Proceedings of the 2016 Winter Simulation Conference
T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, eds.

EXTENSIONS OF THE REGENERATIVE METHOD TO NEW FUNCTIONALS

Zeyu Zheng
Peter W. Glynn

Department of Management Science and Engineering
Stanford University
Stanford, CA 94305, USA

ABSTRACT

This paper briefly reviews the regenerative method for steady-state simulation, and then shows how regenerative structure can be used computationally to develop new estimators for the spectral density, moments of hitting times, and both discounted and average reward value functions. All our estimators typically exhibit the Monte Carlo method's usual "square root" convergence rate. This is in contrast to the usual sub-square root rate exhibited by, for example, spectral density estimators in the absence of regenerative structure.

1 INTRODUCTION

In this paper, we discuss the use of regenerative structure in the construction of efficient simulation-based algorithms for computing various functionals that arise in stochastic modeling and statistics. Our starting point is that of steady-state simulation. In particular, suppose that $Y = (Y(t) : t \geq 0)$ is a real-valued stochastic process representing the output of a simulation, for which three exists a (deterministic) constant α such that

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

as $t \to \infty$, where $\Rightarrow$ denotes weak convergence. The *steady-state simulation problem* is concerned with the efficient computation of α, and with the development of associated confidence intervals for α. (We phrase our results in this paper in continuous time only for notational convenience. To handle discrete time sequences $(Y_n : n \geq 0)$, we can always embed in continuous time via the relation $Y(t) \triangleq Y_{\lfloor t \rfloor}$ for $t \geq 0$.)

To construct confidence intervals for α, the standard starting point is a central limit theorem (CLT) for $\bar{Y}(t)$, in which one strengthens (1) to the limit theorem

$$t^{\frac{1}{2}}(\bar{Y}(t) - \alpha) \Rightarrow \sigma N(0,1) \tag{2}$$

as $t \to \infty$, for some (deterministic) constant α, where $N(0,1)$ is a normal random variable (rv) with mean 0 and unit variance. The quantity σ^2 is called the *time-average variance constant* (TAVC) associated with Y. Hence, if one can construct an estimator $s(t)$ from the simulation output collected up to time t for which

$$s(t) \Rightarrow \sigma \tag{3}$$

as $t \to \infty$, it follows that

$$\left[\bar{Y}(t) - z\frac{s(t)}{\sqrt{t}}, \bar{Y}(t) + z\frac{s(t)}{\sqrt{t}}\right]$$

is an asymptotic $100(1-\delta)\%$ confidence interval for α, provided that $\sigma^2 > 0$ and that we select z so that $P(-z \leq N(0,1) \leq z) = 1 - \delta$.

Unfortunately, constructing such estimators $s(t)$ is not easy in general. In particular, it is usual that when (1) and (2) hold, then

$$(Y(t+u) : u \geq 0) \to (Y^*(u) : u \geq 0)$$

as $t \to \infty$, where $(Y^*(u) : u \geq 0)$ is a stationary process known as the "stationary version" of Y. (The convergence in (4) is usually stated in terms of total variation convergence; see Chapter 5 of Thorisson (2000) for such results.) This suggests that for t large,

$$t^{\frac{1}{2}}(\bar{Y}(t) - \alpha) \overset{D}{\approx} t^{\frac{1}{2}} \left(\frac{1}{t} \int_0^t Y^*(s)ds - \alpha \right),$$

where $\overset{D}{\approx}$ denotes "has approximately the same distribution as" (and is intended to have no rigorous meaning). This, in turn, suggests that

$$\begin{aligned} \sigma^2 &= \lim_{t\to\infty} \frac{1}{t} Var \left(\int_0^t Y^*(s)ds \right) \\ &= 2 \int_0^\infty Cov(Y^*(0), Y^*(s))ds \end{aligned} \tag{4}$$

The formula (4) for the TAVC σ^2 can be rigorously verified in many different settings: see, for example, Meyn and Tweedie (2009), p.422; the same result shows that $\alpha = EY^*(t)$.

The challenge in (4) is that while $Cov(Y^*(0), Y^*(t))$ can typically be easily estimated for each fixed t, computing the integrated covariance quantity σ^2, as given by (4), is statistically hard. The TAVC σ^2 is intimately connected to spectral density estimation for the stationary process $Y^* = (Y^*(t) : t \geq 0)$. Specifically, $\sigma^2 = 2\pi f(0)$, where $f(\cdot)$ is the spectral density of Y^*. In general, the best spectral density estimators converge at rates slower than of order $t^{-\frac{1}{2}}$ in the simulated time horizon t; see Anderson (1971), Chapter 9. This "sub square root" convergence rate reflects the fact that infinitely many covariance terms must be summed together to compute σ^2.

The regenerative method for TAVC estimation takes advantage of the presence of regenerative structure in Y to obtain an estimator $s(t)$ for σ that converges at a rate of order $t^{-\frac{1}{2}}$. This beautiful idea, due independently to Fishman (1974) and Crane and Iglehart (1974), exploits the independent and identically distributed (iid) cycle structure of (classically) regenerative processes, so that one can efficiently truncate the integral appearing in (4) to include only covariance effects between Y^* values collected within the same cycle. This leads to estimation algorithms for σ^2 that converge much more rapidly than do those associated with general spectral density estimation theory or other related statistical procedures.

This paper is organized as follows. Section 2 reviews the key facts about regenerative methodology in its historical setting of steady-state simulation. Regenerative methodology is then extended in Section 3 to cover spectral density estimation, where appropriate estimators and associated confidence intervals are developed. Section 4 discusses new regenerative estimators for higher order moments of hitting times (e.g. the variance), while Section 5 deals with new regenerative representations and related estimators for value functions in both the discounted and average reward settings. Finally, Section 6 provides a brief discussion of computational results obtained for the spectral density estimator of Section 3.

2 APPLICATION 1: ESTIMATING THE TAVC

Given that the estimation of the TAVC σ^2 is the classical problem to which regenerative methodology has been applied within the setting of simulation-based algorithms, we start with a brief review of this estimation domain. We presume that $Y = (Y(t) : t \geq 0)$ is a non-delayed *classically regenerative process*. More precisely, we assume that there exist random times $0 = T(0) < T(1) < \cdots$ such that the cycles $(W_i : i \geq 1)$ are iid, where $W_i \triangleq (Y(T(i-1)+s) : 0 \leq s < \tau_i)$ and $\tau_i \triangleq T(i) - T(i-1)$ for $i \geq 1$.

