*Proceedings of the 2001 Winter Simulation Conference*
*B. A. Peters, J. S. Smith, D. J. Medeiros, and M. W. Rohrer, eds.*

# A MARKOV CHAIN PERSPECTIVE ON ADAPTIVE MONTE CARLO ALGORITHMS

Paritosh Y. Desai
Peter W. Glynn

Department of Management Science and Engineering
Stanford University
Stanford, CA 94309, U.S.A.

## ABSTRACT

This paper discusses some connections between adaptive Monte Carlo algorithms and general state space Markov chains. Adaptive algorithms are iterative methods in which previously generated samples are used to construct a more efficient sampling distribution at the current iteration. In this paper, we describe two such adaptive algorithms, one arising in a finite-horizon computation of expected reward and the other arising in the context of solving eigenvalue problems. We then discuss the connection between these adaptive algorithms and general state space Markov chain theory, and offer some insights into some of the technical difficulties that arise in trying to apply the known theory for general state space chains to such adaptive algorithms.

## 1 INTRODUCTION

This paper discusses the connection between adaptive Monte Carlo algorithms and Markov chain theory in general state space. An adaptive algorithm is basically a *learning algorithm* which proceeds iteratively by using previously generated samples to essentially build a current *best guess* of the solutions. The current *best guess* is then used to create a more effective sample on the current iteration. The hope is that by *learning* from previous iterations, such an adaptive algorithm can obtain faster rates of convergence than those enjoyed by conventional sampling based algorithms.

As the above informal discussion suggests, such an algorithm can be viewed as a Markov chain taking values in a general state space. In particular, the *best guess* of the solution at the next iteration is basically just a (complicated) function of the current *best guess* and the newly generated sample that is required for the next iteration. This suggests that the sequence of *best guesses* is Markovian. Of course, the space required to describe the current *best guess* generally needs to be continuous. Consequently , the Markov chain of *best guesses* is a general state space chain

This paper is organized as follows. In Section 2, we discuss an adaptive algorithm for computing expected reward to absorption in the finite state Markov chain setting. Section 3 describes a new adaptive algorithm for solving eigenvalue problems for non-negative matrices. Finally, Section 4 discusses our two examples through the prism of general state space Markov chain theory. In particular, we show that much of the standard tool-set that is available for the analysis of general state space chains is inapplicable to the class of chains that arise in the adaptive Monte Carlo context. This means that one must work hard to create specific techniques that can handle the chains that arise in the adaptive context; see Kollman et al. (1996) and Desai (2001) for more discussion of the specific tool-set that is effective in analyzing such algorithms.

## 2 AN ADAPTIVE ALGORITHM FOR COMPUTING EXPECTED REWARD TO ABSORPTION

Suppose that $Y = \{Y_n : n \geq 0\}$ is a Markov chain taking values in a finite state space $S$ and possessing a transition matrix $\tilde{P} = (\tilde{P}(x, y) : x, y \in S)$. For a non-empty *target* set $B^c \subseteq S$ and a nonnegative real-valued function $f : S \to \mathbb{R}$, consider the random variable

$$R = \sum_{n=0}^{T-1} f(Y_n),$$

where $T = \inf\{n \geq 0 : Y_n \in B^c\}$ is the *first hitting time* of $B^c$. If we interpret $f(x)$ as the reward earned by spending one unit of time in $x \in S$, then $R$ is the cumulative reward accumulated by $Y$ up to *absorption* in the set $B^c$

For $x \in S$, let $u^*(x) = \tilde{E}_x R$, where $\tilde{E}_x(\cdot)$ is the expectation operator on the path-space of $Y$ under which $Y_0 = x$. It is well known that $(u^*(x) : x \in S)$ is the minimal non-negative solution to the linear system of equations

$$u = f + \tilde{P}u \tag{1}$$

subject to the boundary condition $u(x) = 0$ for $x \in B^c$.

