S.G. Henderson and B.L. Nelson (Eds.), *Handbook in OR & MS, Vol. 13*

DOI: 10.1016/S0927-0507(06)13016-9

Chapter 16

Simulation Algorithms for Regenerative Processes

Peter W. Glynn

Department of Management Science and Engineering, Stanford University, USA
E-mail: glynn@stanford.edu

Abstract

This chapter is concerned with reviewing the basic ideas and concepts underlying the use of regenerative structure in the development of efficient simulation algorithms. While it has long been known that discrete state space Markov chains exhibit regenerative structure, we argue that well-behaved discrete-event simulations typically also contain an embedded sequence of regeneration times. However, algorithmic identification of the corresponding regeneration times turns out to be a nontrivial problem. We discuss the theoretical and implementation issues involved in identifying the corresponding regeneration times, and describe how regenerative methodology supplies effective solutions to several difficult simulation problems. In particular, we discuss the use of regeneration in the context of steady-state simulation as a means of efficiently computing confidence intervals and correcting for initial bias. We also point out that regeneration has the potential to offer significant algorithmic efficiency improvements to the simulationist, and illustrate this idea via discussion of steady-state gradient estimation and computation of infinite horizon expected discounted reward.

1 Introduction

Let $V = (V(t)\colon\ t \geqslant 0)$ be a real-valued stochastic process in which $V(t)$ represents the simulation output collected at time t. Roughly speaking, the process V is said to be (classically) regenerative if there exist random times $T(0) < T(1) < \cdots$ at which the process "regenerates" (in the sense that V probabilistically starts afresh at each time $T(i)$, $i \geqslant 0$, and evolves independently of the process prior to time $T(i)$). Such regenerative structure implies that V can be viewed conceptually as a sequence of independent "cycles" $(V(s)\colon\ T(i-1) \leqslant s < T(i))$ that are "pasted together" (where we adopt the convention that $T(-1) = 0$). Thus, the infinite time behavior of V over $[0, \infty)$ is implicitly captured in the behavior of V over a cycle. Hence, in principle, virtually any expectation of V over $[0, \infty)$ can be alternatively described as

an expectation involving cycle-related quantities. This observation is the key insight that underlies regenerative simulation.

The use of regenerative structure as an algorithmic tool in the simulation setting has primarily focused on its use in the context of steady-state simulation. The first suggestion that regenerative cycles could play a useful role in steady-state simulation output analysis came from Cox and Smith (1961), and the idea was further developed in Kabak (1968). However, the first comprehensive development of the regenerative method for steady-state simulation output analysis came in a series of papers of Crane and Iglehart (1974a, 1974b, 1975), as well as concurrent work by Fishman (1973, 1974). The great majority of subsequent work on algorithmic exploitation of regeneration has followed the historic tradition of focusing on its application to steady-state simulation output analysis.

In this chapter we focus our discussion on the key theoretical and algorithmic issues underlying the use of regeneration in the steady-state simulation context. We start, in Section 2, by describing the key challenges that confront a simulationist in analyzing steady-state simulation output, while Section 3 discusses the basic regenerative approach to forming an estimator for the so-called "time-average variance constant". Section 4 offers some discussion of how the particular choice of regenerative structure influences the efficiency of the method, and Section 5 describes the regenerative solution to the initial transient problem and the construction of low-bias steady-state estimators. In Sections 6–8 we discuss the theoretical issue of when a simulation is regenerative, with a particular focus on when a discrete-event simulation contains algorithmically identifiable regenerative structure. Section 9 then discusses steady-state regenerative analysis from the perspective of martingale theory.

The last two sections of the chapter are intended to give the reader a hint of the role that regeneration can play in the development of computationally efficient algorithms for other simulation problems. In particular, we show that in computing either steady-state gradients or infinite-horizon expected discounted reward that regeneration offers the simulationist the opportunity to not only construct asymptotically valid confidence statements but to also improve computational efficiency. While regeneration is primarily understood within the simulation community as offering a vehicle for analysis of simulation output, our two examples are intended to argue that regeneration has the potential to also play a significant role in the variance reduction context.

2 The steady-state simulation problem

Let $V = (V(t)\colon\ t \geqslant 0)$ be a real-valued stochastic process in which $V(t)$ represents the value of the simulation output process at (simulated) time t. For example, $V(t)$ could represent the total work-in-process at time t in a production context or the inventory position at time t in a supply chain setting. Throughout this chapter, we use a continuous time formulation to describe

the relevant theory. (Note that any discrete-time sequence $(V_n\colon\ n \geqslant 0)$ can be embedded into continuous time via the definition $V(t) = V_{\lfloor t \rfloor}$ for $t \geqslant 0$, where $\lfloor x \rfloor$ is the greatest integer less than or equal to x.)

Many simulation applications demand that the simulationist compute a steady-state performance measure. To be specific, suppose that V satisfies a law of large numbers (LLN), so that there exists a (deterministic) constant α for which

$$\frac{1}{t}\int_0^t V(s)\,\mathrm{d}s \Rightarrow \alpha \tag{1}$$

as $t \to \infty$, where "$\Rightarrow$" denotes weak convergence. The constant α appearing in (1) is known as the steady-state mean of V. Computing α is the central problem in steady-state simulation.

Given the LLN (1), the time-average $\overline{V}(t) \triangleq t^{-1}\int_0^t V(s)\,\mathrm{d}s$ is the natural estimator for α. However, in view of the fact that the simulation of V will usually be initialized in a state that is atypical of equilibrium behavior, the process V will at best exhibit the stationarity associated with steady-state dynamics only in an approximate sense. As a consequence, such a simulation of V over $[0, t]$ will necessarily include some "initial transient period" over which the simulation outputs will be biased as estimators of steady-state performance. This, in turn, induces bias in the estimator $\overline{V}(t)$ (known as "initial bias"). While the effect of the initial transient can be expected to dissipate as $t \to \infty$, it can have a significant "small sample" impact on the quality of the estimator $\overline{V}(t)$.

To reduce the effect of the initial transient on the steady-state estimation algorithm, it is commonly recommended that the simulationist expend his computer budget on one (long) replication of V (for which the time horizon t can be made large), rather than multiple short replications. Because of the fact that only one realization of the process V is then simulated, estimating the variance of the associated estimator can then be challenging.

In particular, it is usually the case that a process V satisfying the LLN (1) will also satisfy a central limit theorem (CLT). Specifically, there exists a (deterministic) constant $\sigma \in (0, \infty)$ for which

$$t^{1/2}\bigl(\overline{V}(t) - \alpha\bigr) \Rightarrow \sigma\mathrm{N}(0, 1) \tag{2}$$

as $t \to \infty$, where $\mathrm{N}(0, 1)$ denotes a normal random variable (r.v.) having mean zero and unit variance. The constant σ^2 is called the time-average variance constant (TAVC) of V. In view of (2), it is easily verified that

$$\left[\overline{V}(t) - \frac{z\sigma}{\sqrt{t}}, \overline{V}(t) + \frac{z\sigma}{\sqrt{t}}\right] \tag{3}$$

is an (asymptotic) $100(1 - \delta)\%$ confidence interval for α, provided that z is chosen so that $\mathrm{P}(-z \leqslant \mathrm{N}(0, 1) \leqslant z) = 1 - \delta$. Of course, (3) can be computed at a practical level only if the TAVC σ^2 is known.