It is known that if Y is non-negative and $E\tau_1 < \infty$ (so that Y is *positive recurrent*), then

$$\bar{Y}(t) \Rightarrow \alpha \tag{5}$$

for a finite α as $t \to \infty$ if and only if $E\int_0^{T(1)} Y(s)ds < \infty$, in which case

$$\alpha = \frac{E\int_0^{T(1)} Y(s)ds}{E\tau_1}; \tag{6}$$

see Glynn and Whitt (1993). In addition,

$$t^{\frac{1}{2}}(\bar{Y}(t) - \alpha) \Rightarrow \sigma N(0,1) \tag{7}$$

as $t \to \infty$ holds if and only if $E\left(\int_0^{T(1)} Y_c(s)ds\right)^2 < \infty$, where $Y_c(s) \triangleq Y(s) - \alpha$, in which case

$$\sigma^2 = \frac{E\left(\int_0^{T(1)} Y_c(s)ds\right)^2}{E\tau_1}; \tag{8}$$

see Glynn and Whitt (2002). In some sense, the representation (8) truncates the covariance integral (4) according to the "cycle boundary" $T(1)$, as suggested in the Introduction. In view of (7), the natural estimator for σ^2 is then

$$s^2(t) = \frac{1}{t}\sum_{i=1}^{N(t)} \left(\int_{T(i-1)}^{T(i)} (Y(s) - \bar{Y}(t))ds\right)^2,$$

where $N(t) = \max(n \geq 0 : T(n) \leq t)$ is the number of regenerative cycles completed by time t. It turns out that $s^2(t) \Rightarrow \sigma^2$ as $t \to \infty$, precisely when (7) holds; see Glynn and Iglehart (1993). In other words, the regenerative method for estimating σ^2 is consistent under the weakest possible conditions. Furthermore, $s(t)$ converges to σ at rate $t^{-\frac{1}{2}}$ under modest additional conditions; see Glynn and Iglehart (1987).

The above theory covers classically regenerative processes. From a modeling viewpoint, this allows one to easily apply the regenerative method to irreducible positive recurrent discrete state space Markov chains and processes. For such examples, the $T(i)$'s are typically defined as the consecutive times at which the underlying Markov chain process enters some fixed (regeneration) state, say z.

But some modeling environments require that one develop steady-state simulation methodology for continuous state space examples. Regenerative methods apply there, as well, at a significant level of generality. In particular, a large class of such models form either Harris recurrent Markov chains (when formulated in discrete time) or Harris recurrent Markov processes (when formulated in continuous time).

Consider first the case of a Harris recurrent Markov chain $X = (X_n : n \geq 0)$, in which case we presume that $Y_n = h(X_n)$ for some non-negative h (see Meyn and Tweedie (2009) for a complete discussion of this class of chains.) In this setting, $Y = (Y(t) : t \geq 0)$ can always be simulated as a non-delayed *wide-sense regenerative process*, so that there exist random times $0 = T(0) < T(1) < \cdots$ such that:

i) $(\tau_i : i \geq 1)$ is an iid sequence;
ii) $T(i)$ is independent of $(Y(T(i)+u) : u \geq 0)$ for $i \geq 0$;
iii) $(Y(T(i)+u) : u \geq 0) \stackrel{D}{=} (Y(u) : u \geq 0)$ for $i \geq 0$ (where $\stackrel{D}{=}$ denotes equality in distribution).

It is then easily seen that if (5) holds, then α is given by (6). With the identity (6) in place, one has now reduced α to a "finite horizon" computation, in the sense that α can be computed by generating iid copies $(\beta_1, \tau_1), (\beta_2, \tau_2), \ldots$ of (β, τ), where $\tau = T(1)$ and $\beta = \int_0^{T(1)} Y(s)ds$, and forming the sample mean estimator

$$\frac{\bar{\beta}_n}{\bar{\tau}_n} = \frac{\sum_{i=1}^n \beta_i/n}{\sum_{i=1}^n \tau_i/n}.$$

If $E(\beta_1^2+\tau_1^2)<\infty$, then

$$n^{\frac{1}{2}}\left(\frac{\bar{\beta}_n}{\bar{\tau}_n}-\alpha\right)\Rightarrow\sqrt{\frac{E(\beta_1-\alpha\tau_1)^2}{(E\tau_1)^2}}N(0,1) \tag{9}$$

as $n\to\infty$, and

$$s_n^2\triangleq\frac{1}{n}\sum_{i=1}^{n}\left(\beta_i-\left(\frac{\bar{\beta}_n}{\bar{\tau}_n}\right)\tau_i\right)^2/\bar{\tau}_n{}^2\to\frac{E(\beta_1-\alpha\tau_1)^2}{(E\tau_1)^2}\quad a.s. \tag{10}$$

as $n\to\infty$. Again, confidence intervals for α can be easily constructed, given the above results.

To handle the Harris recurrent Markov process setting, we appeal to the fact that if $Y(t)=h(X(t))$ for some non-negative h, then $Y=(Y(t):t\geq 0)$ can be simulated as a non-delayed *one-dependent regenerative process*, so that there exist random times $0=T(0)<T(1)<\cdots$ such that:

i) $(W_i:i\geq 1)$ is an identically distributed sequence;
ii) $(W_i:i\geq 1)$ is a one-dependent sequence, so that $(W_{n+j}:j\geq 2)$ is independent of $(W_i:i\leq n)$ for each $n\geq 1$.

As observed in Glynn (1994), (5) implies that (6) holds in this setting, so that confidence intervals for α can again be constructed via (9) and (10).

It is important to recognize that the estimator $\bar{\beta}_n/\bar{\tau}_n$ does not in general have the same distribution as $\bar{Y}(T(n))$, as occurs when Y is classically regenerative. In the classical setting, simulating n iid copies of a cycle is equivalent to simulating Y to time $T(n)$. On the other hand, the correlation structure of the cycles associated with both wide-sense regenerative and one-dependent regenerative processes is more complex than in the classical case, so that simulating n iid copies of (β_1,τ_1) does not yield the path of the original process Y to time $T(n)$. This does not affect the validity of (9) and (10), but it does mean that we can not use such simulations to simultaneously compute (for example) transient quantities, since the iid structure of the (β_i,τ_i)'s may result in modification of the marginal distributions of $Y(t)$ for $t\geq 0$.