The conventional Monte-Carlo algorithm for computing $u^*(x)$ involves simulating the Markov chain $Y$, conditional on $Y_0 = x$, up to the random time $T$. From such a simulation, $R$ can be computed. By independently replicating this experiment $n$ independent and identically distributed (iid) times, we obtain iid copies $R_1, R_2, \ldots, R_n$ of the random variable $R$. The quantity $u^*(x)$ can then be estimated via the sample mean of the $R_i$'s. We now describe an ingenious adaptive algorithm for computing $u^*$ that is due to Booth (1985) and has been analyzed by Kollman et al. (1996). Let $u = (u(x) : x \in B)$ be our current *best guess* of the solution. (Note that $u^*(x) = 0$ for $x \in B^c$, so only those components of $u^*$ corresponding to $B$ need to be computed.) We will use $u$ to construct an *importance distribution* that will drive our sampling of the chain on the next iteration.

Given a current (non-negative) guess $u$, we let $Q(u) = (Q(x, y, u) : x, y \in S)$be the transition matrix given by

$$Q(x, y, u) = \begin{cases} \frac{\tilde{P}(x,y)(f(x)+u(y))}{N(x)}, & x \in B, y \in B \\ \frac{\tilde{P}(x,y)}{N(x)}, & x \in B, y \in B^c \\ \tilde{P}(x, y), & x \in B^c, y \in S \end{cases} \quad (2)$$

where $N(x)$ is the *normalization* factor given by

$$N(x) = \sum_{y \in B} \tilde{P}(x, y)(f(x) + u(y)) + \sum_{y \in B^c} \tilde{P}(x, y) f(x). \quad (3)$$

Rather than simulating $Y$ up to time $T$ using the original dynamics associated with $\tilde{P}$, we instead simulate $Y$ up to time $T$ using the transition matrix $Q(u)$, and adjust the simulation output using the appropriate likelihood ratio (see Glynn and Iglehart (1989) for background on *importance sampling*):

$$\tilde{R} = \sum_{n=0}^{T-1} f(Y_n) \prod_{i=0}^{n-1} \frac{\tilde{P}(Y_i, Y_{i+1})}{Q(Y_i, Y_{i+1}, u)}. \quad (4)$$

By replicating $\tilde{R}$ $k$ iid times, conditional on $Y_0 = x$, we can estimate $u^*(x)$ via the corresponding sample mean. By repeating this process for each $x \in B$, we obtain an estimator for the entire solution vector $u^* = (u^*(x) : x \in B)$. This estimator for $u^*$ can be used to compute the transition matrix $Q(u)$ for the next iteration of the algorithm.

This *best guess* at iteration $n$ is the estimator of $u^*$ obtained at that iteration. Let $X_n$ be the estimator at iteration $n$. Note that $X_n \in \mathbb{R}_+^d$, where $d$ is the number of states in $B$. It is clear intuitively (and easily proved rigorously) that $X = (X_n : n \geq 0)$ is a (time-homogeneous) Markov chain.

Obviously, the goal in designing such an algorithm is that $X_n$ converges to $u^*$ as $n \to \infty$. If such a convergence is to occur, it seems reasonable to expect that the solution $u^* = (u^*(x) : x \in B)$ is an *absorbing point* for $X$.

**Proposition 1.** *Suppose $\tilde{P}$ is irreducible. Then, if $X_0 = u^*$, it follows that $X_1 = u^*$ a.s.*

**Proof** The irreducibility of $\tilde{P}$ ensures that $u^*$ is finite-valued. Note that if $u = u^*$, then the normalization factor $N(x)$ appearing in the definition of $Q(u^*)$ can be written as

$$\begin{aligned} N(x) &= f(x) \sum_{y \in S} \tilde{P}(x, y) + \sum_{y \in B^c} \tilde{P}(x, y) u^*(y) \\ &= f(x) + \sum_{y \in B^c} \tilde{P}(x, y) u^*(y) \\ &= u^*(x) \end{aligned}$$

for $x \in B$. Consequently, $\tilde{R}$ takes the form

$$\tilde{R} = \sum_{n=0}^{T-1} f(Y_n) \prod_{i=0}^{n} \frac{u^*(Y_i)}{f(Y_i) + u^*(Y_{i+1})}.$$

The proposition follows if we can show that $\tilde{R} = u^*(Y_0)$ a.s. We prove this by induction on the length $T$ of the path. Note that if $T = 1$, then

$$\tilde{R} = f(Y_0) \cdot \frac{u^*(Y_0)}{f(Y_0) + u^*(Y_1)}. \quad (5)$$

But, if $T = 1$, then $Y_1 \in B^c$ and $u^*(Y_1) = 0$, proving that $\tilde{R} = u^*(Y_0)$.

