\ell = \text{cov}(V_\ell).$$

If we define $\eta = \sum_{\ell=1}^d \eta_\ell$, then

$$\begin{aligned}E\left(\left\|\frac{1}{k}\sum_{i=1}^k \hat{X}_{ni}(u) - u^*\right\|^2\right) &= \frac{1}{k}(u - u^*)'(P - I)'\eta(P - I)(u - u^*) \\ &\leq \frac{\|(P - I)'\eta(P - I)\|}{k}\|u - u^*\|^2,\end{aligned}$$

and so the conclusions of Theorems 2 and 3 apply with

$$\sigma^2 = \|(I - P)'\eta(I - P)\|.$$

Hence, for the example presented in Section 2, it is possible to construct estimators based on perfect control variate methods that converge at an exponential rate. For either of the estimators described above, one must ensure that sufficient averaging occurs, either by choosing m large enough for the first estimator, or by choosing w sufficiently close to 1 for the second method. If this requirement is not met, then the estimators may fail to converge. A similar “sufficient averaging” requirement was also needed for the adaptive-importance-sampling scheme described in [14], where the authors recommended guessing a reasonable value for the averaging parameter and observing the first few iterations. If it does not appear that convergence is taking place, then the process is interrupted, and the averaging per iteration is increased. In our case, this would correspond to increasing m or w , depending on which estimator is used. An alternative to this procedure is to either bound, or estimate, σ^2 , perhaps from some “pilot runs”, and then select m or w accordingly.

While these approaches are not unreasonable, one might ask if there is a way to avoid this experimentation. In other words, is there an adaptive estimator that converges rapidly, yet does not require any user intervention to ensure convergence?

5 Adaptive Estimators using Sample Means

In this section we develop an adaptive estimation procedure that requires no user intervention. It has the added advantage that it does not require the condition, imposed in the previous section, that the optimal choice of control variate parameter u^* coincide with the target quantity EX . It has the disadvantage that it no longer gives exponential convergence rates, but nevertheless, the procedure does give faster-than-canonical convergence rates. Since the restriction $u^* = EX$ has been eased, we can take $d' = 1$ without loss of generality, i.e., we can assume that X is real valued.

One can expect faster-than-canonical convergence rates under very general circumstances. Consider again the setup of Section 3. A reasonable estimator of $\mu = EX$ based on \mathcal{F}_n is the sample mean

$$\mathcal{A}_n = n^{-1} \sum_{i=1}^n \hat{X}_i. \quad (10)$$

This estimator is unbiased. Our next result shows that if the sequence of estimates u_n converges in probability to u^* as $n \rightarrow \infty$, then the variance of \mathcal{A}_n converges to 0 faster than at rate n^{-1} . Chebyshev's inequality then implies that \mathcal{A}_n converges to μ at a faster rate than $n^{-1/2}$.

Theorem 4 *If*

(a1) $u_n \xrightarrow{p} u^*$ as $n \rightarrow \infty$,

(b1) $\nu(u)$ is continuous at $u = u^*$, and

(c1) $(\nu(u_n) : n \geq 0)$ is uniformly integrable,

then

$$n \text{var}(\mathcal{A}_n) \rightarrow 0$$

as $n \rightarrow \infty$.

Proof: From (a1) and (b1) we conclude that $\nu(u_n) \xrightarrow{p} 0$ as $n \rightarrow \infty$. Assumption (c1) then yields $E\nu(u_n) \rightarrow 0$ as $n \rightarrow \infty$. Thus,

$$\frac{1}{n} \sum_{i=1}^n E\nu(u_{i-1}) \rightarrow 0$$

as $n \rightarrow \infty$. From Proposition 1 we conclude that

$$\text{var}(\mathcal{A}_n) = n^{-2} \sum_{i=1}^n E\nu(u_{i-1}), \quad (11)$$

and the result follows. ■

From (11) it follows that $\text{var}(\mathcal{A}_n) \geq n^{-2} E\nu(u_0)$, so in general the variance cannot converge faster than at rate n^{-2} .

