

CELEBRATING 50 YEARS OF THE APPLIED PROBABILITY TRUST

Edited by
S. ASMUSSEN, P. JAGERS, I. MOLCHANOV and L. C. G. ROGERS

Part 8. Markov processes and renewal theory

EXACT ESTIMATION FOR MARKOV CHAIN EQUILIBRIUM EXPECTATIONS

PETER W. GLYNN, *Stanford University*

Management Science and Engineering, Stanford University, Stanford, CA 94305-4121, USA.

Email address: glynn@stanford.edu

CHANG-HAN RHEE, *Georgia Institute of Technology*

Biomedical Engineering, Georgia Institute of Technology, U. A. Whitaker Building, 313 Ferst Drive, Atlanta, GA 30332, USA. Email address: rhee@gatech.edu



APPLIED PROBABILITY TRUST
DECEMBER 2014

EXACT ESTIMATION FOR MARKOV CHAIN EQUILIBRIUM EXPECTATIONS

BY PETER W. GLYNN AND CHANG-HAN RHEE

Abstract

We introduce a new class of Monte Carlo methods, which we call exact estimation algorithms. Such algorithms provide unbiased estimators for equilibrium expectations associated with real-valued functionals defined on a Markov chain. We provide easily implemented algorithms for the class of positive Harris recurrent Markov chains, and for chains that are contracting on average. We further argue that exact estimation in the Markov chain setting provides a significant theoretical relaxation relative to exact simulation methods.

Keywords: Unbiased estimation; Markov chain equilibrium expectation; Markov chain stationary expectation; exact estimation; exact sampling; exact simulation; perfect sampling; perfect simulation

2010 Mathematics Subject Classification: Primary 65C05
Secondary 60J05

1. Introduction

A key advance in the development of Monte Carlo algorithms for Markov chains has been the introduction of what are known as *exact simulation* or *perfect simulation* algorithms for equilibrium distributions of such stochastic processes. By assuming a suitable structure for the underlying Markov chain, one can construct an algorithm that draws samples perfectly from the equilibrium distribution of the Markov chain, based only on an ability to generate sample paths of the chain starting from an arbitrary state. In particular, such algorithms have been developed for finite-state Markov chains (see Propp and Wilson (1996)), uniformly recurrent Markov chains (see Asmussen *et al.* (1992)), certain stochastically monotone Markov chains (see Propp and Wilson (1996), and Corcoran and Tweedie (2001)), some queueing models (see Ensor and Glynn (2000), and Blanchet and Wallwater (2014)), and various subclasses of Harris recurrent Markov chains (see Kendall (2004), and Connor and Kendall (2007)). While the idea is powerful, it apparently requires exploiting significant structure within the Markov chain itself. In particular, no universal and practically implementable such perfect sampler has been constructed for Harris recurrent Markov chains, or even for countable state positive recurrent Markov chains. In fact, we shall argue below in Section 2 that exact simulation is inherently restrictive, in the sense that such algorithms can typically be constructed only for Markov chains that are ϕ -irreducible.

In this paper we relax the algorithmic formulation so as to require only that the algorithm should produce unbiased estimators for equilibrium expectations, rather than to demand (as in exact simulation) that such an unbiased estimator is necessarily constructed from an exact sample from the equilibrium distribution. To differentiate this new class of algorithms from exact simulation algorithms, we shall refer to them as *exact estimation* algorithms. Our exact estimation algorithms exploit a recent idea of Rhee and Glynn (2013) that showed how

unbiased estimators can often be constructed from a sequence of biased approximations; see also McLeish (2011), and Rhee and Glynn (2012). As we shall see below, this new class of algorithms can provide unbiased estimators for equilibrium expectations for any positive recurrent Harris chain. In fact, we shall see that exact estimation algorithms can even be developed for non- ϕ -irreducible Markov chains, provided that the Markov chain is ‘contractive’ in a certain sense and that the equilibrium expectation involves a suitably Lipschitz functional (Theorem 2.1). This makes clear that exact estimation is indeed a significant relaxation of exact simulation.

One key additional feature of our exact estimation algorithms is that unlike most existing exact simulation methods, our algorithms do not involve explicitly simulating paths from multiple initial conditions, nor do they require monotonicity. Furthermore, our exact estimation procedures can easily be implemented with a minimal need to store sample paths. Consequently, our proposed exact estimation methods are relatively straightforward to implement.

Our paper is organized as follows. In Section 2 we illustrate exact estimation in the setting of contractive Markov chains, and exploit the fact that Markov chains can, in great generality, be viewed as a sequence of random iterated functions. In Section 3 we develop exact estimation algorithms in the context of Harris recurrent Markov chains, thereby establishing that our proposed relaxation is indeed a generalization of exact simulation. In Section 4 we prove a variant of the Glivenko–Cantelli theorem for our newly developed estimator, and in Section 5 we provide a brief computational discussion.

2. Exact estimation for contracting Markov chains

Given an S -valued Markov chain $X = \{X_n : n \geq 0\}$, we wish to develop exact estimation algorithms for computing $\mathbb{E}f(X_\infty)$, where f is real valued and X_∞ has the equilibrium distribution π of X (assumed to exist uniquely). We start by briefly describing and (slightly) generalizing the algorithms and analysis presented in Rhee and Glynn (2012, 2013), and McLeish (2011), following related earlier work by Rychlik (1990, 1995).

Suppose that we wish to compute $\mathbb{E}Y$. We assume that we have available to us a sequence $\{Y_k : k \geq 0\}$ of approximations for which

$$\mathbb{E}Y_k \rightarrow \mathbb{E}Y \quad \text{as } k \rightarrow \infty. \tag{2.1}$$

If $\{\Delta_k : k \geq 0\}$ is a sequence of random variables (RVs) for which $\mathbb{E}\Delta_k = \mathbb{E}(Y_k - Y_{k-1})$ for $k \geq 0$ (with $Y_{-1} := 0$) and

$$\mathbb{E} \sum_{k=0}^{\infty} |\Delta_k| < \infty, \tag{2.2}$$

then it is easy to verify that

$$Z := \sum_{k=0}^N \frac{\Delta_k}{\mathbb{P}\{N \geq k\}}$$

is an unbiased estimate of $\mathbb{E}Y$, provided that N is a \mathbb{Z}_+ -valued RV independent of $\{\Delta_k : k \geq 0\}$. Of course, one implication of (2.2) is that $\Delta_k \Rightarrow 0$ as $k \rightarrow \infty$, where ‘ \Rightarrow ’ denotes weak convergence. Note that in our Markov chain setting, the most natural choice of approximating sequence $\{Y_k : k \geq 0\}$ is to choose $Y_k = f(X_k)$ (with $Y = f(X_\infty)$). However, the obvious choice for Δ_k , namely, $\Delta_k = f(X_k) - f(X_{k-1})$, then fails to satisfy $\Delta_k \Rightarrow 0$ as $k \rightarrow \infty$. Thus, the key to the development of an exact estimation algorithm for X is the construction of a computationally implementable coupling between X_{k-1} and X_k that forces Δ_k to converge

to 0 (hopefully rapidly); see Giles (2008) for the development of a closely related (biased) estimator.