3 APPLICATION 2: THE SPECTRAL DENSITY

As noted in Section 2, the TAVC constant σ^2 is a special case of a spectral density computation. Specifically, if $Y^*=(Y^*(t):t\geq 0)$ is a real-valued stationary process for which

$$\int_0^\infty|Cov(Y^*(0),Y^*(t))|dt<\infty, \tag{11}$$

the *spectral density* $f(\cdot)$ is defined by

$$f(\lambda)=\frac{1}{2\pi}\int_{-\infty}^{\infty}e^{i\lambda t}Cov(Y^*(0),Y^*(t))dt;$$

Note that the TAVC can be recovered from $f(\cdot)$ via the relation $\sigma^2=2\pi f(0)$ (Compare with (4)). The spectral density plays a fundamental role in many different applications areas, and is a key concept in the statistical analysis of time series. In general, estimating $f(\lambda)$ is statistically difficult, and (nonparametric) spectral density estimators exhibit a sub-square root convergence rate in the sampling effort t. However, we will now show how classical regenerative structure can be exploited to obtain an estimator for $f(\lambda)$ having square root convergence rate.

Let $Y=(Y(t):t\geq 0)$ be a non-delayed classically regenerative process with $E\tau_1<\infty$. Then, there exists a unique stationary version Y^* of Y for which

$$P(Y^*\in\cdot)=\frac{E\int_0^{T(1)}I(Y_s\in\cdot)ds}{E\tau_1}, \tag{12}$$

where $Y_s = (Y(s+u) : u \geq 0)$; see, for example, p.342 Thorisson (2000). The following result provides a regenerative representation for $f(\cdot)$. Recall that a rv is *spread-out* if its n-fold convolution has a density component for some $n \geq 1$.

Theorem 1 Suppose that $T(1)$ is a spread-out random variable and

$$E\left(\left(\max_{0\leq s\leq T(1)} |Y(s)|^2 + T(1)^{2+p}\right) T(1)\right) < \infty, \tag{13}$$

for some $p > 0$. Then, for $\lambda \neq 0$,

$$\begin{aligned} 2\pi f(\lambda) = & \frac{E\int_0^{T(1)}\int_s^{T(1)} Y_c(s)Y_c(t)\left(e^{i\lambda(t-s)} + e^{-i\lambda(t-s)}\right)dtds}{E\tau_1} \\ & + \frac{E\int_0^{T(1)} Y_c(s)e^{i\lambda(T(1)-s)}ds}{E\tau_1} \cdot \frac{E\int_0^{T(1)} e^{i\lambda t}Y_c(t)dt}{1-Ee^{i\lambda T(1)}} \\ & + \frac{E\int_0^{T(1)} Y_c(s)e^{-i\lambda(T(1)-s)}ds}{E\tau_1} \cdot \frac{E\int_0^{T(1)} e^{-i\lambda t}Y_c(t)dt}{1-Ee^{-i\lambda T(1)}}, \end{aligned}$$

whereas for $\lambda = 0$,

$$\pi f(0) = \frac{E\int_0^{T(1)}\int_s^{T(1)} Y_c(s)Y_c(t)dtds}{E\tau_1}.$$

Proof. We start by verifying that (11) holds, so that $f(\cdot)$ is well-defined. Set $Y_c^*(t) = Y^*(t) - \alpha$ and

$$\begin{aligned} Cov(Y^*(0), Y^*(t)) &= EY_c^*(0)Y_c^*(t) \\ &= EY_c^*(0)Y_c^*(t)I(T^*(0) > t) + \int_{[0,t]} EY_c^*(0)I(T^*(0)\in ds)EY_c(t-s), \end{aligned} \tag{14}$$

Recalling that $T(1)$ is the first regeneration time for the non-delayed process Y, (12) implies that,

$$\begin{aligned} \int_0^\infty E|Y_c^*(0)Y_c^*(t)|I(T^*(0) > t)dt &= E\int_0^{T(1)} |Y_c(s)| \int_s^{T(1)} |Y_c(t)|dtds/E\tau_1 \\ &= \frac{1}{2}E\left(\int_0^{T(1)} |Y_c(s)|ds\right)^2 /E\tau_1 < \infty, \end{aligned} \tag{15}$$

in view of (13). Also, (12) yields

$$\begin{aligned} \int_0^\infty \int_{[0,t]} E|Y_c^*(0)|I(T^*(0)\in ds)|EY_c(t-s)|dt &= \int_{[0,\infty)} E|Y_c^*(0)|I(T^*(0)\in ds)\cdot \int_0^\infty |EY_c(t)|dt \\ &= E|Y_c^*(0)|T^*(0)\cdot \int_0^\infty |EY_c(t)|dt \\ &= \frac{E\int_0^{T(1)} |Y_c(s)|(T(1)-s)ds}{E\tau_1}\cdot \int_0^\infty |EY_c(t)|dt, \end{aligned} \tag{16}$$

which (13) establishes is finite, provided that we show that $\int_0^\infty |EY_c(t)|dt < \infty$.

Since $T(1)$ is spread-out with $ET(1) < \infty$, $Y_c(\cdot)$ can be coupled to its stationary version $Y_c^*(\cdot)$. Consequently, if β is the associated coupling time,

$$\begin{aligned}
EY(t) - \alpha &= E(Y(t) - Y^*(t))I(\beta \le t) + E(Y(t) - Y^*(t))I(\beta > t) \\
&= E(Y(t) - Y^*(t))I(\beta > t) \\
&\le 2E^{\frac{1}{2}}\left(Y^2(t) + Y^*(t)^2\right)P^{\frac{1}{2}}(\beta > t) \\
&= 2\left(EY^2(t) + EY^*(0)^2\right)^{\frac{1}{2}} o(t^{-2-p})^{\frac{1}{2}} \\
&= 2\left(EY^2(t) + \frac{E\int_0^{T(1)} Y(s)^2 ds}{E\tau_1}\right)^{\frac{1}{2}} o(t^{-1-\frac{p}{2}}) \\
&\le 2\left(EY^2(t) + \frac{E\max_{0\le s\le T(1)} Y^2(s)\cdot T(1)}{E\tau_1}\right)^{\frac{1}{2}} o(t^{-1-\frac{p}{2}}) \qquad (17)
\end{aligned}$$

as $t \to \infty$, where we used the fact that $P(\beta > t) = o(t^{-2-p})$ as $t \to \infty$ when $ET(1)^{3+p} < \infty$; see p.417 of Thorisson (2000). Hence, $|EY_c(\cdot)|$ is integrable, provided we show that $EY^2(\cdot)$ is a bounded function.