We now attempt to more precisely identify the convergence rate of \mathcal{A}_n through a central limit theorem. The form of the central limit theorem depends on the method used to compute the sequence u_n of estimators of u^* . In this section we make the assumption that $u^* = E\xi_1$ and u_n is computed by averaging i.i.d. random elements, $\{\xi_n\}$, where $\xi_n \in \mathcal{F}_n$ for all $n \geq 1$. For example, in the “expected cost accumulated until absorption” problem described in Section 2, $u^* = EX$, and so one could apply this estimation procedure in that setting.

Consider the following adaptive simulation design.

1. Initialize $n = 0$, $\mathcal{A}_0 = 0$, and choose $u_0 \in U$.
2. Increment n , and simulate \hat{X}_n , and ξ_n .
3. Set

$$\mathcal{A}_n = n^{-1} \sum_{i=1}^n \hat{X}_i, \quad \text{and} \quad u_n = n^{-1} \sum_{i=1}^n \xi_i.$$

4. Go to step 2, or terminate.

The following theorem describes the convergence of \mathcal{A}_n in the special case when $U = \mathbb{R}$. We will briefly discuss the more general case where $U = \mathbb{R}^d$ for $d \geq 1$ after proving this result. Let $N(a, b)$ denote a normally-distributed random variable with mean a and variance b .

Theorem 5 *Let $U = \mathbb{R}$, and $u_n = \frac{1}{n} \sum_{i=1}^n \xi_i$, where*

$$E(\xi_1) = u^*, \quad \text{var}(\xi_1) = \sigma^2, \quad \text{and} \quad E((\xi_1 - u^*)^4) < \infty.$$

If ν'' is bounded and continuous, and for some $\delta, \gamma > 0$ and $c < \infty$,

$$E(|X + Y(u) - \mu|^{2+\delta}) \leq c|u - u^*|^{2+\gamma} \quad (12)$$

then

$$\frac{1}{\sqrt{\ln n}} \sum_{i=1}^n (\hat{X}_i - \mu) \Rightarrow N\left(0, \frac{\nu''(u^*)\sigma^2}{2}\right) \quad (13)$$

as $n \rightarrow \infty$. In addition, if u_0 is chosen deterministically, then

$$\frac{n^2}{\ln n} \text{var}(\mathcal{A}_n) \rightarrow \frac{\nu''(u^*)\sigma^2}{2}. \quad (14)$$

Remark 2 The requirement that u_0 be chosen deterministically can be relaxed. This condition is imposed to avoid situations where \mathcal{A}_1 has infinite variance.

The proof of the theorem relies on a series of results, stated as lemmas.

Lemma 6 Under the conditions of Theorem 5,

$$\frac{1}{\ln n} \sum_{i=1}^n \frac{1}{i^2} \sum_{j=1}^i (\xi_j - u^*)^2 \rightarrow \sigma^2$$

as $n \rightarrow \infty$ almost surely.

Proof: First, note that

$$\sum_{i=j}^n \frac{1}{i^2} = \frac{1}{j} - \frac{1}{n} + \varepsilon(j)$$

where

$$|\varepsilon(j)| \leq \frac{K}{j^2}$$

for some $K < \infty$. Therefore,

$$\begin{aligned} & \frac{1}{\ln n} \sum_{i=1}^n \frac{1}{i^2} \sum_{j=1}^i (\xi_j - u^*)^2 \\ &= \frac{1}{\ln n} \sum_{j=1}^n (\xi_j - u^*)^2 \sum_{i=j}^n \frac{1}{i^2} \\ &= \frac{1}{\ln n} \sum_{j=1}^n (\xi_j - u^*)^2 \left(\frac{1}{j} - \frac{1}{n} + \varepsilon(j) \right) \\ &= \frac{1}{\ln n} \sum_{j=1}^n \frac{(\xi_j - u^*)^2}{j} + \frac{1}{\ln n} \sum_{j=1}^n (\xi_j - u^*)^2 \varepsilon(j) - \frac{1}{n \ln n} \sum_{j=1}^n (\xi_j - u^*)^2 \\ &= [A1] + [A2] + [A3]. \end{aligned}$$

The strong law of large numbers yields

$$\frac{1}{n} \sum_{j=1}^n (\xi_j - u^*)^2 \rightarrow \sigma^2$$

as $n \rightarrow \infty$ almost surely, and so $[A3] \rightarrow 0$ as $n \rightarrow \infty$ a.s. To show $[A2] \rightarrow 0$ as $n \rightarrow \infty$ a.s., write

$$\alpha_j = (\xi_j - u^*)^2 \varepsilon(j),$$

and

$$S_n = \sum_{j=1}^n \alpha_j.$$

We have

$$E(\alpha_j) = \sigma^2 \varepsilon(j), \text{ and } \text{var}(\alpha_j) \leq E((\xi - u^*)^4) \varepsilon(j)^2.$$