Suppose that the result is correct for all paths of length less than or equal to $n$. If $T = n + 1$, then

$$\begin{aligned} \tilde{R} &= f(Y_0) \frac{u^*(Y_0)}{f(Y_0) + u^*(Y_1)} \\ &+ \frac{u^*(Y_0)}{f(Y_0) + u^*(Y_1)} \sum_{j=0}^{\tilde{T}-1} f(Y_{1+j}) \prod_{i=0}^{j} \frac{u^*(Y_{1+i})}{f(Y_{1+i}) + u^*(Y_{1+i+1})}, \end{aligned}$$

where $\tilde{T} = \inf\{n \geq 0 : Y_{1+n} \in B^c\}$. If $T = n + 1$, then $\tilde{T} = n$ and the induction hypothesis guarantees that

$$\tilde{R} = f(Y_0) \frac{u^*(Y_0)}{f(Y_0) + u^*(Y_1)} + \frac{u^*(Y_0)}{f(Y_0) + u^*(Y_1)} \cdot u^*(Y_1), \quad (6)$$

from which we may conclude that $\tilde{R} = u^*(Y_0)$. □

Consequently, the Markov chain $X$ has the property that the solution vector $u^*$ is a *fixed point* of the algorithm.

## 3 AN ADAPTIVE ALGORITHM FOR MONTE CARLO SOLUTION OF EIGENVALUE PROBLEMS

Let $G = (G(x, y) : x, y \in S)$ be a finite non-negative irreducible matrix. The Perron-Frobenius theorem (see, for example, Chapter 1 of Seneta (1973)) then guarantees the existence of a unique strictly positive eigenvalue $\lambda$ and a corresponding strictly positive column eigenvector $u^*$ such that

$$Gu^* = \lambda u^*. \tag{7}$$

The *eigenvalue problem* associated with $G$ involves solving for $\lambda$ and $u^*$. Such eigenvalue problems arise in many different application settings; (see Chapter 1 of Desai (2001)).

We will now describe an adaptive Monte Carlo algorithm for solving such eigenvalue problems; a complete discussion can be found in Desai and Glynn (2001). We start by describing a probabilistic representation for $\lambda$ and $u^*$ that can be found in Glynn (1996). We assume, without loss of generality, that $G$ has been *normalized* so as to be substochastic. We expand the state space $S$ to $\tilde{S} = S \bigcup \{\Delta\}$ where $\Delta$ is a *cemetery state* not in $S$, and let $\tilde{P} = (\tilde{P}(x, y) : x, y \in \tilde{S})$ be the stochastic matrix with entries

$$\tilde{P}(x, y) = \begin{cases} G(x, y), & x, y \in S \\ 1 - \sum_{y \in S} G(x, y), & x \in S, y = \Delta \\ 1, & x = y = \Delta. \end{cases}$$

Let $Y = (Y_n : n \geq 0)$ be the $\tilde{S}$-valued Markov chain having transition matrix $\tilde{P}$. Select a *return state* $z \in S$, and let $\tau = \inf\{n \geq 1 : Y_n = z\}$ be the first time $Y$ enters $z$. Also, let $T = \inf\{n \geq 0 : Y_n = \Delta\}$ be the *first hitting time* of $\Delta$. As in Section 2, let $\tilde{E}_x(\cdot)$ be the expectation operator on the path-space of $Y$ under which $Y_0 = x$. For $x \in S$ and $\beta > 0$, put

$$v(x) = \tilde{E}_x \beta^{-\tau} \mathbf{I}(\tau < T), \tag{8}$$

where $\mathbf{I}(A)$ is an indicator random variable that is one or zero depending upon whether or not $A$ occurs. By conditioning on $Y_1$, we find that for $x \in S$,

$$v(x) = \beta^{-1} \sum_{y \neq z} G(x, y) v(y) + \beta^{-1} G(x, x). \tag{9}$$

Suppose we select $\beta_*$ so that

$$v(z) = \tilde{E}_z \beta_*^{-\tau} \mathbf{I}(\tau < T) = 1. \tag{10}$$

Then, equation (9) may be written as

$$v(x) = \beta_*^{-1} \sum_y G(x, y) v(y)$$

or, equivalently,

$$\beta_* v = Gv.$$

In other words, if $\beta_*$is the root of equation (10), then necessarily $\beta_* = \lambda$ and $v = u^*$. Consequently, equations (8) and (10) together provide a probabilistic representation of the solution of our eigenvalue problem.