Note that if $a(t) \triangleq EY^2(t)$, $a(\cdot)$ satisfies the renewal equation

$$a(t) = EY^2(t)I(T(1) > t) + \int_{[0,t]} a(t-s)P(T(1) \in ds).$$

Since $EY^2(t)I(T(1) > t) \le E\max_{0\le s\le T(1)} Y^2(s)I(T(1) > t)$ and this upper bound is non-increasing and integrable (due to 13), the renewal theorem establishes that $EY^2(t)$ converges to a finite limit as $t \to \infty$; see p.147 of Asmussen (2003). So, $EY^2(\cdot)$ is bounded, and (14)-(17) imply (11) is satisfied.

We now turn to verifying the regenerative representation for $f(\cdot)$. We focus on $\lambda \ne 0$, since $\lambda = 0$ is exactly the identity (8). In view of (14), we see that

$$\begin{aligned}
&\int_{-\infty}^{\infty} e^{i\lambda t} Cov(Y^*(0), Y^*(t))dt \\
&= \int_0^{\infty} (e^{i\lambda t} + e^{-i\lambda t})EY_c^*(0)Y_c^*(t)I(T^*(0) > t) + \int_0^{\infty}\int_{[0,t]} EY_c^*(0)I(T^*(0) \in ds)e^{i\lambda s}EY_c(t-s)e^{i\lambda(t-s)}dt \\
&\quad + \int_0^{\infty}\int_{[0,t]} EY_c^*(0)I(T^*(0) \in ds)e^{-i\lambda s}EY_c(t-s)e^{-i\lambda(t-s)}dt \\
&= EY_c^*(0)\int_0^{T^*(0)} (e^{i\lambda t} + e^{-i\lambda t})Y_c^*(t)dt + EY_c^*(0)e^{i\lambda T^*(0)}\cdot\int_0^{\infty} EY_c(t)e^{i\lambda t}dt \\
&\quad + EY_c^*(0)e^{-i\lambda T^*(0)}\cdot\int_0^{\infty} EY_c(t)e^{-i\lambda t}dt \\
&= \frac{E\int_0^{T(1)} Y_c(s)\int_s^{T(1)}\left(e^{i\lambda(t-s)} + e^{-i\lambda(t-s)}\right)Y_c(t)dtds}{E\tau_1} + \frac{E\int_0^{T(1)} Y_c(s)e^{i\lambda(T(1)-s)}ds}{E\tau_1}\cdot\int_0^{\infty} EY_c(t)e^{i\lambda t}dt \\
&\quad + \frac{E\int_0^{T(1)} Y_c(s)e^{-i\lambda(T(1)-s)}ds}{E\tau_1}\cdot\int_0^{\infty} EY_c(t)e^{-i\lambda t}dt,
\end{aligned}$$

where (12) was again used for the final step.

Finally, if $\tilde{a}(t) \triangleq EY_c(t)$,

$$\tilde{a}(t) = EY_c(t)I(T(1) > t) + \int_{[0,t]} \tilde{a}(t-s)P(T(1) \in ds)$$

so

$$\int_0^\infty EY_c(t)e^{i\lambda t}dt = \int_0^\infty e^{i\lambda t}EY_c(t)I(T(1)>t)dt + \int_{[0,\infty)} e^{i\lambda s}P(T(1)\in ds)\cdot\int_0^\infty e^{i\lambda t}\tilde{a}(t)dt$$
$$= E\int_0^{T(1)} e^{i\lambda t}Y_c(t)dt + Ee^{i\lambda T(1)}\cdot\int_0^\infty e^{i\lambda t}\tilde{a}(t)dt$$

and hence

$$\int_0^\infty EY_c(t)e^{i\lambda t}dt = \frac{E\int_0^{T(1)} e^{i\lambda t}Y_c(t)dt}{1-Ee^{i\lambda T(1)}}.$$

In the last step, we take advantage of the fact that $|E\exp(i\lambda T(1))|<1$ for $\lambda\neq 0$ when $T(1)$ is spread-out; see p.170 of Chung (1979). This completes the proof. □

In view of Theorem 1, we can write the spectral density $f(\lambda)$ for $\lambda\neq 0$ as

$$f(\lambda) = \frac{1}{2\pi}\left(\frac{EA_1(\lambda)}{E\tau_1} + \frac{EB_1(\lambda)EC_1(\lambda)}{E\tau_1 ED(\lambda)} + \frac{EB_1(-\lambda)EC_1(-\lambda)}{E\tau_1 ED(-\lambda)}\right)$$

where

$$A_i(\lambda) = \int_{T(i-1)}^{T(i)}\int_s^{T(i)} Y_c(s)Y_c(t)\left(e^{i\lambda(t-s)} + e^{-i\lambda(t-s)}\right)dtds,$$
$$B_i(\lambda) = \int_{T(i-1)}^{T(i)} Y_c(s)e^{i\lambda(T(i)-s)}ds,$$
$$C_i(\lambda) = \int_{T(i-1)}^{T(i)} Y_c(s)e^{i\lambda(s-T(i-1))}ds,$$
$$D_i(\lambda) = 1-e^{i\lambda\tau_i}.$$

Of course, these rv's involve $Y_c(\cdot)$, which depends on the unknown "nuisance parameter" α. Hence, in estimating $f(\lambda)$, we replace $A_i(\lambda)$, $B_i(\lambda)$ and $C_i(\lambda)$ with

$$A_i(\lambda,t) = \int_{T(i-1)}^{T(i)}\int_s^{T(i)} (Y(s)-\bar{Y}(t))(Y(u)-\bar{Y}(t))\left(e^{i\lambda(u-s)} + e^{-i\lambda(u-s)}\right)duds,$$
$$B_i(\lambda,t) = \int_{T(i-1)}^{T(i)} (Y(s)-\bar{Y}(t))e^{i\lambda(T(i)-s)}ds,$$
$$C_i(\lambda,t) = \int_{T(i-1)}^{T(i)} (Y(s)-\bar{Y}(t))e^{i\lambda(s-T(i-1))}ds.$$