Hence,

$$\frac{|E(S_n)|}{\ln n} \leq \frac{K\sigma^2}{\ln n} \sum_{j=1}^n \frac{1}{j^2} \rightarrow 0,$$

as $n \rightarrow \infty$, and

$$\sum_{j=1}^{\infty} \text{var}(\alpha_j) \leq K^2 E((\xi - u^*)^4) \sum_{j=1}^{\infty} \frac{1}{j^4} < \infty.$$

By Kolmogorov's convergence criterion (e.g., [19, p. 212]),

$$\frac{S_n}{\ln n} \rightarrow 0$$

as $n \rightarrow \infty$ almost surely. To show that $[A1] \rightarrow \sigma^2$ as $n \rightarrow \infty$ almost surely, we redefine

$$\alpha_j = \frac{1}{j} (\xi_j - u^*)^2.$$

This time

$$\frac{E(S_n)}{\ln n} = \frac{\sigma^2}{\ln n} \sum_{j=1}^n \frac{1}{j} \rightarrow \sigma^2$$

as $n \rightarrow \infty$ and

$$\sum_{j=1}^{\infty} \text{var}(\alpha_j) \leq E((\xi - u^*)^4) \sum_{j=1}^{\infty} \frac{1}{j^2} < \infty.$$

Kolmogorov's convergence criterion then implies that

$$\frac{S_n}{\ln n} \rightarrow \sigma^2$$

as $n \rightarrow \infty$ a.s. ■

Lemma 7 *Under the conditions of Theorem 5*

$$\frac{1}{\ln n} \sum_{i=1}^n (u_i - u^*)^2 \xrightarrow{p} \sigma^2$$

as $n \rightarrow \infty$.

Proof:

$$\begin{aligned} \frac{1}{\ln n} \sum_{i=1}^n (u_i - u^*)^2 &= \frac{1}{\ln n} \sum_{i=1}^n \left(\frac{1}{i} \sum_{j=1}^i (\xi_j - u^*) \right)^2 \\ &= \frac{1}{\ln n} \sum_{i=1}^n \frac{1}{i^2} \sum_{j=1}^i (\xi_j - u^*)^2 + \frac{1}{\ln n} \sum_{i=1}^n \frac{1}{i^2} \sum_{\substack{j_1, j_2 = 1 \\ j_1 \neq j_2}}^i (\xi_{j_1} - u^*)(\xi_{j_2} - u^*). \end{aligned}$$

The first term converges to σ^2 as $n \rightarrow \infty$ a.s. by Lemma 6, so we must show

$$\chi_n = \frac{1}{\ln n} \sum_{i=1}^n \frac{1}{i^2} \sum_{\substack{j_1, j_2 = 1 \\ j_1 \neq j_2}}^i (\xi_{j_1} - u^*)(\xi_{j_2} - u^*) \xrightarrow{p} 0$$

as $n \rightarrow \infty$. It suffices to show that $E(\chi_n^2) \rightarrow 0$. Let $a \wedge b = \min\{a, b\}$. Since $\{(\xi_i - u^*)\}$ is a zero-mean i.i.d. sequence,