Since knowledge of σ^2 is virtually never available, the simulationist must instead estimate σ^2 from the observed simulation up to time t. If V is stationary,

σ^2 can (in great generality) be represented in terms of the spectral density $\tilde{f}(\cdot)$. Specifically, $\sigma^2 = 2\pi\tilde{f}(0)$, where

$$\tilde{f}(\lambda) = \frac{1}{\pi}\int_0^\infty \cos(\lambda u)\,\mathrm{cov}\big(V(0), V(u)\big)\,\mathrm{d}u. \tag{4}$$

Spectral density estimation for stationary processes has been well studied in the literature; see, for example, Chapter 9 of Anderson (1971). Such estimators converge at a rate of $t^{-1/3}$ or $t^{-2/5}$, with the specific rate depending on the degree of differentiability of $\tilde{f}$. While virtually all steady-state simulations involve simulating nonstationary stochastic processes that contain an initial transient period, one would expect that the best possible convergence rate for an estimator of the TAVC σ^2 will be no faster than that which is achievable in the stationary setting. Hence, more problem structure must be assumed in order to obtain a TAVC that converges (for example) at rate $t^{-1/2}$.

In the next section, we show how regenerative structure can be exploited to obtain an estimator for the TAVC σ^2 that converges at rate $t^{-1/2}$ in the simulation time horizon t. Given the substantial body of theory establishing that $t^{-1/2}$ is typically an optimal rate of convergence for statistical estimators (see, for example, Chapter 2 of Ibragimov and Has'minskii, 1981), this suggests that regenerative structure permits the simulationist to obtain TAVC estimators that converge at the best possible rate.

3 The regenerative estimator for the TAVC

To obtain a TAVC estimator that converges to σ^2 at rate $t^{-1/2}$, one needs to assume additional structure about the process V. To illustrate this idea, suppose that the simulation output process V is a (continuous-time) autoregressive process satisfying

$$\mathrm{d}V(s) = -\gamma V(s)\,\mathrm{d}s + \mathrm{d}W(s), \tag{5}$$

where $\gamma > 0$ and $W = (W(s)\colon s \geqslant 0)$ is a square integrable process with stationary independent increments for which $\mathrm{E}W(s) = \mu s$ and $\mathrm{var}\,W(s) = \eta^2 s$ for $s \geqslant 0$. It is easily verified that

$$V(t) = \mathrm{e}^{-\gamma t}V(0) + \int_0^t \mathrm{e}^{-\gamma(t-s)}\,\mathrm{d}W(s)$$

and that V satisfies (2) with $\alpha = \mu/\gamma$ and $\sigma^2 = \eta^2/\gamma$. Hence, if the simulation output process can be assumed to evolve according to (5), we can estimate σ^2 via $\hat{\eta}^2/\hat{\gamma}$, where $\hat{\eta}^2$ and $\hat{\gamma}$ are appropriately chosen estimators for the parameters η^2 and γ underlying (5). If V satisfies (5), it can be shown (in great generality) that the resulting TAVC estimator converges at a rate $t^{-1/2}$. The problem, of course, is that it is rarely (if ever) the case that the simulation output process V evolves precisely according to (5). As a consequence, a TAVC

estimation algorithm based on assuming that V's dynamics are governed by (5) can, at best, provide only an approximation to the true σ^2. Nevertheless, this autoregressive approach to estimating σ^2 can, when appropriately applied, offer an effective means of estimating σ^2.

The key idea exploited above is the fact that the TAVC σ^2 can be easily and exactly computed for the class of processes described by (5). One then uses a "plug-in" estimator to estimate the unknown quantities appearing in the corresponding expression for the TAVC.

The importance of the regenerative method lies in the fact that a large class of interesting and useful steady-state simulations fall into the class of regenerative processes, and that a simple expression for the TAVC of such processes can be derived. For example, suppose that $V(t) = g(X(t))$, where $X = (X(t)\colon\ t \geqslant 0)$ is an irreducible positive recurrent continuous-time Markov chain (CTMC) living on discrete state space S, and where $g : S \to \mathbb{R}$ is a given performance measure. Fix a state $z \in S$. Then, V is regenerative with cycles defined by the consecutive times $(T(n)\colon\ n \geqslant 0)$ at which X enters z. The class of CTMCs (and its discrete-time cousin, the class of discrete-time Markov chains (DTMCs)) form an important class of models that are commonly simulated and that enjoy regenerative structure.

Simple expressions for the steady-state mean α and TAVC σ^2 can be derived in the regenerative setting. For α, note that (1) suggests that

$$\frac{1}{T(n)} \int_0^{T(n)} V(s)\,\mathrm{d}s \Rightarrow \alpha \tag{6}$$

as $n \to \infty$. But the left-hand side of (6) equals

$$\frac{n^{-1} \sum_{i=1}^{n} Y_i}{n^{-1} \sum_{i=1}^{n} \tau_i},$$

where

$$Y_i \triangleq \int_{T(i-1)}^{T(i)} V(s)\,\mathrm{d}s \quad \text{and}$$
$$\tau_i \triangleq T(i) - T(i-1).$$

Since $(Y_i\colon\ i \geqslant 1)$ is a sequence of independent and identically distributed (i.i.d.) r.v.'s, $\overline{Y}_n \triangleq n^{-1} \sum_{i=1}^{n} Y_i \Rightarrow \mathrm{E}Y_1$ as $n \to \infty$ (provided $\mathrm{E}|Y_1| < \infty$). Similarly, we expect that $\bar{\tau}_n \triangleq n^{-1} \sum_{i=1}^{n} \tau_i \Rightarrow \mathrm{E}\tau_1$ as $n \to \infty$, so that the identity

$$\alpha = \frac{\mathrm{E}Y_1}{\mathrm{E}\tau_1} \tag{7}$$

must hold.