We now turn to describing an adaptive Monte Carlo algorithm that takes advantage of this representation. As in Section 2, the algorithm involves use of importance sampling. In particular, given a current strictly positive *best guess* $u$ for the eigenvector $u^*$, let $Q(u) = (Q(x, y, u) : x, y \in \tilde{S})$ be the transition matrix with the entries given by

$$Q(x, y, u) = \begin{cases} \frac{P(x,y)u(y)}{\sum_{x \in S} P(x,z)u(z)}, & x, y \in S \\ 1, & x = y = \Delta. \end{cases} \tag{11}$$

Rather than simulate $Y$ under the dynamics of $\tilde{P}$, we instead simulate $Y$ up to $T \wedge \tau (\stackrel{\Delta}{=} min(T, \tau))$ using the transition matrix $Q(u)$ and adjust the simulation output using the appropriate likelihood ratio:

$$\tilde{R}(\beta) = \beta^{-\tau} \mathbf{I}(\tau < T) \prod_{i=1}^{\tau - 1} \frac{P(Y_i, Y_{i+1})}{Q(Y_i, Y_{i+1}, u)}.$$

By replicating $\tilde{R}(\beta)$ $k$ iid times, conditional on $Y_0 = x$, we can estimate $v_\beta(x) \stackrel{\Delta}{=} \tilde{E}_x \beta^{-\tau} \mathbf{I}(\tau < T)$ via the corresponding sample mean. In particular for $x = z$, we can compute the empirical root $\lambda_n$ to estimate the root $\lambda$ satisfying $\tilde{E}_z \lambda^{-\tau} \mathbf{I}(\tau < T) = 1$. Having obtained $\lambda_n$ we can then simulate $\tilde{R}(\lambda_n)$ for each of the other states $x \neq z$ in $S$, so that the $\tilde{R}(\lambda_n)$'s are (conditional on $\lambda_n$) unbiased estimators for $v_{\lambda_n}(x), x \in S$. These estimators then form our *best guess* of the eigenvector $u^*$ for the next iteration of the algorithm.

As in Section 2, the sequence of *best guesses* $X = (X_n : n \geq 0)$ is a general state space Markov chain. The state space of $X$ is the space in which the candidate *best guess* solutions live, namely $\mathbb{R}_+^d$ (where $d$ is the number of states in $S$).

Again, as in Section 2, one hopes that $X_n$ converges to $u^*$ as $n \to \infty$ and that $u^*$ is a *fixed point* of the algorithm.

**Proposition 2.** *Suppose that $u^*$ is the unique eigenvector solving the eigenvalue problem for which $u^*(z) = 1$. Then, if $X_0 = u^*$, it follows that $X_1 = u^*$ a.s.*

**Proof** If $X_0 = u^*$ (with $u^*(z) = 1$), then $\sum_{z\in S} P(x,z)u^*(z) = \lambda u^*(x)$ for $x \in S$. Consequently, $Q(x,y,u^*) = (P(x,y)u^*(y))/(\lambda u^*(x))$ for $x, y \in S$. Noting that $T < \tau$ is impossible under $Q(u^*)$, it follows that

$$\begin{aligned}
\tilde{R}(\beta) &= \beta^{-\tau}\mathbf{I}(\tau < T)\prod_{i=0}^{\tau-1}\frac{P(Y_i,Y_{i+1})}{Q(Y_i,Y_{i+1},u^*)} \\
&= \beta^{-\tau}\prod_{i=0}^{\tau-1}(\frac{\lambda u^*(Y_i)}{u^*(Y_{i+1})}) \\
&= (\beta/\lambda)^{-\tau}u^*(Y_0)/u^*(Y_\tau) \\
&= (\beta/\lambda)^{-\tau}u^*(Y_0)/u^*(z) \\
&= (\beta/\lambda)^{-\tau}u^*(Y_0).
\end{aligned}$$

Since $u^*(z) = 1$, the empirical root $\lambda_n$ estimated by simulating $Y$ under $Q(u^*)$ with initial condition $Y_0 = z$ will equal $\lambda$ a.s. Thus, the random variables

$$\begin{aligned}
\tilde{R}(\lambda_n) &= (\lambda_n/\lambda)^{-\tau}u^*(Y_0) \\
&= u^*(Y_0)
\end{aligned}$$

generated under $Q(u^*)$ for all the other initial states $x \in S$ will also almost surely equal $u^*(Y_0)$, proving that the next *best guess* $X_1$ will again equal $u^*$. □

Thus, our second adaptive algorithm again has the characteristic that the desired solution $u^*$ is a *fixed point* of the corresponding general state space Markov chain $X$.