$$\begin{aligned} E(\chi_n^2) &= \frac{1}{\ln^2 n} \sum_{i_1=1}^n \sum_{i_2=1}^n \frac{1}{i_1^2 i_2^2} E \left(\sum_{\substack{j_1, j_2 = 1 \\ j_1 \neq j_2}}^{i_1} (\xi_{j_1} - u^*)(\xi_{j_2} - u^*) \sum_{\substack{j'_1, j'_2 = 1 \\ j'_1 \neq j'_2}}^{i_2} (\xi_{j'_1} - u^*)(\xi_{j'_2} - u^*) \right) \\ &= \frac{2}{\ln^2 n} \sum_{i_1=1}^n \sum_{i_2=1}^n \frac{1}{i_1^2 i_2^2} \sum_{\substack{j_1, j_2 = 1 \\ j_1 \neq j_2}}^{i_1 \wedge i_2} E \left((\xi_{j_1} - u^*)^2 (\xi_{j_2} - u^*)^2 \right) \\ &\leq \frac{2\sigma^4}{\ln^2 n} \sum_{i_1=1}^n \sum_{i_2=1}^n \frac{1}{i_1^2 i_2^2} (i_1 \wedge i_2)^2 \\ &\leq \frac{2\sigma^4}{\ln^2 n} \left(\sum_{i_1=1}^n \sum_{i_2=1}^{i_1} \frac{1}{i_1^2} + \sum_{i_1=1}^n \sum_{i_2=i_1+1}^n \frac{1}{i_2^2} \right) \\ &= \frac{4\sigma^4}{\ln^2 n} \sum_{i=1}^n \left(\frac{1}{i} - \frac{1}{n} + \varepsilon(i) \right) \rightarrow 0 \end{aligned}$$

as $n \rightarrow \infty$. ■

Lemma 8 *In addition to the conditions of Theorem 5, suppose that r_i lies between u_i and u^* for each $i \geq 0$. Then*

$$\frac{1}{\ln n} \sum_{i=1}^n (\nu''(r_i) - \nu''(u^*)) (u_i - u^*)^2 \xrightarrow{P} 0$$

Proof: The Cauchy-Schwarz inequality yields, for $i \geq 1$,

$$E \left(|\nu''(r_i) - \nu''(u^*)| (u_i - u^*)^2 \right) \leq \sqrt{E \left((\nu''(r_i) - \nu''(u^*))^2 \right)} \sqrt{E \left((u_i - u^*)^4 \right)}.$$

Now, $u_i \rightarrow u^*$ as $i \rightarrow \infty$ a.s. and so the same is true of r_i . Since ν'' is bounded and continuous, the bounded convergence theorem ensures that

$$a_i = \sqrt{E \left[(\nu''(r_i) - \nu''(u^*))^2 \right]} \rightarrow 0$$

as $i \rightarrow \infty$. Now, $E(\xi_i - u^*) = 0$, and so

$$\begin{aligned} E \left((u_i - u^*)^4 \right) &= \frac{1}{i^4} \sum_{j_1, j_2, j_3, j_4=1}^i E \left((\xi_{j_1} - u^*) (\xi_{j_2} - u^*) (\xi_{j_3} - u^*) (\xi_{j_4} - u^*) \right) \\ &= \frac{1}{i^4} \left[i E \left((\xi_1 - u^*)^4 \right) + 6i(i-1) E \left((\xi_1 - u^*)^2 (\xi_2 - u^*)^2 \right) \right] \\ &\leq K^2 i^{-2} \end{aligned}$$

for some $K < \infty$. Thus,

$$E \left(\left| \frac{1}{\ln n} \sum_{i=1}^n (\nu''(r_i) - \nu''(u^*)) (u_i - u^*)^2 \right| \right) \leq \frac{K}{\ln n} \sum_{i=1}^n i^{-1} a_i \rightarrow 0 \quad (15)$$

as $n \rightarrow \infty$, since $a_i \rightarrow 0$ as $i \rightarrow \infty$. Since convergence in mean implies convergence in probability, the lemma is proved. \blacksquare

Proof of Theorem 5: The fact that the \hat{X}_i 's are uncorrelated ensures that $\{\sum_{i=1}^n (\hat{X}_i - \mu) : n \geq 0\}$ is a martingale with respect to the filtration $(\mathcal{F}_n : n \geq 0)$. We apply the martingale central limit theorem as given, for example, in [17, p. 442]. It suffices to show that

$$\sum_{i=1}^n E(W_{ni}^2 | \mathcal{F}_{i-1}) \xrightarrow{P} \frac{\nu''(u^*) \sigma^2}{2} \quad (16)$$

as $n \rightarrow \infty$ and for any $\epsilon > 0$,

$$\sum_{i=1}^n E(W_{ni}^2 1_{|W_{ni}| > \epsilon} | \mathcal{F}_{i-1}) \xrightarrow{P} 0 \quad (17)$$