We now illustrate one potential coupling that can be applied in the setting of contractive Markov chains. Assume (for the purpose of this section) that S is a complete separable metric space with metric $\rho: S \times S \rightarrow \mathbb{R}_+$. In this setting, X can be represented in terms of a sequence of independent and identically distributed (i.i.d.) random functions $\{\varphi_i: i \geq 1\}$ mapping S into S , independent of X_0 , so that

$$X_i = \varphi_i(X_{i-1})$$

for $i \geq 1$; see, e.g. Borovkov and Foss (1992). For example, a standard representation of such Markov chains is in terms of a stochastic recursion in which $X_i = \psi(X_{i-1}, \beta_i)$ for some i.i.d. sequence $\{\beta_i: i \geq 1\}$ and deterministic function ψ , so that $\varphi_i(x)$ may be taken as $\psi(x, \beta_i)$. In any case, given this random function representation, $X_n = (\varphi_n \circ \varphi_{n-1} \circ \dots \circ \varphi_1)(X_0)$, conditional on X_0 . An obvious means of coupling X_{n-1} and X_n is then to set

$$\tilde{X}_{n-1} = (\varphi_n \circ \varphi_{n-1} \circ \dots \circ \varphi_2)(x)$$

for $n \geq 2$, with $\tilde{X}_0 := x$. Clearly, $\tilde{X}_{n-1} \stackrel{D}{=} X_{n-1}$ for $n \geq 1$, where ‘ $\stackrel{D}{=}$ ’ denotes equality in distribution.

For $y, z \in S$, set $r(y, z) = \mathbb{E}\rho^2(\varphi_1(y), \varphi_1(z))$ and assume that there exists $b < 1$ for which

$$r(y, z) \leq b\rho^2(y, z) \tag{2.3}$$

for $y, z \in S$, so that X is ‘contractive on average’. Suppose further that f is Lipschitz with respect to the metric ρ , so that there exists $\kappa < \infty$ for which

$$|f(y) - f(z)| \leq \kappa\rho(y, z) \tag{2.4}$$

for $y, z \in S$. Evidently,

$$\begin{aligned} \mathbb{E}\Delta_k^2 &\leq \kappa^2\mathbb{E}\rho^2(X_k, \tilde{X}_{k-1}) \\ &= \kappa^2\mathbb{E}\rho^2(\varphi_k(X_{k-1}), \varphi_k(\tilde{X}_{k-2})) \\ &\leq \kappa^2 b\mathbb{E}\rho^2(X_{k-1}, \tilde{X}_{k-2}) \\ &\leq \dots \\ &\leq \kappa^2 b^{k-1}\mathbb{E}\rho^2(X_1, x). \end{aligned}$$

Hence, if

$$\mathbb{E}\rho^2(\varphi_1(x), x) < \infty \tag{2.5}$$

for $x \in S$, it follows that $\mathbb{E}\Delta_k^2 \rightarrow 0$ geometrically fast, so that (2.2) holds. Also, Theorem 2.1 of Diaconis and Freedman (1999) applies, in the presence of (2.3) and (2.5), so that X has a unique stationary distribution π . In fact, their proof makes clear that $\mathbb{E}f(X_n) \rightarrow \mathbb{E}f(X_\infty)$ geometrically fast (where X_∞ has distribution π) when f is Lipschitz. Consequently, (2.1) is valid, thereby proving that Z is an unbiased estimate of $\mathbb{E}f(X_\infty)$.

But more can be said. Note that

$$\mathbb{E}Z^2 = \sum_{k=0}^{\infty} \frac{\mathbb{E}\Delta_k^2 + 2 \sum_{j=k+1}^{\infty} \mathbb{E}\Delta_k \Delta_j}{\mathbb{P}\{N \geq k\}}. \tag{2.6}$$

By virtue of the Cauchy–Schwarz inequality, it follows that $\text{var } Z < \infty$, provided that we choose the distribution for N so that

$$\sum_{k=0}^{\infty} \frac{b^k}{\mathbb{P}\{N \geq k\}} < \infty. \tag{2.7}$$

Under condition (2.7), the central limit theorem asserts that if we generate i.i.d. copies Z_1, Z_2, \dots of Z and form the sample mean $\bar{Z}_n = (Z_1 + \dots + Z_n)/n$, then \bar{Z}_n converges weakly to $\mathbb{E}f(X_\infty)$ at rate $n^{-1/2}$ in the number n of samples that are generated.

Of course, the amount of computer time needed to generate each Z_i could be excessive. To take this effect into account, we let ξ_i be the computer time needed to generate Z_i . In view of the fact that computing Z requires us to generate $\varphi_1, \dots, \varphi_N$, it seems natural to assess the computer time as being equal to N . Hence, we put $\xi_i = N_i$, where N_i is the corresponding randomization RV N associated with Z_i . If $\Gamma(c)$ is the number of Z_i generated in c units of computer time, $\Gamma(c) = \max\{k \geq 0: \xi_1 + \dots + \xi_k \leq c\}$ and the estimator for $\mathbb{E}f(X_\infty)$ available after expending c units of computer time is $\bar{Z}_{\Gamma(c)}$. If

$$\mathbb{E}\xi_i = \mathbb{E}N = \sum_{k=0}^{\infty} \mathbb{P}\{N \geq k\} < \infty \tag{2.8}$$

and $\text{var } Z < \infty$, it is well known (see, for example, Glynn and Whitt (1992)) that

$$c^{1/2}(\bar{Z}_{\Gamma(c)} - \mathbb{E}f(X_\infty)) \Rightarrow \sqrt{\mathbb{E}N \cdot \text{var } Z} N(0, 1) \text{ as } c \rightarrow \infty,$$

where $N(0, 1)$ denotes a normal RV with mean 0 and variance 1.