## 4 A MARKOV CHAIN PERSPECTIVE

In Sections 2 and 3, we provided two examples of adaptive algorithms in which the algorithmic progress, as a function of the iteration count, can be viewed as a general state space Markov chain $X = (X_n : n \geq 0)$. In addition, we showed that for both examples, the desired solution $u^*$ is a *fixed point* of the chain $X$.

Typically, in discussing the long-run behavior of a Markov chain, the key concept is that of a stationary distribution or *steady-state* distribution. However, the notion of a *fixed point* is easily expressed in terms of such distributions. In particular, asserting that $u^*$ is a fixed point for $X$ (in the sense that if $X_0 = u^*$, then $X_1 = u^*$) is equivalent to requiring that $\delta_{u^*}$ be a stationary distribution for $X$, where $\delta_x$ is a unit point mass distribution at $x$. Thus, the chains $X$ that typically arise in the analysis of adaptive Monte Carlo algorithms have the interesting (and unusual) property that they possess deterministic stationary distributions.

Given that $\delta_{u^*}$ is the stationary distribution for $X$, we can now proceed to hypothesize the likely implications for the chain. In substantial generality, nicely behaved general state space Markov chains are aperiodic and geometrically ergodic; see Chapter 15 of Meyn and Tweedie (1993) for details. In particular, if $\pi$ is the associated stationary distribution for $X$ and if $P_x(\cdot)$ is the probability on the path-space of $X$ under which $X_0 = x$, then such chains have the characteristic that there exists $\rho \in (0,1)$ and $a > 0$ such that

$$\|P_x(X_n \in \cdot) - \pi(\cdot)\| \leq a\rho^n$$

for $n \geq 1$, where $\|\cdot\|$ is the so-called *total variation* norm on the space of probabilities. (See Meyn and Tweedie (1993) for the definition.) In our current setting, the stationary distribution is the probability $\delta_{u^*}$, so this suggests that we can reasonably expect

$$\|P_x(X_n \in \cdot) - \delta_{u^*}(\cdot)\| \leq a\rho^n \tag{12}$$

to hold for $n \geq 1$. The relation (12) states that the distribution of $X_n$ converges to a point mass at $u^*$at a geometrically fast rate as $n \to \infty$.

This conclusion is very suggestive of what is currently known about adaptive algorithms. Specifically both Kollman et al. (1996) and Desai (2001) have shown that their respective algorithms generate a sequence of $X_n$'s that converge almost surely to $u^*$ at a geometric rate. As a consequence, one might reasonably expect that a natural approach to establishing such geometric convergence results is to verify the geometric ergodicity of $X$, and to then invoke the general theory that is known for such chains.

Unfortunately, while such an approach seems reasonable, it cannot be applied to the chains that arise from the analysis of adaptive algorithms. The problem is that the Markov chain $X$ is essentially never geometrically ergodic. In fact, such chains are essentially never Harris recurrent. This notion of recurrence, which we shall discuss further momentarily, is much more general than that of geometric ergodicity. Thus, the almost sure exponential convergence verified by Kollman et al. (1996) and Desai (2001) must be established via very different methods of proof. The need to look for alternative proof methods in analyzing adaptive algorithms is one of the main messages of the current paper.

**Definition 4.1.** *A $\mathbb{R}_+^d$-valued Markov chain $X = (X_n : n \geq 0)$ is said to be* Harris recurrent *if there exists a set $A$, a positive number $\lambda$, an integer $m > 1$, and a probability $\varphi$ such that:*

1. $P_x(X_n \in A \text{ infinitely often}) = 1$ *for* $x \in \mathbb{R}_+^d$;
2. $P_x(X_m \in \cdot) \geq \lambda\varphi(\cdot)$ *for* $x \in A$.