To heuristically derive a corresponding expression for σ^2, note that a regenerative process is (in great generality) asymptotically stationary. Given (4), we

expect to be able to represent σ^2 as

$$\sigma^2 = 2\int_0^\infty \operatorname{cov}\bigl(V^*(0), V^*(r)\bigr)\,dr, \tag{8}$$

where $V^* = (V^*(r)\colon r \geqslant 0)$ is a stationary version of V. Put $V_c^*(r) = V^*(r) - \alpha$ and rewrite (8) as

$$\sigma^2 = 2\int_0^\infty \mathrm{E}V_c^*(0)V_c^*(r)\,dr. \tag{9}$$

The stationary version V^* is itself regenerative, with regeneration times $0 < T^*(0) < T^*(1) < \cdots$. In view of the independence across cycles and the fact that $\mathrm{E}V_c^*(r) = 0$, we might hope that the right-hand side of (9) simplifies to

$$2\,\mathrm{E}\int_0^{T^*(0)} V_c^*(0)V_c^*(r)\,dr. \tag{10}$$

Given the approximate stationarity of V, (10) should be approximately equal to

$$2\,\mathrm{E}\int_s^{T(N(s)+1)} V_c(s)V_c(r)\,dr$$

when s is large, where $N(s) = \max\{n \geqslant -1\colon T(n) \leqslant s\}$. Averaging over $s \in [0, T(n)]$, this suggests that

$$\mathrm{E}\frac{2}{T(n)}\int_0^{T(n)}\int_s^{T(N(s)+1)} V_c(s)V_c(r)\,dr\,ds$$

should be close to σ^2 when t is large. But

$$\begin{aligned}
&2\int_0^{T(n)}\int_s^{T(N(s)+1)} V_c(s)V_c(r)\,dr\,ds \\
&\quad = 2\sum_{i=0}^{n}\int_{T(i-1)}^{T(i)} V_c(s)\int_s^{T(i)} V_c(r)\,dr\,ds \\
&\quad = \sum_{i=0}^{n} Z_i^2,
\end{aligned}$$

where $Z_i \triangleq Y_i - \alpha\tau_i$. The i.i.d. cycle structure implies that $n^{-1}\sum_{i=1}^n Z_i^2 \Rightarrow \mathrm{E}Z_1^2$ as $n \to \infty$. Equation (10) therefore suggests that the equality

$$\sigma^2 = \frac{\mathrm{E}Z_1^2}{\mathrm{E}\tau_1} \tag{11}$$

should hold.

Given the equality (11), the obvious "plug-in" estimator for the TAVC σ^2 based on simulating V over $[0, t]$ is

$$\frac{\sum_{i=1}^{N(t)}(Y_1 - \overline{V}(t)\tau_i)^2/N(t)}{\sum_{i=1}^{N(t)}\tau_i/N(t)}$$

or its asymptotically equivalent variant

$$\hat{\sigma}^2(t) = \frac{1}{t}\sum_{i=1}^{N(t)}\bigl(Y_1 - \overline{V}(t)\tau_i\bigr)^2.$$

The following theorem makes rigorous the validity of the regenerative TAVC estimator $\hat{\sigma}^2(t)$. For the proof, see Glynn and Whitt (1993, 2002) and Glynn and Iglehart (1993).

Theorem 1. *Suppose that V is regenerative with respect to the regeneration time sequence $(T(n)\colon\ n \geqslant 0)$. Assume that $\mathrm{E}\tau_1 < \infty$. Then, there exist (deterministic) constants α and σ^2 such that*

$$t^{1/2}\bigl(\overline{V}(t) - \alpha\bigr) \Rightarrow \sigma\mathrm{N}(0, 1) \tag{12}$$

as $t \to \infty$ if and only if $\mathrm{E}|Y_1| < \infty$ and $\mathrm{E}Z_1^2 < \infty$, in which case

$$\alpha = \frac{\mathrm{E}Y_1}{\mathrm{E}\tau_1}, \qquad \sigma^2 = \frac{\mathrm{E}Z_1^2}{\mathrm{E}\tau_1}.$$

Furthermore, if (12) holds, then

$$\hat{\sigma}^2(t) \Rightarrow \sigma^2$$

as $t \to \infty$.

Theorem 1 shows that $\hat{\sigma}^2(t)$ is consistent as an estimator for σ^2 precisely when the CLT (12) holds. When (12) holds with $\sigma^2 > 0$, then

$$\left[\overline{V}(t) - z\sqrt{\frac{\hat{\sigma}^2(t)}{t}}, \overline{V}(t) + z\sqrt{\frac{\hat{\sigma}^2(t)}{t}}\right]$$

is an approximate $100(1 - \delta)\%$ confidence interval for α.

4 Choice of the optimal regeneration state

Given a simulation of V over the time interval $[0, t]$, the natural point estimator for the steady-state mean α is, of course, the time-average $\overline{V}(t)$. While it may be desirable to modify $\overline{V}(t)$ to deal with initial transient or initial bias effects, one would expect such modifications to be of small order asymptotically.

Hence, any reasonable point estimator for α will either be exactly equal to $\overline{V}(t)$ or asymptotically equivalent to $\overline{V}(t)$. Of course, the r.v. $\overline{V}(t)$ is not influenced in any way by the choice of the regeneration or return state z.

On the other hand, the TAVC estimator $\hat{\sigma}^2(t)$ is defined relative to a specific choice of the regeneration or return state $z \in S$. A natural question that then arises is the determination of the state z^* that is the "best" choice of return state for estimating σ^2. This question can be resolved by studying the rate of convergence of $\sqrt{\hat{\sigma}^2(t)}$ to $\sqrt{\sigma^2}$.

Theorem 2. *Suppose that* $\mathrm{E}[Y_1^4 + \tau_1^4] < \infty$. *Then*

$$t^{1/2}\big(\overline{V}(t) - \alpha, \hat{\sigma}(t) - \sigma\big) \Rightarrow \mathrm{N}(0, D)$$

as $t \to \infty$, *where* $\mathrm{N}(0, D)$ *is a bivariate normal r.v. with mean* 0 *and covariance matrix* D *given by*

$$D = \frac{1}{\mathrm{E}\tau_1} \begin{pmatrix} \mathrm{E}Z_1^2 & \frac{\mathrm{E}A_1 Z_1 - \lambda \mathrm{E}Z_1^2}{2\sigma} \\ \frac{\mathrm{E}A_1 Z_1 - \lambda \mathrm{E}Z_1^2}{2\sigma} & \frac{\mathrm{E}A_1^2 - 2\lambda \mathrm{E}A_1 Z_1 + \lambda^2 \mathrm{E}Z_1^2}{4\sigma^2} \end{pmatrix},$$

where $A_i = Z_i^2 - \sigma^2 \tau_i$ *and* $\lambda = 2\,\mathrm{E}Z_1\tau_1/\mathrm{E}\tau_1$.

See Glynn and Iglehart (1987) for the proof. Theorem 2 establishes that the TAVC estimator does indeed converge at rate $t^{-1/2}$. It further describes the asymptotic variance of $\hat{\sigma}^2(t)$ in terms of the given regenerative cycle structure. The asymptotic variance can be explicitly computed for certain CTMC models; see Glynn and Iglehart (1986). These examples make clear that there is, unfortunately, no simple guidance available for how to choose the best possible regeneration state. In particular, the examples make clear that choosing the regeneration state $\tilde{z}$ that minimizes the mean return time is not necessarily the choice that minimizes the asymptotic variance of $\hat{\sigma}^2(t)$.

One odd characteristic of Theorem 2 is that the covariance entry D_{12} $(= D_{21})$ of the matrix D appearing there turns out to be independent of the choice of regeneration state. This result, due to Calvin (1994), has no obvious and apparent simple explanation, and is a consequence of a direct computation. By contrast, the entry D_{11} must clearly be independent of the choice of the regeneration state z, since it is the asymptotic variance of the r.v. $\overline{V}(t)$ that is defined independently of z.