We summarize our discussion with the following result, which establishes that the exact estimation algorithm exhibiting ‘square root’ convergence rates can be obtained for any suitably contractive chain and Lipschitz function f .

Theorem 2.1. *Assume that (2.3)–(2.5) hold. If the distribution of N is chosen so that (2.7) and (2.8) hold (e.g. $\mathbb{P}\{N \geq k\} = ck^{-\alpha}$ for $\alpha > 1$), then*

$$c^{1/2}(\bar{Z}_{\Gamma(c)} - \mathbb{E}f(X_\infty)) \Rightarrow \sqrt{\mathbb{E}N \cdot \text{var } Z} N(0, 1) \text{ as } c \rightarrow \infty.$$

In view of Theorem 2.1, a natural question arises as to the optimal choice for the distribution of N . According to Proposition 1 of Rhee and Glynn (2013), the optimal choice is to set

$$\mathbb{P}\{N \geq k\} = \left(\frac{\mathbb{E}\Delta_k^2 + 2 \sum_{j=k+1}^{\infty} \mathbb{E}\Delta_k \Delta_j}{\mathbb{E}\Delta_0^2 + 2 \sum_{j=1}^{\infty} \mathbb{E}\Delta_0 \Delta_j} \right)^{1/2} \tag{2.9}$$

for $k \geq 0$, provided that the right-hand side is nonincreasing; see Theorem 3 of Rhee and Glynn (2013) for details of the form of the optimal distribution when the right-hand side fails to be nonincreasing. Given this result and the geometric decay of the $\mathbb{E}\Delta_k^2$, it therefore seems reasonable to expect that requiring the tail of N to be geometric will often be a good choice in this setting.

We turn next to a slightly different implementation of our coupling idea in the S -valued metric space contractive setting. Given the independence of N from the φ_i , an alternative coupling for (X_{i-1}, X_i) is to set

$$X_j^* = (\varphi_N \circ \varphi_{N-1} \circ \dots \circ \varphi_{N-j+1})(x)$$

for $0 \leq j \leq N$. Clearly, conditional on N , $X_j^* \stackrel{D}{=} X_j$ and $\Delta_j^* := X_j^* - X_{j-1}^* \stackrel{D}{=} \Delta_j$ for $j \geq 0$. Because $\mathbb{E}(\Delta_j^*)^2 = \mathbb{E}\Delta_j^2$ for $j \geq 0$, the same argument as for Z shows that

$$Z^* = \sum_{j=0}^N \frac{\Delta_j^*}{\mathbb{P}\{N \geq j\}}$$

is unbiased for $\mathbb{E}f(X_\infty)$. Furthermore, the estimator for $\mathbb{E}f(X_\infty)$ corresponding to computing a sample average of i.i.d. copies of Z^* satisfies Theorem 2.1, under conditions (2.3)–(2.8).

The estimator based on Z^* is slightly more complicated to implement, because X_i^* cannot be computed recursively from X_{i-1}^* in this setting (whereas (X_i, \tilde{X}_i) can be computed recursively from $(X_{i-1}, \tilde{X}_{i-1})$). Of course, the estimator based on Z^* will have a different variance than does Z , because $\mathbb{E}\Delta_k\Delta_j \neq \mathbb{E}\Delta_k^*\Delta_j^*$ for $k < j$. In particular, while all four of the quantities $f(X_k^*), f(X_{k-1}^*), f(X_j^*), f(X_{j-1}^*)$ appearing in $\mathbb{E}\Delta_k^*\Delta_j^*$ will be close to one another when k is large, $f(X_j) - f(X_k)$ will exhibit significant variability, regardless of the magnitude of k .

We close this section by noting that the construction of an exact simulation algorithm typically requires the underlying Markov chain to be ϕ -irreducible. Since contractive chains need not be ϕ -irreducible, this discussion serves to illustrate the fact that exact estimation does indeed cover Markov chains to which the theory of exact simulation does not apply.

Recall that an S -valued Markov chain $X = \{X_n : n \geq 0\}$ is ϕ -irreducible if there exists a σ -finite (nonnegative) measure ϕ such that whenever $\phi(A) > 0$ for some (measurable) A ,

$$R(x, A) := \sum_{n=0}^{\infty} 2^{-n} \mathbb{P}_x\{X_n \in A\} > 0$$

for all $x \in S$, where $\mathbb{P}_x\{\cdot\} := \mathbb{P}\{\cdot \mid X_0 = x\}$. Equivalently,

$$\phi(\cdot) \ll R(x, \cdot)$$

for each $x \in S$, where ‘ \ll ’ denotes ‘is absolutely continuous with respect to’.

A typical exact simulation algorithm involves simulating X from multiple initial states x_1, x_2, \dots , thereby yielding a family of RVs $\{X_{ij} : i \geq 1, j \geq 0\}$ such that

$$\mathbb{P}\{\{X_{ij} : j \geq 0\} \in \cdot\} = \mathbb{P}\{\{X_j : j \geq 0\} \in \cdot \mid X_0 = x_i\}$$

for $i \geq 1$; the exact simulation algorithm then outputs X_{IJ} for some appropriately chosen pair of RVs (I, J) . Exact simulation demands that if X has an equilibrium distribution $\pi(\cdot)$ then

$$\mathbb{P}\{X_{IJ} \in \cdot\} = \pi(\cdot).$$

The probability $\mathbb{P}\{X_{IJ} \in \cdot\}$ is mutually absolutely continuous with respect to

$$\mathbb{E}[2^{-I-J} \mathbf{1}\{X_{IJ} \in \cdot\}].$$

But

$$\begin{aligned} \mathbb{E}[2^{-I-J} \mathbf{1}\{X_{IJ} \in \cdot\}] &= \sum_{i=1}^{\infty} \sum_{j=0}^{\infty} 2^{-i-j} \mathbb{P}\{X_{ij} \in \cdot, I = i, J = j\} \\ &\leq \sum_{i=1}^{\infty} \sum_{j=0}^{\infty} 2^{-i-j} \mathbb{P}\{X_{ij} \in \cdot\} \\ &= \sum_{i=1}^{\infty} 2^{-i} R(x_i, \cdot). \end{aligned}$$

Hence, it follows that

$$\pi(\cdot) \ll \sum_{i=1}^{\infty} 2^{-i} R(x_i, \cdot). \tag{2.10}$$

Thus, the existence of an exact simulation algorithm requires one to have *a priori* knowledge of a set of states x_1, x_2, \dots satisfying (2.10). Without additional structure on the chain, the only way to guarantee this is to require that

$$\pi(\cdot) \ll R(x, \cdot) \quad \text{for each } x \in S.$$

In other words, X must be ϕ -irreducible, with $\phi = \pi$. (Note that when a ϕ -irreducible Markov chain has a stationary distribution, one choice for ϕ is always $\phi = \pi$.) Consequently, ϕ -irreducibility and exact simulation are tightly connected concepts.