If such a Harris recurrent Markov chain $X$ has a stationary distribution $\pi$, it must be that for $x \in \mathbb{R}_+^d$,

$$\|n^{-1}\sum_{j=0}^{n-1}P_x(X_j \in \cdot) - \pi(\cdot)\| \to 0$$

as $n \to \infty$; see Meyn and Tweedie (1993). In particular, if $\pi = \delta_{u^*}$, then

$$n^{-1} \sum_{j=0}^{n-1} P_x(X_j = u^*) \to 1 \tag{13}$$

as $n \to \infty$. Let $\Gamma = \inf\{n \geq 0 : X_n = u^*\}$. Relation (13) implies that $P_x(\Gamma < \infty) = 1$ for $x \in \mathbb{R}_+^d$, so that $u^*$ is visited in finite time by $X$. Of course, once $u^*$ is visited, $X$ remains in $u^*$. Thus, if the chain $X$ is Harris recurrent, evidently the associated adaptive algorithm converges to the exact solution $u^*$ in a finite number of iterations. Intuitively, however, it seems clear that such adaptive algorithms only converge to $u^*$ in the limit.

To verify rigorously that $X$ is not Harris recurrent on $\mathbb{R}_+^d$, note that the support of $X_1$ is a countably infinite discrete set $D_1(x)$, depending on the initial state $X_0 = x$. If $D_\infty(x)$ is the union (over $n \in \{1, 2, \ldots\}$) of the supports of the random vectors $X_n$, $D_\infty(x)$ is also a countably infinite set. If $P_x(\Gamma < \infty) = 1$ for $x \in \mathbb{R}_+^d$, then $u^* \in D_\infty(x)$ for $x \in \mathbb{R}_+^d$. But it is easily seen that there exist distinct points $x_1, x_2 \in \mathbb{R}_+^d$ for which $D_\infty(x_1)$ and $D_\infty(x_2)$ are disjoint. This contradiction implies that $X$ can not be a Harris recurrent chain.

Given that the chain $X$ is not Harris recurrent, it is natural to search for alternative proof techniques for establishing the desired almost sure convergence of $X_n$ to $u^*$. We now discuss one such alternative.

By using Lyapunov function methods, one might hope that one could establish that $X$ visits every $\epsilon$-neighborhood of $u^*$ almost surely infinitely often; see Appendix B of Meyn and Tweedie (1993) for details on Lyapunov functions and their applicability to such a problem. Furthermore, $u^*$ is a *fixed point* for the chain $X$. It therefore appears reasonable to expect that if the transition dynamics of $X$ are suitably continuous in a neighborhood of $u^*$, then $X$ should converge to $u^*$.

**Definition 4.2.** *A $\mathbb{R}_+^d$-valued Markov chain $X = (X_n : n \geq 0)$ is said to be* weakly continuous *if $P_{x_n}(X_1 \in \cdot) \Longrightarrow P_x(X_1 \in \cdot)$ whenever $x_n \to x \in \mathbb{R}_+^d$ (where $\Longrightarrow$ denotes weak convergence).*

In other words, we would like to conclude that if $X = (X_n : n \geq 0)$ is a weakly continuous Markov chain on $\mathbb{R}_+^d$ for which $P_x(|\, X_n - u^* \,| < \epsilon \text{ infinitely often}) = 1$ for $x \in \mathbb{R}_+^d$ with $\delta_{u^*}$ a stationary distribution, then $X_n \to u^*$ a.s. as $n \to \infty$. Unfortunately, as we shall see in the example below, this conclusion is invalid in general. Thus, we again conclude that different tools are needed to establish convergence for such algorithms.

**Example 1**

Consider the Markov chain $X = (X_n : n \geq 0)$ on $[0, \infty)$ with the following transition structure:

$$P_x(X_1 \in \cdot) = (1 - p(x))\delta_{x/2}(\cdot) + p(x)\delta_1(\cdot)$$

for $x > 0$, and

$$P_0(X_1 \in \cdot) = \delta_0(\cdot).$$

Assume that $0 \leq p(x) \leq 1$ and that $p(\cdot)$ is continuous on $[0, \infty)$ with $p(0) = 0$. Then, $\delta_0$ is a stationary distribution for $X$ and $X$ is weakly continuous on $\mathbb{R}_+^d$.