as $n \rightarrow \infty$, where

$$W_{ni} = \frac{1}{\sqrt{\ln n}}(\hat{X}_i - \mu).$$

To verify (16) we first recall, from Proposition 1 that

$$E\left((\hat{X}_i - \mu)^2 \mid \mathcal{F}_{i-1}\right) = \nu(u_{i-1}).$$

Now, ν is twice continuously differentiable. Since $\nu(\cdot)$ is also nonnegative, it follows that $\nu'(u^*) = 0$. Hence, for $i \geq 1$,

$$\nu(u_{i-1}) = \frac{\nu''(r_{i-1})}{2}(u_{i-1} - u^*)^2$$

for some r_{i-1} between u_{i-1} and u^* . Thus,

$$\begin{aligned} \sum_{i=1}^n E(W_{ni}^2 \mid \mathcal{F}_{i-1}) &= \frac{1}{\ln n} \sum_{i=1}^n \frac{\nu''(r_{i-1})}{2}(u_{i-1} - u^*)^2 \\ &= \frac{1}{2 \ln n} \sum_{i=1}^n (\nu''(r_{i-1}) - \nu''(u^*)) (u_{i-1} - u^*)^2 + \frac{\nu''(u^*)}{2 \ln n} \sum_{i=1}^n (u_{i-1} - u^*)^2. \end{aligned}$$

Condition (16) now follows from Lemmas 7 and 8.

To verify (17) we write

$$\begin{aligned} \sum_{i=1}^n E(W_{ni}^2 1_{|W_{ni}| > \epsilon} \mid \mathcal{F}_{i-1}) &= \frac{1}{\ln n} \sum_{i=1}^n E([\hat{X}_i - \mu]^2 1_{|\hat{X}_i - \mu| > \epsilon \sqrt{\ln n}} \mid \mathcal{F}_{i-1}) \\ &\leq \frac{1}{\ln n} \sum_{i=1}^n E([\hat{X}_i - \mu]^2 \frac{|\hat{X}_i - \mu|^\delta}{\epsilon^\delta (\ln n)^{\delta/2}} \mid \mathcal{F}_{i-1}) \\ &\leq \frac{1}{\epsilon^\delta (\ln n)^{\delta/2}} \frac{1}{\ln n} \sum_{i=1}^n E(|\hat{X}_i - \mu|^{2+\delta} \mid \mathcal{F}_{i-1}) \\ &\leq \frac{c}{\epsilon^\delta (\ln n)^{\delta/2}} \frac{1}{\ln n} \sum_{i=1}^n |u_{i-1} - u^*|^{2+\gamma}. \end{aligned} \tag{18}$$

The law of the iterated logarithm ensures that for i sufficiently large

$$|u_i - u^*| \leq \left(\frac{3\sigma^2 \log \log i}{i} \right)^{1/2}$$

almost surely, and so

$$|u_i - u^*|^{2+\gamma} \leq \frac{c_0}{i}$$

for some constant $c_0 > 0$ and i sufficiently large, almost surely. Hence, (18) converges to 0 as $n \rightarrow \infty$ almost surely, and the central limit theorem (13) is proved.

To complete the proof of the theorem note, from Proposition 1, that

$$\begin{aligned} \frac{n^2}{\ln n} \text{var}(\mathcal{A}_n) &= \frac{1}{\ln n} \sum_{i=1}^n E(\nu(u_{i-1})) \\ &= \frac{1}{\ln n} \sum_{i=1}^n E\left(\frac{\nu''(r_{i-1})}{2}(u_{i-1} - u^*)^2\right). \end{aligned} \quad (19)$$

The arguments of Lemma 8 can then be applied to show that (19) converges to $\nu''(u^*)\sigma^2/2$ as $n \rightarrow \infty$. ■

Remark 3 It is straightforward to extend Theorem 5 to the case where $U = \mathbb{R}^d$ with $d < \infty$. In that case

$$\frac{1}{\sqrt{\ln n}} \sum_{i=1}^n (\hat{X}_i - \mu) \Rightarrow N(0, \eta^2)$$

as $n \rightarrow \infty$, where the variance constant

$$\eta^2 = \frac{1}{2} \sum_{i=1}^d \sum_{j=1}^d \nabla^2 \nu_{i,j}(u^*) \Lambda_{ij}$$

and Λ is the covariance matrix of the vector ξ_1 .