3. Exact estimation for Harris recurrent Markov chains

In this section we prove that exact estimation algorithms can be constructed for any positive recurrent Harris chain, under the assumption that S is a separable metric space. In the presence of such separability, it is well known that there exists a so-called small set A , i.e. there exists $m \geq 1$, a probability ν , and $\lambda > 0$, for which

$$\mathbb{P}\{X_m \in \cdot \mid X_0 = x\} \geq \lambda\nu(\cdot) \tag{3.1}$$

for all $x \in A$. Given (3.1), we can express the m -step transition probability on A via the mixture representation

$$\mathbb{P}\{X_m \in \cdot \mid X_0 = x\} = \lambda\nu(\cdot) + (1 - \lambda)Q(x, \cdot), \tag{3.2}$$

where $Q(x, \cdot)$ is a probability on S for each $x \in A$. In view of (3.2), we can construct regeneration times $T(1), T(2), \dots$ for X by first running the chain until it hits A . Once it hits A , at time T say, we distribute the chain at time $T + m$ according to ν with probability λ and according to $Q(X_T, \cdot)$ with probability $1 - \lambda$; we then ‘condition in’ the values of $X_{T+1}, \dots, X_{T+m-1}$, conditional on X_T and X_{T+m} . Then if we succeed in distributing X according to ν at time $T + m$, set $T(1) = T + m$; otherwise, continue simulating X forward from time $T + m$, and continue attempting to distribute X according to ν at successive visits to A until we are successful, thereby defining the first regeneration time $T(1)$. We then successively follow the same procedure from time $T(1)$ forward to construct $T(2), T(3), \dots$. The Markov chain X is *wide-sense regenerative* with respect to the sequence of random times $T(1), T(2), \dots$. In particular, the random element $((X_{T(i)+j}, T(i+j+1) - T(i)) : j \geq 0)$ is identically distributed and independent of $T(i)$ for $i \geq 1$; see Meyn and Tweedie (2009) and Thorisson (2000) for details. As noted in Asmussen and Glynn (2007), one can implement the above algorithm using acceptance/rejection so that explicit generation from the conditional distribution (given X_T and X_{T+m}) can be avoided.

We now explain our exact estimation algorithm in the special case that X is aperiodic; later we generalize to the periodic case. As in Section 2, the key is to construct a coupling of (X_{n-1}, X_n) that makes Δ_n small. Specifically, alongside $\{X_n : n \geq 0\}$ we construct another sequence $\{X'_n : n \geq 0\}$ and a random time τ such that $\{X'_n : n \geq 0\} \stackrel{D}{=} \{X_n : n \geq 0\}$, and then attempt to (distributionally) couple the X'_n to the X_{n+1} in such a way that $X_\tau \stackrel{D}{=} X'_{\tau-1}$.

Start by drawing X_0 from the distribution ν and set $X'_0 = X_0$. We have already discussed the simulation of X and the construction of the associated $T(n)$. Conditional on X_0 , simulate $\{X'_n : n \geq 1\}$ independently of $\{X_n : n \geq 1\}$, thereby producing an associated sequence of regeneration times $0 = T'(0) < T'(1) < T'(2) < \dots$. Now let the (distributional) coupling time τ be the first time at which one of the $(T(j) - 1)$ coincides with one of the $T'(i)$, specifically,

$$\tau = \inf\{T(n) : n \geq 1, \text{ there exists } m \geq 0 \text{ such that } T'(m) = T(n) - 1\}.$$

With this definition of τ ,

- (a) $\{X_{\tau+j} : j \geq 0\} \stackrel{D}{=} \{X'_{\tau+j-1} : j \geq 0\}$; and
- (b) $\{(X_{\tau+j}, X'_{\tau+j-1}) : j \geq 0\}$ is independent of τ .

Consequently, setting $\Delta_k = [f(X_k) - f(X'_{k-1})] \mathbf{1}\{\tau > k\}$ (with f bounded), it follows that

$$\begin{aligned} \mathbb{E}\Delta_k &= \mathbb{E}[(f(X_k) - f(X'_{k-1})) \mathbf{1}\{\tau > k\}] \\ &= \mathbb{E}[(f(X_k) - f(X'_{k-1})) \mathbf{1}\{\tau > k\}] + \sum_{j=1}^k \mathbb{E}[(f(X_{\tau+k-j}) - f(X'_{\tau+k-j-1})) \mathbb{P}\{\tau = j\}] \\ &= \mathbb{E}[(f(X_k) - f(X'_{k-1})) \mathbf{1}\{\tau > k\}] + \sum_{j=1}^k \mathbb{E}[(f(X_{\tau+k-j}) - f(X'_{\tau+k-j-1})) \mathbf{1}\{\tau = j\}] \\ &= \mathbb{E}[f(X_k) - f(X'_{k-1})] \\ &= \mathbb{E}[f(X_k) - f(X_{k-1})], \end{aligned}$$

where the second equality follows because of (a), and the third inequality is a consequence of independence as at (b). Furthermore, the aperiodicity of X and the boundedness of f imply that $\mathbb{E}f(X_k) \rightarrow \mathbb{E}f(X_\infty)$ as $k \rightarrow \infty$, where X_∞ has the distribution π , the unique stationary distribution of the Harris chain.