Suppose $X$ visits 1 infinitely often. On such sample paths, $X_n \to 0$ a.s. as $n \to \infty$, so that $X$ visits each $\epsilon$-neighborhood of the origin. On the other hand, for $\epsilon > 0$, there exists $n = n(\epsilon)$ such that $2^{-n} < \epsilon$, so that $P_1(X_n < \epsilon) > 0$. If $X$ visits 1 infinitely often, a *geometric trials* argument therefore establishes that $X$ will eventually visit the $\epsilon$-neighborhood. It follows that for each initial state $x \in \mathbb{R}_+$, $X$ visits the $\epsilon$-neighborhood of the origin infinitely often.

We now show that $X_n$ fails to converge to 0 a.s. as $n \to \infty$. To see this, we apply the conditional Borel-Cantelli lemma (p. 323 of Doob (1953)). Note that the conditional Borel-Cantelli lemma ensures that $\sum_{n=1}^\infty \mathbf{I}(X_n = 1)$ converges/diverges on precisely the same set of sample outcomes as does $\sum_{n=1}^\infty P_x(X_n = 1 \mid X_{n-1})$. But $\sum_{n=1}^\infty \mathbf{I}(X_n = 1) < \infty$ implies that $L$, the last time at which $X$ visits 1, is finite. Also,

$$\begin{aligned} \sum_{n=1}^\infty P_x(X_n = 1 \mid X_{n-1}) &\geq \sum_{n=L}^\infty P_x(X_n = 1 \mid X_{n-1}) \\ &= \sum_{n=0}^\infty p(2^{-n} X_L). \end{aligned}$$

So, if we choose $p(\cdot)$ so that $p(x) \downarrow 0$ with $x \downarrow 0$ in such a way that $\sum_{n=0}^\infty p(2^{-n}) = \infty$ (e.g. $p(2^{-j}) = 1/j$), the latter sum diverges and it follows that $P_x(X_n = 1 \mid X_{n-1})$ is diverging on a set of outcomes on which $\sum_{n=1}^\infty \mathbf{I}(X_n = 1)$ is converging. This contradiction of the conditional Borel-Cantelli lemma forces us to conclude that $\sum_{n=1}^\infty \mathbf{I}(X_n = 1) = \infty$ a.s. Hence, we conclude that $X_n$ does not converge almost surely to zero.

The above counterexample shows that our proposed proof technique fails. This example illustrates the need for the set of special tools developed in Kollman et al. (1996) and Desai (2001) to analyze adaptive Monte Carlo algorithms. The reader is referred to those works for a full exposition of the theory.

## REFERENCES

Booth, T. E. 1985. Exponential convergence for Monte Carlo particle transport. *Transactions of the American Nuclear Society* 50: 267–268

Desai, P. Y. 2001. Adaptive Monte Carlo methods for solving eigenvalue problems. Ph.D. thesis, Department of Management Science and Engineering, Stanford University, Stanford, California.

Desai, P. Y. and Glynn, P. W. 2001. Adaptive Monte carlo algorithms and general state space Markov chains. Unpublished working paper, Stanford University, Stanford CA.

Doob, J. L. 1953. *Stochastic Processes.* John Wiley and Sons, New York.

Glynn, P. W., and Iglehart D. L. 1989. Importance sampling for stochastic simulations. *Mangement Science* 35: 1367–1392

Glynn, P. W. 1996. Numerical computation of large deviations exponents via simulation. *Allerton Conference.*

Kollman, C., Baggerly, K., Cox, D., Picard, R. 1996. Adaptive importance sampling on discrete Markov chains. *Technical Report LA-UR-96-3998, Los Alamos National Laboratory.*

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

Seneta, E. 1973. *Non-Negative Matrices; An Introduction to Theory and Applications* Wiley, New York.

## AUTHOR BIOGRAPHIES

**PARITOSH Y. DESAI** received his Ph.D. in Operations Research from Stanford University in 2001. His email address is <paritosh.desai@stanfordalumni.org>. His research interests include simulation theory, stochastic processes and applied probability.

**PETER W. GLYNN** received his Ph.D. from Stanford University, after which he joined the faculty of the Department of Industrial Engineering at the University of Wisconsin-Madison. In 1987, he returned to Stanford, where he is the Thomas Ford Professor of Engineering in the Department of Management Science and Engineering. He was a co-winner of the 1993 Outstanding Simulation Publication Award sponsored by the TIMS College on Simulation. His research interests include discrete-event simulation, computational probability, queueing, and general theory for stochastic systems.