5 The regenerative approach to the initial transient and initial bias problems

As discussed in Section 2, one of the major challenges in steady-state simulation is the mitigation of effects due to the initial transient and initial bias. We deal first with the better understood issue of how to reduce biasing effects due to a nonstationary initialization.

It is usual, in the presence of (1), that there exists $\nu > 0$ such that

$$\mathrm{E}V(t) = \alpha + \mathrm{O}\big(\mathrm{e}^{-\nu t}\big) \tag{13}$$

as $t \to \infty$, where $\mathrm{O}(h(t))$ represents a function that is bounded by a constant multiple of $|h(t)|$ as $t \to \infty$. For example, (13) is known to typically hold for geometrically ergodic Markov processes; see Meyn and Tweedie (1993). Given (13),

$$\int_0^t \mathrm{E}\big(V(s) - \alpha\big)\,\mathrm{d}s = b + \mathrm{O}\big(\mathrm{e}^{-\nu t}\big)$$

as $t \to \infty$, so that

$$\mathrm{E}\overline{V}(t) = \alpha + t^{-1}b + \mathrm{O}\big(\mathrm{e}^{-\nu t}\big) \tag{14}$$

as $t \to \infty$, where

$$b = \int_0^\infty \mathrm{E}\big(V(s) - \alpha\big)\,\mathrm{d}s.$$

An estimator with lower initialization bias can therefore be constructed if one can find an asymptotically unbiased estimator for b. Deriving such an estimator without imposing additional structure is an impossibility in the single replication context, because only one realization of the process from which to estimate b is available. On the other hand, if the process V is assumed to be regenerative, estimating b should (in principle) be possible, because the i.i.d. cycle structure suggests that the effect of initialization is now implicitly replicated (for example, by permuting the simulated cycles).

Appealing to renewal theory yields the following expression for b; for the proof see Glynn (1994).

Proposition 1. *Suppose that V is regenerative with respect to the regeneration times $0 = T(0) < T(1) < \cdots$. If τ_1 has a density and satisfies*

$$\mathrm{E}\tau_1\left(1 + \int_0^{\tau_1} \big|V(s)\big|\,\mathrm{d}s\right) < \infty,$$

then (14) holds with

$$b = -\frac{1}{\mathrm{E}\tau_1}\left(\mathrm{E}\int_0^{\tau_1} sV(s)\,\mathrm{d}s - \alpha\,\mathrm{E}\frac{\tau_1^2}{2}\right).$$

In view of Proposition 1, it is now clear how one can potentially reduce the effects of initial bias. In particular, consider the "plug-in" estimator for b defined by

$$\hat{b}(t) = -\frac{1}{T(N(t))}\sum_{i=1}^{N(t)}\left(\int_0^{\tau_i} sV\big(T(i-1)+s\big)\,\mathrm{d}s - \overline{V}(t)\frac{\tau_i^2}{2}\right).$$

We expect that

$$\mathrm{E}\big(\overline{V}(t) - t^{-1}\hat{b}(t)\big) = \alpha + \mathrm{o}\left(\frac{1}{t}\right) \tag{15}$$

as $t \to \infty$, where $\mathrm{o}(h(t))$ denotes a function that, when divided by $|h(t)|$, tends to zero as $t \to \infty$. Given that $\alpha_1(t) \triangleq \overline{V}(t) - t^{-1}\hat{b}(t)$ generally satisfies

$$\frac{\operatorname{var}\alpha_1(t)}{\operatorname{var}\overline{V}(t)} \to 1 \tag{16}$$

as $t \to \infty$, (15) establishes that $\alpha_1(t)$ has lower asymptotic bias while suffering no increase in asymptotic variance. For additional discussion, see Glynn (1994).

Consider the estimators:

$$\alpha_2(t) = \overline{V}\big(T\big(N(t)+1\big)\big),$$

$$\alpha_3(t) = \overline{V}\big(T\big(N(t)\big)\big) + \frac{1}{t^2}\sum_{i=1}^{N(t)}\big(Y_i - \overline{V}(t)\tau_i\big)\tau_i,$$

$$\alpha_4(t) = N(t)\overline{V}\big(T\big(N(t)\big)\big) - \frac{N(t)-1}{N(t)}\sum_{i=1}^{N(t)}\frac{\sum_{j\neq i}Y_j}{\sum_{j\neq i}\tau_j}.$$

Each of the above estimators duplicates the performance of $\alpha_1(t)$, in the sense that each satisfies both (15) and (16); see Meketon and Heidelberger (1982) and Glynn and Heidelberger (1990, 1992).

Turning next to the question of identifying the duration of the initial transient period, recall that the sequence of cycle variables $((Y_i, \tau_i)\colon i \geqslant 1)$ is i.i.d. Hence, on the time-scale of regenerative cycles, the initial transient disappears entirely. Furthermore, recall that the steady-state mean α can be expressed as the ratio of expectations (7) defined in terms of cycle-related quantities Y_i and τ_i. Hence, if one simulates V over n regenerative cycles (to time $T(n)$), the natural estimator for α is the ratio estimator $\overline{V}(T(n)) = \overline{Y}_n/\bar{\tau}_n$. The bias of this estimator arises as a consequence of the fact that $\overline{Y}_n/\bar{\tau}_n = h(\overline{Y}_n, \bar{\tau}_n)$, where h is the nonlinear function defined by $h(x_1, x_2) = x_1/x_2$. Thus, on the time-scale of regenerative cycles, initialization bias manifests itself as estimator bias due to nonlinearity. Such nonlinearity bias has long been studied within the statistical literature, and a large number of remedies for dealing with bias of this kind have been developed over the years: Taylor expansion methods (Cramér, 1946), the jack-knife (Miller, 1974), the bootstrap (Efron and Tibshirani, 1993), and sectioning (Lewis and Orav, 1989).

Thus, regenerative structure suggests a variety of different mechanisms for dealing with initial bias (and, on the regenerative cycle time scale, the initial transient).

6 When is a simulation regenerative?

As has been seen in preceding sections, regenerative structure turns out to be algorithmically useful in developing solutions to various aspects of the steady-state simulation problem. Furthermore, regenerative structure can be easily identified in the setting of discrete state space Markov chains, in either discrete or continuous time.

Of course, most real-world discrete-event simulations do not involve simulating a discrete state space Markov processes. Much more complicated models are typically simulated. However, one can persuasively argue that the great majority of discrete event simulations can, from a mathematical standpoint, be viewed as simulations of Markov processes (living on a continuous state space rather than a discrete state space). In particular, by adding supplementary variables to the "physical state" (e.g., the location of each customer in a network) of the system, one typically ends up with a state descriptor that evolves according to a Markov process. For example, one can supplement the physical state by adding the remaining time to completion of each currently scheduled event associated with the currently occupied physical state. Thus, one may abstractly view the typical discrete-event simulation as corresponding to the simulation of a Markov process $X = (X(t)\colon\ t \geqslant 0)$ living on a continuous state space S, where the continuous component of the state records the remaining time to event completion for each of the active event "clocks".