It remains to establish condition (2.2). Observe that the $(T(n) - 1)$ are the regeneration times for the sequence $\{X_{n+1} : n \geq 0\}$, in which X is initialized with the distribution $\mathbb{P}\{X_1 \in \cdot\}$. Equivalently, the $(T(n) - 1)$ are renewal times for the delayed renewal process in which the interrenewal times share the same interrenewal distribution as for the $T'(n)$, but in which the probability mass function for the initial delay is given by $\{q_j : j \geq 0\}$, where $q_j = \mathbb{P}\{T'(2) - T'(1) = j + 1\}$. Thus, τ is the first time that two independent aperiodic renewal processes couple, in which one is nondelayed (corresponding to $\{X'_n : n \geq 0\}$) and the other is delayed with initial delay $(q_j : j \geq 0)$. According to Lindvall (2002, p. 27), $\mathbb{E}\tau^r < \infty$ for $r \geq 1$, provided that $\mathbb{E}[T(2) - T(1)]^r < \infty$. (Note that $\sum_{j=0}^\infty j^r q_j < \infty = \mathbb{E}[T(2) - T(1)]^r$ for the specific delay distribution that arises here.) Of course, the positive recurrence of X implies that $\mathbb{E}[T(2) - T(1)] < \infty$ (see Athreya and Ney (1978)), implying that $\mathbb{E}\tau < \infty$. Hence,

$$\mathbb{E} \sum_{k=0}^\infty |\Delta_j| \leq 2 \|f\| \mathbb{E} \sum_{k=0}^\infty \mathbf{1}\{\tau > k\} = 2 \|f\| \mathbb{E}\tau < \infty,$$

where $\|f\| = \sup\{|f(x)| : x \in S\} < \infty$, validating (2.2). It follows that Z is an unbiased estimator for $\mathbb{E}f(X_\infty)$.

For the periodic case (with period p), we can apply the above algorithm to $(X_{pn} : n \geq 0)$, and apply the coupling τ to coupling the X_{pn} to the $X_{p(n-1)}$. (Note that, by setting $X_0 = X'_0$, we guarantee that both $\{X_n : n \geq 0\}$ and $\{X'_n : n \geq 0\}$ start in the same periodic subclass, so that $\{X_{p(n+1)} : n \geq 0\}$ can successfully couple with $\{X'_{pn} : n \geq 0\}$.) We summarize our discussion thus far in the following result.

Theorem 3.1. *If X is a positive recurrent Harris chain and f is bounded, the RV Z described above is an unbiased estimator for $\mathbb{E}f(X_\infty)$.*

Of course, this estimator may fail to exhibit a ‘square root convergence rate’, because Z may not have finite variance and the expected computation time to generate Z may be infinite. However, we note that because $\Delta_k = 0$ for $k > \tau$ in this setting, the number of time steps of $\{(X_j, X'_j) : j \geq 0\}$ that need to be simulated in order to compute Z is bounded by $2 \min\{\tau, N\}$. (The factor of 2 appears because we need to simulate both the X_j and X'_j .) Hence, if ξ is a measure of the computational effort required to generate Z , $\mathbb{E}\xi$ is automatically finite because $\mathbb{E}\tau < \infty$, regardless of the distribution of N . (In fact, we may set $N = \infty$ almost surely (a.s.) in this setting, if we so wish.)

Turning now to the variance of Z , we note that if f is bounded, $\mathbb{E}\Delta_k^2 = O(\mathbb{P}\{\tau > k\})$ as $k \rightarrow \infty$. Furthermore, if $\mathbb{E}[T(2) - T(1)]^r < \infty$ for $r > 1$, then $\mathbb{E}\tau^r < \infty$, so that $\mathbb{E}\Delta_k^2 = O(k^{-r})$ by virtue of Markov’s inequality. Also, for $k < j$, $\Delta_k \Delta_j = 0$ unless $\tau > j$, so $\mathbb{E}\Delta_k \Delta_j = O(j^{-r})$ as $j \rightarrow \infty$, uniformly in k . Thus,

$$\sum_{k=j+1}^{\infty} \mathbb{E}\Delta_k \Delta_j = O(k^{1-r}) \quad \text{as } k \rightarrow \infty.$$

In order that there exist a probability distribution N so that $\mathbb{E}Z^2 < \infty$, (2.6) implies that it is therefore sufficient that $\alpha > 2$ (in which case we can, for example, choose N so that $\mathbb{P}\{N \geq k\}$ is of order $k^{1-\alpha/2}$ for large k).

We have therefore proved the following theorem, establishing square root convergence (in computational effort c) for our estimator.

Theorem 3.2. *If X is a Harris chain with $\mathbb{E}[T(2) - T(1)]^r < \infty$ for $r > 2$ and f is bounded, then*

$$c^{1/2}(\bar{Z}_{\Gamma(c)} - \mathbb{E}f(X_\infty)) \Rightarrow \sqrt{\mathbb{E}\xi \cdot \text{var } Z} N(0, 1) \quad \text{as } c \rightarrow \infty.$$

An improvement to the above coupling is easily implemented. In the above algorithm, τ occurs whenever X and X' m time steps earlier were in A , and both X and X' independently chose at that time to distribute themselves according to ν m time units later. But an alternative coupling is to generate (X_{T+m}, X'_{T+m-1}) as follows, whenever $(X_T, X'_{T-1}) \in A \times A$. As in the previous algorithm, distribute X_{T+m} according to ν with probability λ , and according to $Q(X_T, \cdot)$ with probability $1 - \lambda$. Now modify the dynamics for X' . Whenever X_{T+m} is distributed according to ν , set $X'_{T+m-1} = X_{T+m}$. On the other hand, whenever X_{T+m} is distributed according to $Q(X_T, \cdot)$, independently generate X'_{T+m-1} according to $Q(X'_{T-1}, \cdot)$. This coupling preserves the marginal distribution of X and X' , but the time τ' at which X and X' couple (so that $X_{\tau'} = X'_{\tau'-1}$) is a.s. smaller than under the previous ‘independent coupling’. Consequently, $\mathbb{P}\{\tau' \geq k\} \leq \mathbb{P}\{\tau \geq k\}$ for $k \geq 0$, so $\mathbb{E}(\tau')^r \leq \mathbb{E}\tau^r$ for $r > 0$, thereby establishing that this coupling can be used in place of τ in proving Theorem 3.2. Since $\tau' \leq \tau$, this coupling is computationally preferable to τ .