Then, if we set $\bar{A}(\lambda,t) = \sum_{i=1}^{N(t)} A_i(\lambda,t)/t$, $\bar{B}(\lambda,t) = \sum_{i=1}^{N(t)} B_i(\lambda,t)/t$, $\bar{C}(\lambda,t) = \sum_{i=1}^{N(t)} C_i(\lambda,t)/t$ and $\bar{D}(\lambda) = \sum_{i=1}^{N(t)} D_i(\lambda)/t$, our estimator for $f(\lambda)$ is defined via

$$2\pi\hat{f}_t(\lambda) = \bar{A}(\lambda,t) + \bar{B}(\lambda,t)\frac{\bar{C}(\lambda,t)}{\bar{D}(\lambda)} + \bar{B}(-\lambda,t)\frac{\bar{C}(-\lambda,t)}{\bar{D}(-\lambda)}. \tag{18}$$

To obtain confidence intervals for $f(\lambda)$, we need a CLT. To develop a CLT for $\hat{f}_t(\lambda)$, note that

$$2\pi\hat{f}_t(\lambda) - 2\pi f(\lambda) = 2\pi\hat{f}_t(\lambda) - \frac{1}{t}\sum_{j=1}^{N(t)}\left(\frac{EA_1(\lambda)}{E\tau_1} + \frac{EB_1(\lambda)EC_1(\lambda)}{E\tau_1 ED_1(\lambda)} + \frac{EB_1(-\lambda)EC_1(-\lambda)}{E\tau_1 ED_1(-\lambda)}\right)\tau_j + o_p(t^{-\frac{1}{2}})$$

as $t \to \infty$, where $o_p(t^{-\frac{1}{2}})$ denotes a stochastic process for which $t^{\frac{1}{2}} o_p(t^{-\frac{1}{2}}) \Rightarrow 0$ as $t \to \infty$. Suppose that

$$E\left(\int_0^{T(1)} (|Y(s)|+1)ds\right)^4 < \infty. \tag{19}$$

Then, it is easily seen that

$$\begin{aligned}
&\bar{A}(\lambda,t) - \frac{1}{t}\sum_{j=1}^{N(t)} \frac{EA_1(\lambda)}{E\tau_1}\tau_j \\
&= \frac{1}{t}\sum_{j=1}^{N(t)}\left[A_j(\lambda) - \frac{EA_1(\lambda)}{E\tau_1}\tau_j\right] - (\bar{Y}(t)-\alpha)\frac{1}{t}\sum_{j=1}^{N(t)}\int_{T(j-1)}^{T(j)}\int_s^{T(j)}(Y(u)+Y(s)) \\
&\quad \cdot \left(e^{i\lambda(u-s)} + e^{-i\lambda(u-s)}\right) duds + o_p(t^{-\frac{1}{2}}) \\
&= \frac{1}{t}\sum_{j=1}^{N(t)}\left[A_j(\lambda) - \frac{EA_1(\lambda)}{E\tau_1}\tau_j - Z_j\frac{E\tilde{A}_1(\lambda)}{E\tau_1}\right] + o_p(t^{-\frac{1}{2}})
\end{aligned}$$

as $t \to \infty$, where $Z_j = \int_{T(j-1)}^{T(j)} Y(s)ds - \alpha\tau_j$ and

$$\tilde{A}_j(\lambda) = \int_{T(j-1)}^{T(j)}\int_s^{T(j)}(Y(u)+Y(s))\left(e^{i\lambda(u-s)} + e^{-i\lambda(u-s)}\right)duds.$$

Similarly,

$$\bar{B}(\lambda,t) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{EB_1(\lambda)}{E\tau_1}\tau_j = \frac{1}{t}\sum_{j=1}^{N(t)}\left[B_j(\lambda) - \frac{EB_1(\lambda)}{E\tau_1}\tau_j - Z_j\frac{E\tilde{B}_1(\lambda)}{E\tau_1}\right] + o_p(t^{-\frac{1}{2}}),$$

$$\bar{C}(\lambda,t) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{EC_1(\lambda)}{E\tau_1}\tau_j = \frac{1}{t}\sum_{j=1}^{N(t)}\left[C_j(\lambda) - \frac{EC_1(\lambda)}{E\tau_1}\tau_j - Z_j\frac{E\tilde{C}_1(\lambda)}{E\tau_1}\right] + o_p(t^{-\frac{1}{2}}),$$

and

$$\frac{1}{\bar{D}(\lambda)} - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{E\tau_1}{ED(\lambda)}\tau_j = -\left(\frac{E\tau_1}{ED_1(\lambda)}\right)^2\frac{1}{t}\sum_{j=1}^{N(t)}\left[D_j(\lambda) - \frac{ED_1(\lambda)}{E\tau_1}\tau_j\right] + o_p(t^{-\frac{1}{2}}),$$

Consequently,

$$\begin{aligned}
2\pi\hat{f}_t(\lambda) - 2\pi f(\lambda) &= \left(\bar{A}(\lambda,t) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{EA_1(\lambda)}{E\tau_1}\tau_j\right) + \frac{EB_1(-\lambda)}{ED_1(-\lambda)}\left(\bar{C}(-\lambda,t) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{EC_1(-\lambda)}{E\tau_1}\tau_j\right) \\
&\quad + \frac{EB_1(\lambda)}{ED_1(\lambda)}\left(\bar{C}(\lambda,t) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{EC_1(\lambda)}{E\tau_1}\tau_j\right) - \frac{EB_1(\lambda)EC_1(\lambda)}{(ED_1(\lambda))^2}\left(\bar{D}(\lambda) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{ED_1(\lambda)}{E\tau_1}\tau_j\right) \\
&\quad + \frac{EC_1(-\lambda)}{ED_1(-\lambda)}\left(\bar{B}(-\lambda,t) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{EB_1(-\lambda)}{E\tau_1}\tau_j\right) + \frac{EC_1(\lambda)}{ED_1(\lambda)}\left(\bar{B}(\lambda,t) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{EB_1(\lambda)}{E\tau_1}\tau_j\right) \\
&\quad - \frac{EB_1(-\lambda)EC_1(-\lambda)}{(ED_1(-\lambda))^2}\left(\bar{D}(-\lambda) - \frac{1}{t}\sum_{j=1}^{N(t)}\frac{ED_1(-\lambda)}{E\tau_1}\tau_j\right) + o_p(t^{-\frac{1}{2}}) \\
&= \frac{1}{t}\sum_{j=1}^{N(t)} W_1(\lambda) + o_p(t^{-\frac{1}{2}})
\end{aligned}$$