Remark 4 The conditions of Theorem 5 can be verified for the cost to absorption problem given in Section 2. The vector ξ_1 is given by X . The ℓ th component X^ℓ of X is bounded in absolute value by $\|g\|T_\ell$, and the stopping time T_ℓ has a finite moment generating function in a neighborhood of 0. Hence the moment conditions on ξ_1 are satisfied. The result (9) ensures that (12) holds with $\delta = \gamma$ for any $\delta > 0$. Hence, the estimator \mathcal{A}_n satisfies the CLT.

Note that \mathcal{A}_n weights all of the values \hat{X}_n equally, but we expect that \hat{X}_n is more accurate for larger n . Consider the slightly-modified unbiased estimator

$$\tilde{\mathcal{A}}_n = \sum_{i=1}^n \frac{2i}{n(n+1)} \hat{X}_i,$$

where u_n is as in Theorem 5. The weighting terms are chosen to minimize the asymptotic variance.

Theorem 9 *Under the conditions of Theorem 5, (and assuming that u_0 is chosen deterministically),*

$$n^2 \text{var}(\tilde{\mathcal{A}}_n) \rightarrow \sigma^2 \nu''(u^*)$$

as $n \rightarrow \infty$.

Proof:

$$\begin{aligned} \text{var}(\tilde{\mathcal{A}}_n) &= \sum_{i=1}^n \frac{4i^2}{n^2(n+1)^2} E\left(\frac{\nu''(r_{i-1})}{2} (u_{i-1} - u^*)^2\right) \\ &= \sum_{i=1}^n \frac{4i^2}{n^2(n+1)^2} E\left(\frac{\nu''(u^*)}{2} (u_{i-1} - u^*)^2\right) \\ &\quad + \sum_{i=1}^n \frac{4i^2}{n^2(n+1)^2} E\left(\frac{\nu''(r_{i-1}) - \nu''(u^*)}{2} (u_{i-1} - u^*)^2\right) \\ &= A + B. \end{aligned}$$

We have

$$n^2 A = \frac{2\nu''(u^*)}{(n+1)^2} E((u_0 - u^*)^2) + \frac{2\nu''(u^*)}{(n+1)^2} \sum_{i=2}^n i^2 \frac{\sigma^2}{i-1} \rightarrow \nu''(u^*) \sigma^2$$

as $n \rightarrow \infty$. As in the proof of Lemma 8 we can show that $n^2 B \rightarrow 0$ as $n \rightarrow \infty$ and the result follows. \blacksquare

Thus $\tilde{\mathcal{A}}_n$ converges slightly faster (at rate n^{-1}) than \mathcal{A}_n . Since the weighting factors are asymptotically optimal for $\tilde{\mathcal{A}}_n$, no estimator that averages the \hat{X}_n 's can converge faster.

6 Conclusions

We have discussed rates of convergence of several adaptive control variate estimators in the presence of a perfect control variate. In the general case discussed in Section 5, it is possible to obtain a rate of convergence of n^{-1} , which is clearly attractive when compared with the canonical rate $n^{-1/2}$. In the special case discussed in Section 4, an exponential rate of convergence is possible.

The question arises whether exponential convergence, or $O_p(n^{-1}\sqrt{\ln n})$ for that matter, is feasible in practice if d is large or infinite. Even if d is finite, the bound on σ^2 may be

so large that from a practical point of view, the estimators \mathcal{E}_n and $\tilde{\mathcal{E}}_n$ are not feasible. It may be possible to decrease the bound on σ^2 by using a nonEuclidean metric on \mathbb{R}^d . For example, in the expected cost to absorption problem, it seems that if a state j is very rarely visited, then the estimate $u_n(j)$ of $u^*(j)$ need not be as accurate as the corresponding estimate for states that are more frequently visited. A metric reflecting this asymmetry may reduce the bound on σ^2 considerably. In any case, we expect that when perfect control variates are available, one or more of our adaptive estimators are superior to the standard Monte Carlo estimator.

Another question that we are actively pursuing is that of how to proceed when one has a nonlinearly parameterized control variate, but a perfect control variate does not exist within the parameterization. This case is quite likely to arise in applications when one uses a limited class of approximating martingale processes. For example, this problem was considered in the context of steady-state simulation of a class of Markov chains related to multiclass networks in [10], where a scheme involving stochastic approximation to tune the parameter u was described.

Acknowledgments

The work of the first author was partially supported by NSF Grant Number DMI-0230528.