We assume throughout the reminder of the paper that the state space S is a separable metric space (so that, for example, open and closed subsets of $\mathbb{R}^d$ are covered by our theory).

Now that we have established that the typical discrete-event simulation can be viewed as a Markov process, we next argue that any Markov process for which the steady-state simulation problem is well defined necessarily exhibits regenerative structure. This, in turn, will show that any discrete-event simulation for which the steady-state simulation problem makes sense must contain regenerative structure.

We start with a precise definition of "well-posedness" for the steady-state simulation problem. For $x \in S$, let $\mathrm{P}_x(\cdot)$ and $\mathrm{E}_x(\cdot)$ be the probability and expectation operator, respectively, under which $X(0) = x$.

Definition 1. We say that the steady-state simulation problem is *well-posed* for the S-valued Markov process $X = (X(t)\colon\ t \geqslant 0)$ if for each bounded function $g : S \to \mathbb{R}$, there exists $\alpha(g)$ such that for $x \in S$,

$$\frac{1}{t}\int_0^t \mathrm{E}_x g\big(X(s)\big)\,\mathrm{d}s \to \alpha(g)$$

as $t \to \infty$.

According to the definition, the expectation of $V(s) = g(X(s))$ must converge (at least in an average sense) to a number $\alpha = \alpha(g)$ that is independent of the initial state x. This seems a reasonable definition of well-posedness, for

otherwise, there exists a bounded function g for which $t^{-1}\int_0^t \mathrm{E}_x g(X(s))\,\mathrm{d}s$ either does not converge or converges to a limit that depends on the initial state x. In either case, this precludes what is normally understood by the notion of steady-state.

Recall that a sequence $(W_i\colon i \geqslant 0)$ is said to be one-dependent if $(W_j\colon j < i)$ is independent of $(W_j\colon j > i)$ for each $i \geqslant 1$.

Theorem 3. *Let $X = (X(t)\colon t \geqslant 0)$ be an S-valued Markov process for which the steady-state simulation problem is well-posed. Then, there exist random times $0 \leqslant T(0) < T(1) < \cdots$ such that*:

(i) *The probability distribution* $\mathrm{P}_x((X(T(i-1)+s)\colon 0 \leqslant s \leqslant \tau_i) \in \cdot)$ *is independent of both $x \in S$ and $i \geqslant 1$;*
(ii) *The sequence of cycles* $((X(T(i-1)+s)\colon 0 \leqslant s \leqslant \tau_i)\colon i \geqslant 0)$ *is one-dependent*;
(iii) $\mathrm{E}_x\tau_1 < \infty$.

For the proof, see Glynn (1994).

This theorem asserts that any simulation for which the steady-state simulation problem is well-posed necessarily possesses regenerative structure. However, the regenerative structure identified by this result only guarantees the existence of one-dependent identically distributed cycles. Fortunately, much of the theory developed in Sections 3–5 generalizes from the classical regenerative structure (of i.i.d. cycles) to the setting of one-dependent regenerative processes. For example, the one-dependent analog to $\hat{\sigma}^2(t)$ continues to converge to the TAVC σ^2 at rate $t^{-1/2}$ in the one-dependent setting; see Henderson and Glynn (2001).

However, an alternative approach exists that permits one to use conventional regenerative methodology based on i.i.d. cycle structure. For one-dependent regenerative processes, the ratio formula for the steady-state mean $\alpha(g)$ continues to hold:

$$\alpha(g) = \frac{\mathrm{E}\int_0^{\tau_1} g(X(T(0)+s))\,\mathrm{d}s}{\mathrm{E}\tau_1}. \tag{17}$$

To estimate $\alpha(g)$, we simulate X over the cycle corresponding to the interval $[T(0), T(1)]$. At time $T(1)$, rather than continuing the simulation of X, one now terminates the simulation. One then independently draws a new initial condition from $\mathrm{P}(X(T(0)) \in \cdot)$ and simulates a second independent trajectory of X up to completion of its corresponding first cycle. By repeating this process, we are simulating independent realizations of X over its first cycle. By "pasting" these i.i.d. cycles back to back, one is generating a new process $\widetilde{X}$ that is regenerative in the classical sense (with i.i.d. cycles). Given the ratio formula (17), the steady-state of $\widetilde{X}$ exactly coincides with that of the one-dependent process X. Hence, if one simulates $\widetilde{X}$ rather than X, all the methods of Sections 3–5 apply without change.

7 When is a GSMP regenerative?

Section 6 makes clear that regeneration is the rule rather than the exception for well-behaved steady-state simulations. This, however, leaves open the question of when a specific simulation model has the structure necessary to guarantee that the associated steady-state simulation is well behaved.

We shall focus exclusively, in this section, on conditions under which discrete-event simulations possess the required structure. We take the point of view here that a discrete-event simulation is tantamount to simulation of a class of processes known as generalized semi-Markov processes (GSMPs). To describe a GSMP, we make concrete our discussion of Section 6, in which we argued that a discrete-event simulation can be viewed as a Markov process. Let $\mathcal{P}$ be a finite or countably infinite set of physical states and let $\mathcal{E}$ be a finite set corresponding to those events that can trigger physical state transitions. For each of the events $e \in \mathcal{E}$ that are active in a physical state $s \in \mathcal{P}$, we can conceptually imagine that there is an associated clock. When a clock e^* runs down to zero in state s, it triggers a physical state transition to s' with probability $p(s'; s, e^*)$. The clocks e' active in s' correspond to events that were already scheduled in the previous state s (but had not yet run down to zero), in which case the clocks continue running down to zero in s' at rate $r(s', e')$, or correspond to new events that must be scheduled in s'. The clocks associated with such new events in e' are independently scheduled according to distributions $F(\cdot; e', s', e^*, s)$, where e^* is the trigger event that initiated the transition. Experienced simulationists will recognize that the state of the clocks effectively describes the "future event schedule" of the associated discrete-event simulation.

Given that the physical state and future event schedule is precisely the information necessary to evolve a discrete-event simulation forward in time, it is clear that $X(t) = (S(t), C(t))$ is Markov, where $S(t)$ is the physical state occupied at time t (known as the GSMP corresponding to X), and $C(t)$ is the vector of clock readings corresponding to active events.

To develop a sufficient condition under which the steady-state simulation problem for the Markov process $X = (X(t): t \geqslant 0)$ is well posed, one clearly needs to invoke an assumption that ensures that there is a unique stationary distribution for X. This, of course, requires an irreducibility hypothesis of some kind.

Definition 2. The GSMP corresponding to X is said to be *irreducible* if for each pair $(s, s') \in \mathcal{P} \times \mathcal{P}$, there exists a finite sequence of states $s_1, \ldots, s_n$ and events $e_1, \ldots, e_n$ such that for $0 \leqslant i \leqslant n$, e_i is active in s_i ($s_0 \triangleq s$, $s_{n+1} = s'$) and

$$\prod_{i=0}^{n} p(s_{i+1}; s_i, e_i) r(s_i, e_i) > 0.$$

The above assumption induces irreducibility over the physical states of the GSMP. For irreducibility over the clock readings of X, consider the following definition.