as $t \to \infty$, where

$$W_j(\lambda) = \left(A_j(\lambda) - Z_j \frac{E\tilde{A}_1(\lambda)}{E\tau_1}\right) + \frac{EC_1(\lambda)}{ED_1(\lambda)}\left(B_j(\lambda) - Z_j \frac{E\tilde{B}_1(\lambda)}{E\tau_1}\right) - \frac{EB_1(-\lambda)EC_1(-\lambda)}{(ED_1(-\lambda))^2} D_j(-\lambda)$$
$$- \frac{EB_1(\lambda)EC_1(\lambda)}{(ED_1(\lambda))^2} D_j(\lambda) + \frac{EC_1(-\lambda)}{ED_1(-\lambda)}\left(B_j(-\lambda) - Z_j \frac{E\tilde{B}_1(-\lambda)}{E\tau_1}\right) + \frac{EB_1(\lambda)}{ED_1(\lambda)}\left(C_j(\lambda) - Z_j \frac{E\tilde{C}_1(\lambda)}{E\tau_1}\right)$$
$$+ \frac{EB_1(-\lambda)}{ED_1(-\lambda)}\left(C_j(-\lambda) - Z_j \frac{E\tilde{C}_1(-\lambda)}{E\tau_1}\right) - 2\pi f(\lambda)\tau_j$$

and

$$\tilde{B}_j(\lambda) = \int_{T(j-1)}^{T(j)} e^{i\lambda(T(j)-s)} ds, \quad \tilde{C}_j(\lambda) = \int_{T(j-1)}^{T(j)} e^{i\lambda(s-T(j-1))} ds.$$

Under (19), $EW_j(\lambda) = 0$ and $Var(W_j(\lambda)) < \infty$.The CLT for sums of iid rv's, in combination with the random time change theorem (see, for example, p.144 of Billingsley (1968)), yields our next result.

Theorem 2 Under (19),

$$t^{\frac{1}{2}} \left(\hat{f}_t(\lambda) - f(\lambda)\right) \Rightarrow \frac{1}{2\pi} \sqrt{\frac{EW_1(\lambda)^2}{E\tau_1}} N(0,1)$$

as $t \to \infty$.

Note that Theorem 2 asserts that $\hat{f}_t(\lambda)$ converges to $f(\lambda)$ at "square root convergence rate"; see Brockwell and Davis (1991) for a discussion of the "sub-square root" rate associated with conventional spectral density estimation in the non-regenerative setting. We also observe that $EW_1(\lambda)^2/E\tau_1$ can be estimated in a straightforward fashion, in order to produce asymptotically valid confidence intervals for $f(\lambda)$ based on the normal approximation associated with Theorem 2. In particular, while $W_i(\lambda)$ contains the "nuisance parameters" $f(\lambda)$, $E\tilde{A}_1(\lambda)$, $E\tau_1$, $EB_1(\lambda)$, $E\tilde{B}_1(\lambda)$, $EC_1(\lambda)$, $E\tilde{C}_1(\lambda)$, $ED_1(\lambda)$, $EB_1(-\lambda)$, $E\tilde{B}_1(-\lambda)$, $EC_1(-\lambda)$, $E\tilde{C}_1(-\lambda)$ and $ED_1(-\lambda)$, each of these nuisance quantities can be replaced by their corresponding sample means, thereby obtaining $W_i(\lambda,t)$. We then estimate $EW_1(\lambda)^2/E\tau_1$ via $\sum_{j=1}^{N(t)} W_j(\lambda,t)^2/t$.

Furthermore, we observe that $\hat{f}_t(\cdot)$ is differentiable. In this context, we expect that $f_t'(\lambda) \to f'(\lambda)$ almost surely as $t \to \infty$, so our spectral density estimator will typically be well-behaved as a function of λ, leading to good global estimators of the spectral density function.

4 REGENERATIVE ESTIMATION FOR HITTING TIME EXPECTATIONS

We next discuss the use of regenerative methods in efficiently computing moments of hitting times. In particular, given an S-valued classically regenerative process $X = (X(t) : t \geq 0)$, let $T = \inf\{t \geq 0 : X(t) \in A\}$ be the "hitting time" of the set $A \subseteq S$. Our goal here is to compute ET^k for some integer moment index k. Such hitting times are of interest in many applications settings. For example, in the dependability modeling context, the time to system failure is a key performance measure, while in the queueing context, the time to buffer overflow is of interest; such random times can obviously be represented as hitting times.

We note that if X is non-delayed, then the cycle independence implies that

$$E\exp(\theta T) = E\exp(\theta T)I(T < \tau_1) + E\exp(\theta\tau_1)I(T \geq \tau_1) \cdot E\exp(\theta T). \tag{20}$$

For $\theta \in \mathbb{R}_+$, set $\varphi(\theta) = E\exp(\theta T)$, $\eta(\theta) = E\exp(\theta T)I(T < \tau_1)$ and $v(\theta) = E\exp(\theta\tau_1)I(T \geq \tau_1)$. If $\varphi(\theta) < \infty$, then $\eta(\theta)$ is finite-valued also, and $\eta(\theta) > 0$ implies $v(\theta) < 1$, so that

$$\varphi(\theta) = \frac{\eta(\theta)}{1 - v(\theta)}. \tag{21}$$

Furthermore, $\varphi(\cdot)$, $\eta(\cdot)$, and $v(\cdot)$ are then infinitely differentiable at 0, and (20) implies that

$$\varphi^{(k)}(0) = \eta^{(k)}(0) + \sum_{j=0}^{k} \binom{k}{j} v^{(j)}(0)\varphi^{(k-j)}(0).$$

It follows that

$$ET^k = \frac{1}{P(T < \tau_1)} \left(E(\tau_1 \wedge T)^k + \sum_{j=1}^{k-1} \binom{k}{j} E\tau_1^j I(T \geq \tau_1) ET^{k-j} \right); \tag{22}$$

where $a \wedge b = \min(a,b)$. (If we use a direct argument to establish (22) that avoids $E\exp(\theta T)$, one finds that (22) holds whenever $ET^k < \infty$.) Specializing (22) to $k = 1,2$, we find that

$$ET = \frac{E(\tau_1 \wedge T)}{P(T < \tau_1)},$$
$$ET^2 = \left(E(\tau_1 \wedge T)^2 + E\tau_1 I(T \geq \tau_1) \frac{E(\tau_1 \wedge T)}{P(T < \tau_1)} \right) \frac{1}{P(T < \tau_1)}.$$

When A is "rare" (so that $P(T < \tau_1)$ is small), one can use (22) in combination with importance sampling to obtain significant variance reductions. In particular,we would estimate the terms $E(T \wedge \tau_1)^j$ and $E\tau_1^j I(T \geq \tau_1)$ for $1 \leq j \leq k$ from their sample means obtained by repeatedly sampling $(X(s) : 0 \leq s \leq T \wedge \tau_1)$ (without importance sampling), but estimate $P(T < \tau_1)$ by sampling $(X(s) : 0 \leq s \leq T \wedge \tau_1)$ under an appropriately chosen importance distribution.