4. A Glivenko–Cantelli result

In some settings, one may be interested in computing the equilibrium distribution of some real-valued functional f of the Markov chain, rather than merely its expected value $\mathbb{E}f(X_\infty)$. In this section we study the behavior of our unbiased estimator for the equilibrium probability $\mathbb{P}\{f(X_\infty) \leq x\} = \mathbb{E}[\mathbf{1}\{f(X_\infty) \leq x\}]$ as a function of x . Because the mapping $\mathbf{1}\{f(\cdot) \leq x\}$ is not Lipschitz, the theory of Section 2 does not apply. Consequently, we focus here on the case where X is a positive recurrent Harris chain.

Set $Y_k = f(X_k)$ and $Y'_k = f(X'_k)$. Let $\{(Y_{k,j}, Y'_{k,j}) : 0 \leq k \leq \min\{\tau_j, N_j\} : j \geq 1\}$ be a sequence of i.i.d. copies of $\{(Y_k, Y'_k) : 0 \leq k \leq \min\{\tau, N\}\}$, where the Y_k and Y'_k are constructed as in Section 3. The empirical measure (intended to estimate $F_\infty(\cdot) := \mathbb{P}\{f(X_\infty) \leq \cdot\}$) associated with sample size n is then given by the random signed measure

$$\pi_n(\cdot) = \frac{1}{n} \sum_{j=1}^n \sum_{k=0}^{\tau_j \wedge N_j} \frac{\delta_{Y_{k,j}}(\cdot) - \delta_{Y'_{k-1,j}}(\cdot)}{\mathbb{P}\{N \geq k\}},$$

where $a \wedge b := \min\{a, b\}$ and $\delta_y(\cdot)$ is a unit point mass measure at y . Observe that

$$\int_S y \pi_n(dy) \quad \text{and} \quad F_n(x) := \int_S \mathbf{1}\{y \leq x\} \pi_n(dy)$$

are unbiased estimators for $\mathbb{E}f(X_\infty)$ and $F_\infty(x)$ respectively as in Section 3. We may rewrite $F_n(\cdot)$ as

$$F_n(x) = \sum_{k=0}^\infty \frac{1}{n} \sum_{j=1}^n (\mathbf{1}\{Y_{k,j} \leq x\} - \mathbf{1}\{Y'_{k-1,j} \leq x\}) \frac{\mathbf{1}\{\tau_j \wedge N_j \geq k\}}{\mathbb{P}\{N \geq k\}}.$$

Because the sample functions

$$\frac{1}{n} \sum_{j=1}^n \mathbf{1}\{Y_{k,j} \leq x\} \mathbf{1}\{\tau_j \wedge N_j \geq k\} \quad \text{and} \quad \frac{1}{n} \sum_{j=1}^n \mathbf{1}\{Y'_{k-1,j} \leq x\} \mathbf{1}\{\tau_j \wedge N_j \geq k\}$$

are monotone in x , a proof identical to that of the standard Glivenko–Cantelli theorem (see, for example, Chung (2001)) establishes that

$$\sup_x \left| \sum_{j=1}^n \mathbf{1}\{Y_{k,j} \leq x\} \mathbf{1}\{\tau_j \wedge N_j \geq k\} - \mathbb{E}[\mathbf{1}\{Y_k \leq x\} \mathbf{1}\{\tau \geq k\}] \mathbb{P}\{N_j \geq k\} \right| \xrightarrow{\text{a.s.}} 0$$

and

$$\sup_x \left| \sum_{j=1}^n \mathbf{1}\{Y'_{k-1,j} \leq x\} \mathbf{1}\{\tau_j \wedge N_j \geq k\} - \mathbb{E}[\mathbf{1}\{Y'_{k-1} \leq x\} \mathbf{1}\{\tau \geq k\}] \mathbb{P}\{N_j \geq k\} \right| \xrightarrow{\text{a.s.}} 0$$

as $n \rightarrow \infty$, for each fixed $k \geq 0$. Since we proved in Section 3 that

$$\mathbb{E}[(\mathbf{1}\{Y_k \leq x\} - \mathbf{1}\{Y'_{k-1} \leq x\}) \mathbf{1}\{\tau \geq k\}] = \mathbb{P}\{Y_k \leq x\} - \mathbb{P}\{Y_{k-1} \leq x\},$$

it follows that, for any $m \geq 1$,

$$\sup_x \left| \sum_{k=0}^m \frac{1}{n} \sum_{j=1}^n (\mathbf{1}\{Y_{k,j} \leq x\} - \mathbf{1}\{Y'_{k-1,j} \leq x\}) \frac{\mathbf{1}\{\tau_j \wedge N_j \geq k\}}{\mathbb{P}\{N \geq k\}} - \mathbb{P}\{Y_m \leq x\} \right| \xrightarrow{\text{a.s.}} 0 \quad (4.1)$$

as $n \rightarrow \infty$. If X is any aperiodic positive recurrent Harris chain, Y_m converges to Y_∞ in total variation, and, hence,

$$\sup_x |\mathbb{P}\{Y_m \leq x\} - F_\infty(x)| \rightarrow 0 \quad \text{as } m \rightarrow \infty. \quad (4.2)$$

(If X is periodic, adapt (4.2) by restricting m to multiples of p , replacing (4.2) by

$$\sup_x \left| \frac{1}{p} \sum_{i=0}^{p-1} \mathbb{P}\{Y_{m+i} \leq x\} - F_\infty(x) \right| \rightarrow 0$$

as $m \rightarrow \infty$.)