Definition 3. The GSMP corresponding to X satisfies the *positive density assumption* if each distribution $F(\cdot; e', s', e, s)$ has a density component that is strictly bounded away from zero on an interval $[0, \varepsilon]$ with $\varepsilon > 0$.

We are now ready to state our main result for GSMPs.

Theorem 4. *Suppose that*:

(i) *The GSMP corresponding to X is irreducible and satisfies the positive density assumption*;
(ii) $|\mathcal{P}| < \infty$;
(iii) $\int_{[0,\infty)} tF(\mathrm{d}t; e', s', e, s) < \infty$ *for all* (e', s', e, s).

Then, the steady-state simulation problem for the Markov process X is well-posed.

For a proof, see Glynn and Haas (2006). The above conditions are necessary, in the sense that if any of the three conditions above is violated, then there exist counterexamples.

8 Algorithmic identification of regenerative structure

Our discussion of Sections 6 and 7 makes clear that regenerative structure exists within the typical discrete-event steady-sate simulation. On the other hand, the TAVC estimator of Section 3, as well as the low bias estimators of Section 5, all depend upon the ability of the simulationist to identify the associated regeneration times. Of course, this identification is trivial in the setting of discrete state space Markov chains, where the regeneration times can be chosen to be those times at which the chain enters a fixed state. Unfortunately, identification of the regenerative structure guaranteed by Theorem 3 is not algorithmically trivial in general.

The main difficulty is that the regenerative structure of Theorem 3 involves the use of randomized regeneration. This means that the regeneration times cannot be identified purely on the basis of simulating X alone. Some additional random variables (i.e., the "randomization") must also be generated. This means, for example, that a "randomization post-processor" must be added to the typical discrete-event simulation in order to be able to exploit the regenerative structure that is theoretically guaranteed to exist.

We now proceed to describe the randomized regeneration structure that arises in the setting of discrete time Markov chains $X = (X_n: n \geqslant 0)$. The regenerative structure that can be required in continuous time is generally more

complicated to describe; see Sigman (1990) for details. Fortunately, the discrete time theory covers discrete-event simulations. In that setting, the key discrete-time process is the sequence $(X(\Gamma_n)\colon n \geqslant 0)$, where $0 = \Gamma_0 < \Gamma_1 < \cdots$ are the epochs at which physical state transitions occur and $(X(t)\colon t \geqslant 0)$ is the continuous-time Markov process (associated with GSMPs) described earlier.

Let $(X_n\colon n \geqslant 0)$ be an S-valued discrete-time Markov chain satisfying the conditions of Theorem 3. Then, there exists a function $\lambda : S \to [0, 1]$, a subset $A \subseteq S$, an integer $m \geqslant 1$ and a probability distribution φ such that:

(i) $\mathrm{P}_x(X_m \in \cdot) \geqslant \lambda(x)\varphi(\cdot)$, $x \in S$;
(ii) $\mathrm{P}_x(X_n \in A \text{ infinitely often}) = 1$, $x \in S$;
(iii) $\inf\{\lambda(x)\colon x \in A\} > 0$.

To see how this gives rise to randomized regenerative structure, note that condition (i) guarantees that we can write

$$\mathrm{P}_x(X_m \in \cdot) = \lambda(x)\varphi(\cdot) + \big(1 - \lambda(x)\big)Q(x, \cdot), \tag{18}$$

where $Q(x, \cdot)$ is a probability distribution on S for each x. Hence, conditional on $X_n = x$, we can generate X_{n+m} by generating a Bernoulli r.v. having parameter $\lambda(X_n)$. If the Bernoulli r.v. takes on value 1, then we distribute X_{n+m} according to φ; otherwise, we distribute X_{n+m} according to $Q(X_n, \cdot)$. The segment $(X_{n+1}, \ldots, X_{n+m-1})$ is then generated from the conditional distribution, given the starting state X_n and ending state X_{n+m} for the full segment $(X_n, \ldots, X_{n+m})$. Whenever we distribute X_{n+m} according to φ, X_{n+m} is independent of the history of the chain up to and including step n. Conditions (ii) and (iii) guarantee that there exist infinitely many times $T(0) < T(0)+m \leqslant T(1) < T(1)+m \leqslant T(2) < \cdots$ (separated by gaps of at least m steps) at which the chain is distributed according to φ.

The random times $(T(n)\colon n \geqslant 0)$ form cycle boundaries that correspond to a regenerative process with one-dependent identically distributed cycles. In the special case that $m = 1$, the cycles are i.i.d. and the process X is regenerative in the classical sense. One difficulty with this means of identifying regenerative times is that the algorithm is invasive. By invasive, we mean that the algorithm impacts the way we generate sample replications of the process X. In particular, were we to straightforwardly adapt the above mechanism for constructing regeneration times, the basic algorithms used to simulate (for example) discrete-event systems would need to be modified.

In view of this, it is clearly desirable to develop an alternative implementation of the algorithm. Under the conditions above, it can be shown that there exists a function $w : S \times S \to [0, 1]$ for which

$$\lambda(x)\varphi(\mathrm{d}y) = w(x, y)\,\mathrm{P}_x(X_m \in \mathrm{d}y),$$
$$\big(1 - \lambda(x)\big)Q(x, \mathrm{d}y) = \big(1 - w(x, y)\big)\,\mathrm{P}_x(X_m \in \mathrm{d}y).$$

Suppose that one simulates a realization of the process X (using one's algorithm of choice). To identify the regeneration time, we apply the following

"post-processor" to the sampled realization. If $X_n = x$ and $X_{n+m} = y$, generate a Bernoulli r.v. having probability $w(X_n, X_{n+m})$. If the Bernoulli r.v. is 1, then the process has distribution φ at time $n + m$; otherwise, it does not. This algorithm is probabilistically equivalent to the method described earlier, but it is noninvasive and requires only the post-processing step.

This idea can be applied even to discrete state space Markov chains in which consecutive entry times to a fixed state z constitute easily identified regenerations. To illustrate, suppose that $(P(x, y)\colon x, y \in S)$ is the one-step transition matrix of a discrete-time Markov chain to be simulated. Put $\varphi(y) = P(z, y)$ and

$$\lambda(x) = \inf_{y \in S} \frac{P(x, y)}{P(z, y)}$$

so that $P(x, y) \geqslant \lambda(x)P(z, y)$ for $x, y \in S$. Clearly, $\lambda(z) = 1$. However, $\lambda(x)$ typically will be positive for some states $x \neq z$. Hence, our randomized (classical) regenerations occur every time the chain leaves z, but generally occur more frequently. For this special case in which one sequence of regenerations is a superset of another sequence of regenerations, one can prove that the superset provides a statistically more efficient estimator of the TAVC (i.e., the D_{22} term in Theorem 2 is reduced); see Andradóttir et al. (1995) for details.