This idea has been used with great success when $k = 1$; see, for example, Glynn et al. (1993). Our contribution here is to show how it easily extends to $k > 1$ in the presence of regenerative structure. Furthermore, it should be noted that it is straightforward to extend this approach to delayed regenerative processes, in which case the identity

$$E\exp(\theta T) = E\exp(\theta T)I(T < T(0)) + E\exp(\theta T)I(T > T(0)) \frac{E\exp(\theta \tilde{T})I(\tilde{T} < \tau_1)}{1 - E\exp(\theta \tau_1)I(\tilde{T} > \tau_1)}$$

replaces (21), where $\tilde{T} = \inf\{t \geq 0 : X(T(0)+t) \in A\}$.

5 COMPUTING VALUE FUNCTIONS

Suppose that $X = (X(t) : t \geq 0)$ is an irreducible positive recurrent Markov process taking values in a discrete state space S. In numerically calculating optimal controls for such an S via Monte Carlo, a key issue that arises is the computation of the "value function" associated with a given policy. One important class of control problems concerns optimization for infinite-horizon discounted reward/cost, in which case the value function takes the form

$$v(x) \triangleq E_x \int_0^\infty \exp\left(-\int_0^t g(X(s))ds \right) h(X(t))dt, \tag{23}$$

where $h : S \to \mathbb{R}_+$ and $g(x)$ is the instantaneous discount rate associated with state x. Here, $E_x(\cdot) \triangleq E(\cdot|X(0) = x)$. We note that the rv appearing in (23) involves simulating X over an infinite time horizon, so direct sampling is infeasible.

To compute $v(\cdot)$ via regenerative methods, let $\tau(x) = \inf\{t \geq 0 : X(t) = x, X(t-) \neq x\}$ be the first entry time into state x. Then, because $\tau(z)$ is a regeneration time for X,

$$\begin{aligned}
v(x) &= E_x \int_0^{\tau(z)} \exp\left(-\int_0^t g(X(s))ds\right) h(X(t))dt \\
&\quad + E_x \exp\left(-\int_0^{\tau(z)} g(X(s))ds\right) \int_{\tau(z)}^{\infty} \exp\left(-\int_{\tau(z)}^t g(X(s))ds\right) h(X(t))dt \\
&= E_x \int_0^{\tau(z)} \exp\left(-\int_0^t g(X(s))ds\right) h(X(t))dt \\
&\quad + E_x \exp\left(-\int_0^{\tau(z)} g(X(s))ds\right) E_z \int_0^{\infty} \exp\left(-\int_0^t g(X(s+\tau(z)))ds\right) h(X(t+\tau(z)))dt \\
&= E_x \int_0^{\tau(z)} \exp\left(-\int_0^t g(X(s))ds\right) h(X(t))dt + E_x \exp\left(-\int_0^{\tau(z)} g(X(s))ds\right) v(z) \\
&\triangleq E_x\beta(z) + E_x\Gamma(z) \cdot v(z)
\end{aligned} \tag{24}$$

so that

$$v(z) = \frac{E_z\beta(z)}{1 - E_z\Gamma(z)}$$

whenever $v(z) < \infty$. The quantity $v(z)$ can then be computed by generating iid copies of $(\beta(z), 1-\Gamma(z))$ conditional on $X(0) = z$, followed by forming the corresponding ratio estimator; confidence intervals can then be obtained through the standard methodology available for ratio estimators, as in Section 2.

To compute $v(\cdot)$ at multiple points $x_1, x_2, \cdots, x_m$, we can either use the above ratio estimation methodology at each of the m points, or we can fix $z \in S$, and note that (23) implies that

$$v(x_i) = E_{x_i}\beta(z) + E_{x_i}\Gamma(z)v(z).$$

Hence, by simulating $(\beta(z), \Gamma(z))$ conditional on $X(0) = x_i\,(1 \leq i \leq m)$ and combining this with our above estimator for $v(x)$, we obtain estimators for the $v(x_i)'s$. We can additionally stratify our sampling, so as to optimize the proportion of simulation effort expended on runs starting from the x_i's versus the effort expended on simulation effort starting from z; see (Fox and Glynn 1989) for details.

In some applications, a steady-state formulation is more appropriate. In such settings, the analog to (23) is the "relative value" function defined by

$$w(x) \triangleq \int_0^{\infty} E_x h_c(X(t))dt, \tag{25}$$

where $h_c(x) = h(x) - Eh(X(\infty))$ and $X(\infty)$ is a rv having the stationary distribution of X. When $|S| < \infty$, there exists a function k (solving Poisson's equation) such that

$$k(X(t)) + \int_0^t h_c(X(s))ds$$

is a martingale, so that

$$E_x k(X(t)) + E_x \int_0^t h_c(X(s))ds = k(x). \tag{26}$$

If X is aperiodic in discrete time or is a Markov jump process in continuous time, $E_x k(X(t)) \to Ek(X(\infty))$, where $X(\infty)$ has the stationary distribution of X. Sending $t \to \infty$ in (26), we conclude that $w(x) = k(x) - Ek(X(\infty))$. But applying optional sampling to the martingale yields the identity

$$k(z) + E_x \int_0^{\tau(z)} h_c(X(s))ds = k(x). \tag{27}$$

Hence, applying (12), we find that

$$
\begin{aligned}
Ek(X(\infty)) &= \frac{E_z \int_0^{\tau(z)} k(X(s))ds}{E_z \tau(z)} \\
&= \frac{E_z \int_0^{\tau(z)} \int_s^{\tau(z)} h_c(X(u))duds}{E_z \tau(z)} + k(z) \\
&= \frac{E_z \int_0^{\tau(z)} uh_c(X(u))du}{E_z \tau(z)} + k(z)
\end{aligned}
$$

and consequently

$$
w(x) = E_x \int_0^{\tau(z)} h_c(X(s))ds - \frac{E_z \int_0^{\tau(z)} uh_c(X(u))du}{E_z \tau(z)}. \tag{28}
$$

The above expression includes the nuisance parameter $Eh(X(\infty))$, so this needs to be separately estimated (within the runs starting from z) using the ideas of Section 1. The estimation of this nuisance parameter then affects the resulting CLT for the estimator of $w(X)$ (as in our spectral density analysis).