Suppose now that $\mathbb{P}\{N \geq k\} \sim ck^{-\alpha}$ as $k \rightarrow \infty$, where $a_k \sim b_k$ means that $a_k/b_k \rightarrow 1$ as $k \rightarrow \infty$. Note that, for sufficiently large m ,

$$\begin{aligned} & \sum_{k>m} \left| \frac{1}{n} \sum_{j=1}^n (\mathbf{1}\{Y_{kj} \leq x\} - \mathbf{1}\{Y'_{k-1,j} \leq x\}) \frac{\mathbf{1}\{\tau_j \wedge N_j \geq k\}}{\mathbb{P}\{N \geq k\}} \right| \\ & \leq \sum_{k>m} \frac{1}{n} \sum_{j=1}^n \frac{\mathbf{1}\{\tau_j \wedge N_j \geq k\}}{\mathbb{P}\{N \geq k\}} \\ & \leq \frac{2}{c} \sum_{k>m} k^\alpha \frac{1}{n} \sum_{j=1}^n \mathbf{1}\{\tau_j \wedge N_j \geq k\} \\ & \leq \frac{3}{c(\alpha + 1)} \frac{1}{n} \sum_{j=1}^n (\tau_j \wedge N_j + 1)^{\alpha+1} \mathbf{1}\{\tau_j \wedge N_j \geq m\} \\ & \xrightarrow{\text{a.s.}} \frac{3}{c(\alpha + 1)} \mathbb{E}[(\tau \wedge N + 1)^{\alpha+1} \mathbf{1}\{\tau \wedge N \geq m\}] \quad \text{as } n \rightarrow \infty. \end{aligned} \tag{4.3}$$

If $\mathbb{E}(\tau \wedge N)^{\alpha+1} < \infty$, it follows that

$$\mathbb{E}[(\tau \wedge N + 1)^{\alpha+1} \mathbf{1}\{\tau \wedge N \geq m\}] \rightarrow 0 \quad \text{as } m \rightarrow \infty. \tag{4.4}$$

By first fixing m , then letting $n \rightarrow \infty$, and finally sending $m \rightarrow \infty$, (4.1)–(4.4) therefore prove that

$$\sup_x |F_n(x) - F_\infty(x)| \xrightarrow{\text{a.s.}} 0 \quad \text{as } n \rightarrow \infty.$$

It remains to consider the finiteness of $\mathbb{E}(\tau \wedge N)^{\alpha+1}$. Observe that

$$\mathbb{E}(\tau \wedge N)^{\alpha+1} \leq (\alpha + 1) \sum_{k=0}^\infty k^\alpha \mathbb{P}\{\tau \geq k\} \mathbb{P}\{N \geq k\}.$$

But $\sum_{k=0}^\infty \mathbb{P}\{\tau \geq k\} = \mathbb{E}(\tau + 1) < \infty$ and $k^\alpha \mathbb{P}\{N \geq k\} \rightarrow c$ as $k \rightarrow \infty$, thereby proving that $\mathbb{E}(\tau \wedge N)^{\alpha+1}$ is necessarily finite.

This proves the following Glivenko–Cantelli-type theorem for the estimator of Section 3.

Theorem 4.1. *Suppose that $\mathbb{P}\{N \geq k\} \sim ck^{-\alpha}$ as $k \rightarrow \infty$ for $\alpha > 0$. If X is a positive recurrent Harris chain then $F_n(x)$ is an unbiased estimator for $F_\infty(x)$ for each $x \in \mathbb{R}$, and*

$$\sup_x |F_n(x) - F_\infty(x)| \xrightarrow{\text{a.s.}} 0 \quad \text{as } n \rightarrow \infty.$$

5. Numerical results

We present here a brief account of the numerical performance of our exact estimation algorithms.

As an example of a contracting chain, consider the non- ϕ -irreducible Markov chain $X = \{X_n : n \geq 0\}$ given by

$$X_{n+1} = \frac{1}{2}X_n + V_{n+1},$$

TABLE 1: Contracting chain: 10^6 time steps, $X_0 = 1$, $\mathbb{P}\{N \geq n\} = 2^{1-n}$.

$f(x)$	Estimator	90% confidence interval	# Samples
x	Z	$1.013 \pm 1.1 \times 10^{-2}$	6.7×10^4
	Z^*	$0.9974 \pm 7.3 \times 10^{-3}$	5.0×10^4
$\min(1, x)$	Z	$0.7531 \pm 6.2 \times 10^{-3}$	6.7×10^4
	Z^*	$0.7552 \pm 4.7 \times 10^{-3}$	5.1×10^4
x^2	Z	$1.344 \pm 2.3 \times 10^{-2}$	6.7×10^4
	Z^*	$1.334 \pm 1.6 \times 10^{-2}$	5.0×10^4

TABLE 2: Contracting chain: 10^6 time steps, $X_0 = 1$, $\mathbb{P}\{N \geq n\} = 0.95^{n-1}$.

$f(x)$	Estimator	90% confidence interval	# Samples
x	Z	$1.009 \pm 3.3 \times 10^{-2}$	2.7×10^3
	Z^*	$1.006 \pm 6.1 \times 10^{-2}$	2.4×10^2
$\min(1, x)$	Z	$0.743 \pm 1.7 \times 10^{-2}$	2.7×10^3
	Z^*	$0.764 \pm 3.6 \times 10^{-2}$	2.5×10^2
x^2	Z	$1.356 \pm 6.7 \times 10^{-2}$	2.7×10^3
	Z^*	$1.39 \pm 1.3 \times 10^{-1}$	2.5×10^2

where $S = [0, 2]$ and the V_i are i.i.d. with $\mathbb{P}\{V_n = 0\} = \frac{1}{2} = \mathbb{P}\{V_n = 1\}$, with corresponding Lipschitz functions $f_1(x) = x$, $f_2(x) = \min\{1, x\}$, and $f_3(x) = x^2$. For this example, π is uniform on $[0, 2]$,

$$\mathbb{E}f_1(X_\infty) = 1, \quad \mathbb{E}f_2(X_\infty) = \frac{3}{4}, \quad \text{and} \quad \mathbb{E}f_3(X_\infty) = \frac{4}{3}.$$

In Tables 1 and 2 we present results for two different distributions for N . As expected, the algorithm based on Z becomes more attractive when N has a heavier tail, because the computational effort for Z^* increases quadratically in N (because of the nonrecursive computation of the Δ_i), whereas the effort for Z increases linearly in N .