References

- [1] W. A. Al-Qaq, M. Devetsikiotis, and J.-K. Townsend. Stochastic gradient optimization of importance sampling for the efficient simulation of digital communication systems. *IEEE Transactions on Communications*, 43:2975–2985, 1995.
- [2] K. Baggerly, D. Cox, and R. Picard. Exponential convergence of adaptive importance sampling for Markov chains. *J. Appl. Probab.*, 37(2):342–358, 2000.
- [3] P. Billingsley. *Convergence of Probability Measures*. Wiley, New York, 1968.

- [4] T. E. Booth. Exponential convergence for Monte-Carlo particle-transport. *Transactions of the American Nuclear Society*, 50:267–268, 1985.
- [5] T. E. Booth. Zero-variance solutions for linear Monte-Carlo. *Nuclear Science and Engineering*, 102:332–340, 1989.
- [6] T. E. Booth. An approximate Monte Carlo adaptive importance sampling method. *Nuclear Science and Engineering*, 138:96–103, 2001.
- [7] P. Desai and P. W. Glynn. A Markov chain perspective on adaptive Monte Carlo algorithms. In B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, editors, *Proceedings of the 2001 Winter Simulation Conference*, pages 379–384, Piscataway NJ, 2001. IEEE.
- [8] P. Y. Desai. *Adaptive Monte Carlo Methods for Solving Eigenvalue Problems*. PhD thesis, Department of Management Science and Engineering, Stanford University, Stanford, CA, 2001.
- [9] M. Fitzgerald and R. Picard. Accelerated Monte Carlo for particle dispersion. *Communications in Statistics, Part A – Theory and Methods*, 30(11):2459–2471, 2001.
- [10] S. Henderson, S. P. Meyn, and V. Tadic. Performance evaluation and policy selection in multiclass networks. *Discrete Event Dynamic Systems*, 13:149–189, 2003. Special issue on learning and optimization methods.
- [11] S. G. Henderson and P. W. Glynn. Approximating martingales for variance reduction in Markov process simulation. *Mathematics of Operations Research*, 27:253–271, 2002.
- [12] M. Hsieh. Adaptive Monte Carlo methods for rare event simulations. In E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, editors, *Proceedings of the 2002 Winter Simulation Conference*, pages 108–115, Piscataway NJ, 2002. IEEE.
- [13] S. Karlin and H. M. Taylor. *A Second Course in Stochastic Processes*. Academic Press, Boston, 1981.

- [14] C. Kollman, K. Baggerly, D. Cox, and R. Picard. Adaptive importance sampling on discrete Markov chains. *Ann. Appl. Probab.*, 9(2):391–412, 1999.
- [15] A. M. Law and W. D. Kelton. *Simulation Modeling and Analysis*. McGraw-Hill, New York, 3rd edition, 2000.
- [16] D. Lieber, A. Nemirovskii, and R. Y. Rubinstein. A fast Monte Carlo method for evaluating reliability indexes. *IEEE Transactions on Reliability*, 48:256–261, 1999.
- [17] R. S. Liptser and A. N. Shiriyayev. *Theory of Martingales*. Kluwer Academic, Boston, 1989.
- [18] R. R. Picard, M. Fitzgerald, and M. J. Brown. Accelerating convergence in stochastic particle dispersion simulation codes. *Journal of Computational Physics*, 173:231–255, 2001.
- [19] S. I. Resnick. *A Probability Path*. Birkhäuser, Boston, 2001.
- [20] R. Y. Rubinstein. The cross-entropy method for combinatorial and continuous optimization. *Methodology and Computing in Applied Probability*, 1:127–190, 1999.
- [21] Y. Su and M. C. Fu. Importance sampling in derivative securities pricing. In J. A. Joines, R. R. Barton, K. Kang, and P. A. Fishwick, editors, *Proceedings of the 2000 Winter Simulation Conference*, pages 587–596, Piscataway NJ, 2000. IEEE.
- [22] F. Vázquez-Abad and D. Dufresne. Accelerated simulation for pricing Asian options. In D. Medeiros, E. Watson, J. S. Carson, and M. S. Manivannan, editors, *Proceedings of the 1998 Winter Simulation Conference*, pages 1493–1500, Piscataway, NJ, 1998. IEEE.