We further note that (27) proves that w differs from $\tilde{k}$ and $\tilde{l}$ by an additive constant, where

$$
\begin{aligned}
\tilde{k}(x) &= E_x \int_0^{\tau(z)} h_c(X(s))ds, \\
\tilde{l}(x) &= -E_z \int_0^{\tau(x)} h_c(X(s))ds.
\end{aligned} \tag{29}
$$

So, one has two different estimation approaches (one based on $\tilde{k}$ and one based on $\tilde{l}$) for approximating $w(\cdot)$ at multiple points $x_1, x_2, \cdots, x_m$ (up to an additive constant); see Glynn and Meyn (1996) for an earlier derivation of $\tilde{k}(\cdot)$ Presumably, because z can be chosen to be a state to which X is naturally attracted, building an estimation methodology for $w(\cdot)$ on the basis of $\tilde{k}$ is typically better. Again, as elsewhere in this paper, confidence intervals can be constructed by taking advantage of the CLT.

6 A NUMERICAL EXAMPLE

In this section, we present a brief account of the numerical performance of our regenerative estimator for the spectral density (18). Consider an $M/M/1/m$ queue with capacity $m = 50$ and service rate $\mu = 2$. We implement two systems with arrival rates 1 and 1.8. $Y(\cdot)$ denotes the number-in-system process. For each system, we simulate 100 replications of 50000 simulated time units. Within each replication, we compute a regenerative point estimator, the jack-knife estimator for the associated variance, and a corresponding 95% confidence interval (CI) based on our CLT. In the table, the "Ave. Est." stands for the average of all estimators over 100 replications; the "Half-width of CI" stands for the average confidence interval half-width; and "Coverage" stands for the fraction of CIs that cover the true value of the spectral density. The true values are computed numerically using methods in Glynn (1984). We provide additional numerical examples on regenerative estimators of value functions and hitting time expectations as an online supplement.

Table 1: Regenerative estimator for spectral density.

Arrival/Service Rate Ratio	λ	True Value of $f(\lambda)$	Ave. Est.	Half-width of CI	Coverage
$\rho = 0.5$	0.5	3.3424	3.3392	0.0692	97/100
	4	0.1148	0.1147	0.0010	93/100
$\rho = 0.9$	0.5	12.9117	13.2338	0.4822	97/100
	4	0.2217	0.2197	0.0067	99/100

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support of the National Science Foundation under Award Number DMS-1320158.

REFERENCES

Anderson, T. W. 1971. *The Statistical Analysis of Time Series*. John Wiley & Sons, New York.

Asmussen, S. 2003. *Applied Probability and Queues*. Springer, New York.

Billingsley, P. 1968. *Convergence of Probability Measures*. John Wiley & Sons, New York.

Brockwell, P. J., and R. A. Davis. 1991. *Time Series: Theory and Methods*. Springer-Verlag.

Chung, K. L. 1979. *A Course in Probability Theory*. Academic Press, New York.

Crane, M. A., and D. L. Iglehart. 1974. "Simulating Stable Stochastic Systems, I: General Multiserver Queues". *Journal of the ACM* 21 (1): 103–113.

Fishman, G. S. 1974. "Output Analysis of Queueing Simulations". *Management Science* 20:363–369.

Fox, B. L., and P. W. Glynn. 1989. "Simulating Discounted Costs". *Management Science* 35 (11): 1297–1315.

Glynn, P., and D. L. Iglehart. 1987. "A Joint Central Limit Theorem for the Sample Mean and Regenerative Variance Estimator". *Annals of Operations Research* 8 (1): 41–55.

Glynn, P. W. 1984. "Some Asymptotic Formulas for Markov Chains with Applications to Simulation". *Journal of Statistical Computation and Simulation* 19 (2): 97–112.

Glynn, P. W. 1994. "Some Topics in Regenerative Steady-state Simulation". *Acta Applicandae Mathematica* 34 (1-2): 225–236.

Glynn, P. W., P. Heidelberger, V. F. Nicola, and P. Shahabuddin. 1993. "Efficient Estimation of the Mean Time Between Failures in Non-regenerative Dependability Models". In *Proceedings of the 1993 Winter Simulation Conference*, 311–316.

Glynn, P. W., and D. L. Iglehart. 1993. "Conditions for the Applicability of the Regenerative Method". *Management Science* 39 (9): 1108–1111.

Glynn, P. W., and S. P. Meyn. 1996. "A Liapounov Bound for Solutions of the Poisson Equation". *The Annals of Probability*:916–931.

Glynn, P. W., and W. Whitt. 1993. "Limit Theorems for Cumulative Processes". *Stochastic Processes and their Applications* 47 (2): 299–314.

Glynn, P. W., and W. Whitt. 2002. "Necessary Conditions in Limit Theorems for Cumulative Processes". *Stochastic Processes and Their Applications* 98 (2): 199–209.

Meyn, S., and R. L. Tweedie. 2009. *Markov Chains and Stochastic Stability*. Cambridge University Press, Cambridge.

Thorisson, H. 2000. *Coupling, Stationarity, and Regeneration*, Volume 200. Springer, New York.

AUTHOR BIOGRAPHIES

ZEYU ZHENG is currently a Ph.D. candidate in Management Science and Engineering at Stanford University, specializing in Operations Research. He holds a B.Sc. in Mathematics from Peking University. His research interests include stochastic simulation, statistical inference and financial economics. His email address is zyzheng@stanford.edu.

PETER W. GLYNN is the Thomas Ford Professor in the Department of Management Science and Engineering (MS&E) at Stanford University. He is a Fellow of INFORMS and of the Institute of Mathematical Statistics, has been co-winner of Best Publication Awards from the INFORMS Simulation Society in 1993 and 2008, and was the co-winner of the John von Neumann Theory Prize from INFORMS in 2010. In 2012, he was elected to the National Academy of Engineering. His research interests lie in stochastic simulation, queueing theory, and statistical inference for stochastic processes. His email address is glynn@stanford.edu.