We turn next to the Harris chain algorithm, implemented with the coupling τ' of Section 3. Consider the Markov chain $W = \{W_n : n \geq 0\}$ on \mathbb{R}_+ corresponding to the waiting time sequence for the M/M/1 queue, with arrival rate $\frac{1}{2}$ and unit service rate. The equilibrium distribution π here is a mixture of a unit point mass at 0 and an exponential distribution with rate parameter $\frac{1}{2}$, with the same probability $\frac{1}{2}$ for each mixture component. Let the function f be given by $f(x) = \mathbf{1}\{x > 1\}$, so that

$$\mathbb{P}\{W_\infty > 1\} = \frac{1}{2}e^{-1/2} \approx 0.303.$$

As for N , we note that the same proof technique as for Proposition 1 of Rhee and Glynn (2013) establishes that the optimal choice for the distribution of N is to choose $\mathbb{P}\{N \geq k\}$ proportional to

$$\sqrt{\frac{\mathbb{E}\Delta_k^2 + 2 \sum_{j=k+1}^\infty \mathbb{E}\Delta_k \Delta_j}{\mathbb{P}\{\tau \geq k\}}},$$

provided that this sequence is nonincreasing. Since it seems likely that $\mathbb{E}\Delta_k^2$ will frequently be of roughly the same order as $\mathbb{P}\{\tau \geq k\}$ for large k , this suggests that the optimal distribution will often have positive mass at infinity. In view of this observation, we have chosen to use

TABLE 3: Harris chain.

# Steps simulated	90% confidence interval	# Samples
1.0×10^5	$0.283 \pm 8.3 \times 10^{-2}$	3.2×10^4
2.0×10^5	$0.279 \pm 5.8 \times 10^{-2}$	6.5×10^4
5.0×10^5	$0.296 \pm 3.4 \times 10^{-2}$	1.6×10^5
1.0×10^6	$0.329 \pm 2.7 \times 10^{-2}$	3.2×10^5
2.0×10^6	$0.294 \pm 1.8 \times 10^{-2}$	6.5×10^5
5.0×10^6	$0.308 \pm 1.2 \times 10^{-2}$	1.6×10^6
1.0×10^7	$0.2992 \pm 8.4 \times 10^{-3}$	3.2×10^6
2.0×10^7	$0.3089 \pm 5.8 \times 10^{-3}$	6.5×10^6
5.0×10^7	$0.2995 \pm 3.7 \times 10^{-3}$	1.6×10^7
1.0×10^8	$0.3041 \pm 2.6 \times 10^{-3}$	3.3×10^7
2.0×10^8	$0.3024 \pm 1.9 \times 10^{-3}$	6.5×10^7
5.0×10^8	$0.3036 \pm 1.2 \times 10^{-3}$	1.6×10^8

a very heavy-tailed specification for N , namely, $\mathbb{P}\{N \geq k\} = 1/k$ for $k \geq 0$. In Table 3 we present the results of our computations with $\lambda = 1$ and a small set $A = \{0\}$; the results show the ‘square root’ decrease in the width of the confidence interval that is to be expected.

Acknowledgements

The authors gratefully acknowledge the support of the National Science Foundation grant DMS-1320158 and a Samsung Scholarship, respectively. We also thank the anonymous referees for their excellent suggestions for improving the readability of the paper.

References

- ASMUSSEN, S. AND GLYNN, P. W. (2007). *Stochastic Simulation: Algorithms and Analysis* (Stoch. Modelling Appl. Prob. 57). Springer, New York.
- ASMUSSEN, S., GLYNN, P. W. AND THORISSON, H. (1992). Stationarity detection in the initial transient problem. *ACM Trans. Model. Comput. Simul.* **2**, 130–157.
- ATHREYA, K. B. AND NEY, P. (1978). A new approach to the limit theory of recurrent Markov chains. *Trans. Amer. Math. Soc.* **245**, 493–501.
- BLANCHET, J. AND WALLWATER, A. (2014). Exact sampling for the steady-state waiting time of a heavy-tailed single server queue. Submitted.
- BOROVKOV, A. A. AND FOSS, S. G. (1992). Stochastically recursive sequences and their generalizations. *Siberian Adv. Math.* **2**, 16–81.
- CHUNG, K. L. (2001). *A Course in Probability Theory*, 3rd edn. Academic Press, San Diego, CA.
- CONNOR, S. B. AND KENDALL, W. S. (2007). Perfect simulation for a class of positive recurrent Markov chains. *Ann. Appl. Prob.* **17**, 781–808.
- CORCORAN, J. N. AND TWEEDIE, R. L. (2001). Perfect sampling of ergodic Harris chains. *Ann. Appl. Prob.* **11**, 438–451.
- DIACONIS, P. AND FREEDMAN, D. (1999). Iterated random functions. *SIAM Rev.* **41**, 45–76.
- ENSOR, K. B. AND GLYNN, P. W. (2000). Simulating the maximum of a random walk. *J. Statist. Planning Infer.* **85**, 127–135.
- GILES, M. B. (2008). Multilevel Monte Carlo path simulation. *Operat. Res.* **56**, 607–617.
- GLYNN, P. W. AND WHITT, W. (1992). The asymptotic efficiency of simulation estimators. *Operat. Res.* **40**, 505–520.
- KENDALL, W. S. (2004). Geometric ergodicity and perfect simulation. *Electron. Commun. Probab.* **9**, 140–151.
- LINDVALL, T. (2002). *Lectures on the Coupling Method*. Dover Publications, Mineola, NY.
- MCLEISH, D. (2011). A general method for debiasing a Monte Carlo estimator. *Monte Carlo Meth. Appl.* **17**, 301–315.
- MEYN, S. AND TWEEDIE, R. L. (2009). *Markov Chains and Stochastic Stability*, 2nd edn. Cambridge University Press.
- PROPP, J. G. AND WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures Algorithms* **9**, 223–252.

- RHEE, C.-H. AND GLYNN, P. W. (2012). A new approach to unbiased estimation for SDE's. In *Proc. 2012 Winter Simulation Conference*, eds. C. Laroque *et al.*, 7pp.
- RHEE, C.-H. AND GLYNN, P. W. (2013). Unbiased estimation with square root convergence for SDE models. Submitted.
- RYCHLIK, T. (1990). Unbiased nonparametric estimation of the derivative of the mean. *Statist. Prob. Lett.* **10**, 329–333.
- RYCHLIK, T. (1995). A class of unbiased kernel estimates of a probability density function. *Appl. Math.* **22**, 485–497.
- THORISSON, H. (2000). *Coupling, Stationarity, and Regeneration*. Springer, New York.

PETER W. GLYNN, *Stanford University*

Management Science and Engineering, Stanford University, Stanford, CA 94305-4121, USA.

Email address: glynn@stanford.edu

CHANG-HAN RHEE, *Georgia Institute of Technology*

Biomedical Engineering, Georgia Institute of Technology, U. A. Whitaker Building, 313 Ferst Drive, Atlanta, GA 30332, USA. Email address: rhee@gatech.edu