Returning to the discrete-event context, the key to practical implementation is computing the function w. This, in turn, involves computing $\mathrm{P}_x(X_m \in \cdot)$. This can easily be done for $m = 1$, but is nontrivial for values of $m \geqslant 2$. Henderson and Glynn (2001) argue that for a discrete-event simulation in which no "event cancellation" is possible, the minimal value of m compatible with conditions (i)–(iii) is m^*, where m^* is the minimal number of events that are ever simultaneously scheduled by the simulation. Hence, any discrete-event simulation possessing a so-called "single state" (i.e., a state in which only one event is scheduled) is easily amenable to regenerative methodology.

When $m^* > 1$, algorithmic identification of regeneration times is substantially more complicated. Suppose, however, that there is no event cancellation in the GSMP and that each "speed" $r(s, e)$ equals one. Assume the GSMP is initially in a state $x = (s, c)$ in which m clocks are active and in which event $\tilde{e} = \tilde{e}(x)$ is the one that has the most remaining time until its clock runs down to zero. Then, the m-step transition probability for $(X(\Gamma_i)\colon i \geqslant 0)$ can be easily computed on that part of the sample space on which the (random) trigger event e_m^* corresponding to transition m equals $\tilde{e}$. In other words, $\mathrm{P}_x(X(\Gamma_m) \in \cdot, e_m^* = \tilde{e}(x))$ can be written down explicitly, because the event $\{e_m^* = \tilde{e}(X_0)\}$ is precisely the event that all of the clocks scheduled prior to time m were set at values greater than that associated with $\tilde{e}(X_0)$. Thus, if we find a function λ and a probability φ for which

$$\mathrm{P}_x\big(X_m \in \cdot, e_m^* = \tilde{e}(X_0)\big) \geqslant \lambda(x)\varphi(\cdot),$$

we can implement the same approach as desired above, by letting w be chosen so that

$$w(x, y)\, \mathrm{P}_x\big(X_m \in \mathrm{d}y,\, e^*_m = \tilde{e}(X_0)\big) = \lambda(x)\varphi(\mathrm{d}y).$$

The only significant change is that we generate the Bernoulli r.v. having parameter $w(X_n, X_{n+m})$ only when $e^*_{n+m} = \tilde{e}(X_n)$; see Henderson and Glynn (1999) for a related idea.

This method (and analogous ideas for dealing with nonunit speeds and event cancellation) can be used to construct regenerations for a large class of GSMPs. Practically speaking, however, the method is difficult to apply for GSMPs for which m^* is large. The difficulty is that condition (i) becomes more demanding when m^* is large, because one must lower bound a joint density on a "clock space" having dimension at least m^*. Such lower bounds typically degenerate badly in high-dimensional contexts, leading to implementations in which regeneration occurs (very) infrequently.

Thus, the question of developing easily implemented and practically useful methods for algorithmic identification of regeneration times remains largely open at this point.

9 A martingale perspective on regeneration

To illustrate the connection between martingales and regeneration, we focus here on the case in which $V(t) = g(X(t))$, where $X = (X(t)\colon\ t \geqslant 0)$ is an irreducible finite state continuous time Markov chain with rate matrix $Q = (Q(x, y)\colon\ x, y \in S)$. Given the performance measure $g : S \to \mathbb{R}$ (where we choose to encode $g = (g(x)\colon\ x \in S)$ as a column vector), the linear system

$$Qh = -(g - \alpha e) \tag{19}$$

has a solution h. Here, e is the column vector in which all entries equal 1, and (19) is called Poisson's equation.

It is a standard fact in the theory of Markov processes that

$$h\big(X(t)\big) - \int_0^t (Qh)\big(X(s)\big)\,\mathrm{d}s$$

is then a martingale; see, for example, Karlin and Taylor (1981). In other words,

$$h\big(X(t)\big) + \int_0^t g\big(X(s)\big)\,\mathrm{d}s - \alpha t$$

enjoys the martingale property. Suppose that T is a stopping time adapted to X (so that $I(T \leqslant t)$ is a deterministic function of the path $(X(s)\colon\ 0 \leqslant s \leqslant t)$) having finite expectation. Then, the optional sampling theorem for martingales (Breiman, 1968) can be applied, yielding

$$\mathrm{E}h\big(X(T)\big) + \mathrm{E}\int_0^T g\big(X(s)\big)\,\mathrm{d}s - \alpha\,\mathrm{E}T = \mathrm{E}h\big(X(0)\big). \tag{20}$$

The identity (20) holds in great generality for Markov processes in discrete and continuous time, provided that one suitably generalizes the rate matrix Q in an appropriate way; see Ethier and Kurtz (1986).

Note that if the stopping time T can be chosen so that $\mathrm{E}h(X(T)) = \mathrm{E}h(X(0))$, then α can be viewed as the ratio of expectations

$$\alpha = \frac{\mathrm{E}\int_0^T g(X(s))\,\mathrm{d}s}{\mathrm{E}T} \triangleq \frac{\mathrm{E}Y}{\mathrm{E}\tau}. \tag{21}$$

Hence, by simulating independent copies of the process X over the time interval $[0, T]$, α can be computed via a ratio estimator that enjoys precisely the same central limit and bias properties as the conventional regenerative estimator described earlier in this chapter. In particular, if $(Y_1, \tau_1), \ldots, (Y_n, \tau_n)$ are n i.i.d. copies of (Y, τ), then $\alpha_n = (Y_1 + \cdots + Y_n)/(\tau_1 + \cdots + \tau_n)$ satisfies the CLT

$$n^{1/2}(\alpha_n - \alpha) \Rightarrow \sigma \mathrm{N}(0, 1)$$

as $n \to \infty$, where $\sigma^2 = \mathrm{var}(Y - \alpha\tau)/(\mathrm{E}\tau)^2$, and the bias expansion

$$\mathrm{E}\alpha_n = \alpha - \frac{1}{n}\frac{\mathrm{E}(Y - \alpha\tau)\tau}{(\mathrm{E}\tau)^2} + \mathrm{o}\left(\frac{1}{n}\right)$$

holds (thereby providing a straightforward construction of a "low bias" estimator having bias $\mathrm{o}(1/n)$).

Of course, the key is to find a distribution for $X(0)$ and a random time T so that $\mathrm{E}h(X(T)) = \mathrm{E}h(X(0))$.

Since the simulationist does not know the solution to Poisson's equation, one simply chooses $X(0)$ and T so that $X(0)$ and $X(T)$ have the same distribution. Of course, the easiest way to guarantee this is to let T be the first time that X returns to the state occupied at time $t = 0$. In this case, the above estimation procedure just reduces to the conventional regenerative method based on successive returns to a single state. Thus, the martingale perspective offers a complementary viewpoint regarding the regenerative method for computing steady-state expectations.

However, the optional sampling formula (20) offers the potential for developing new steady-state estimation algorithms. As an example, consider the successive times $\beta_1, \beta_2, \ldots$ at which X enters some fixed nonempty subset $A \subset S$. To be precise, let $\beta_0 = 0$ and put

$$\beta_i = \inf\bigl\{t > \beta_{i-1}\colon X(t) \in A, X(t-) \notin A\bigr\}.$$

Then, $(X(\beta_i)\colon i \geqslant 1)$ is an irreducible A-valued discrete time Markov chain with stationary distribution ν_A. If $\mathrm{E}_{\nu_A}(\cdot)$ is the expectation operator under which $X(0)$ has distribution ν_A, then (20) yields the equality

$$\mathrm{E}_{\nu_A} h\bigl(X(T)\bigr) + \mathrm{E}_{\nu_A}\int_0^{\beta_1} g\bigl(X(s)\bigr)\,\mathrm{d}s - \alpha\,\mathrm{E}_{\nu_A}\beta_1 = \mathrm{E}_{\nu_A} h\bigl(X(0)\bigr).$$

Because $\mathrm{E}_{\nu_A} h(X(T)) = \mathrm{E}_{\nu_A} h(X(0))$, we obtain the identity

$$\alpha = \frac{\mathrm{E}_{\nu_A} \int_0^{\beta_1} g(X(s))\,\mathrm{d}s}{\mathrm{E}_{\nu_A} \beta_1}. \tag{22}$$

Hence, the regenerative ratio formula (7) generalizes beyond the conventional setting in which cycles are defined in terms of i.i.d. cycles. (Note that the β_i's split the sample path for X into identically distributed cycles having a complex dependency structure.)

One might hope to algorithmically exploit (22) in the same way as for (21). If we could generate variates from ν_A, the algorithm would be clear. Just generate $X(0)$ from ν_A, simulate to time β_1, and compute (Y, τ), where

$$\begin{aligned} Y &= \int_0^{\beta_1} g\big(X(s)\big)\,\mathrm{d}s, \\ \tau &= \beta_1. \end{aligned}$$

By simulating i.i.d. copies of (Y, τ), we can estimate α via $\alpha_n = (Y_1 + \cdots + Y_n)/(\tau_1 + \cdots + \tau_n)$. Precisely the same CLT and low bias estimation procedure as before can thus be utilized. Because a "multi-state" A is hit more frequently than any single state, we can expect this estimation procedure to be more efficient than the conventional regenerative method based on returns to a single fixed state.

The difficulty, of course, is that we typically are unable to generate variates from ν_A. However, by conditioning on $X(0)$ and $X(\beta_1)$ in (22), we can rewrite (22) as

$$\alpha = \frac{\sum_{x,y} \mathrm{E}_x Y(y) \nu_A(x)}{\sum_{x,y} \mathrm{E}_x \tau(y) \nu_A(x)} \triangleq \frac{\sum_{x,y} \mathrm{E}_x u(x, y)}{\sum_{x,y} \ell(x, y)}, \tag{23}$$

where

$$\begin{aligned} Y(y) &= \int_0^{\beta_1} g\big(X(s)\big)\,\mathrm{d}s\, I\big(X(\beta_1) = y\big), \\ \tau(y) &= \beta_1 I\big(X(\beta_1) = y\big). \end{aligned}$$

Each term in the numerator and denominator can be estimated via simulation of X over $[0, t]$, namely

$$\hat{u}(t, x, y) = \frac{1}{J(t)} \sum_{i=1}^{J(t)} \int_{\beta_{i-1}}^{\beta_i} g\big(X(s)\big)\,\mathrm{d}s\, I\big(X(\beta_{i-1}) = x, X(\beta_i) = y\big),$$

$$\hat{\ell}(t, x, y) = \frac{1}{J(t)} \sum_{i=1}^{J(t)} (\beta_i - \beta_{i-1}) I\big(X(\beta_{i-1}) = x, X(\beta_i) = y\big),$$

where $J(t)$ is the number of times the process X enters A over $[0, t]$. The representation (23) leads to a point estimator

$$\frac{\sum_{x,y} \hat{u}(t, x, y)}{\sum_{x,y} \hat{\ell}(t, x, y)}$$

for α. A corresponding "plug-in" TAVC estimator can be implemented in a straightforward fashion when $|A| < \infty$. The plug-in estimator takes advantage of the fact that the TAVC estimator for the A-valued discrete-time Markov chain $(X(\beta_i)\colon\ i \geqslant 1)$ can be computed easily (by solving a linear system of $|A|$ equations in $|A|$ unknowns). This new estimator for the TAVC associated with simulation of X over $[0, t]$ is essentially the semi-regenerative TAVC estimator proposed by Calvin et al. (2006).

10 Efficiency improvement via regeneration: Computing steady state gradients

In many applications settings, it is of interest to compute the sensitivity of the system's performance to perturbations in an underlying parameter. For example, it may be that the arrival rate to a queue is only approximately known, so that computing the change in performance that corresponds to changing the arrival rate is relevant. In particular, computing the derivative (or, more generally, the gradient) of a steady-state performance measure with respect to the arrival rate is a computational problem of significant importance. Of course, such derivatives also play a key role in both simulation-based stochastic optimization and statistical analysis of complex stochastic systems; see Glynn (1990).

To be more precise, suppose that the probability distribution underlying the simulation of V depends on a parameter $\theta \in \mathbb{R}^d$. Let P_θ be the distribution corresponding to θ. Then, the steady-state mean of V depends on θ, so that $\alpha = \alpha(\theta)$. As noted earlier, the time average $\overline{V}(t)$ satisfies

$$\overline{V}(t) \Rightarrow \alpha(\theta)$$

under P_θ, suggesting that

$$\mathrm{E}_\theta \overline{V}(t) \to \alpha(\theta). \tag{24}$$

In significant generality, there exists a random process $L(\theta, t)$ (typically, a martingale), known as the "likelihood ratio process", such that

$$\mathrm{E}_\theta \overline{V}(t) = \mathrm{E}_{\theta_0} \overline{V}(t) L(\theta, t) \tag{25}$$

for $t \geqslant 0$, where $\mathrm{E}_{\theta_0}(\cdot)$ is the expectation operator corresponding to P_{θ_0}. Assuming that the gradient can be interchanged with the expectation operator $\mathrm{E}_{\theta_0}(\cdot)$, we find that

$$\nabla \mathrm{E}_\theta \overline{V}(t)|_{\theta=\theta_0} = \mathrm{E}_{\theta_0} \overline{V}(t) \nabla L(\theta, t)|_{\theta=\theta_0}.$$

In view of (24), it seems reasonable to expect that

$$\nabla \mathrm{E}_\theta \overline{V}(t)|_{\theta=\theta_0} \to \nabla\alpha(\theta_0)$$

as $t \to \infty$, so that

$$\mathrm{E}_{\theta_0} \overline{V}(t) \nabla L(\theta_0, t) \to \nabla\alpha(\theta_0)$$

as $t \to \infty$. In particular, assume that a bias expansion similar to that derived in Section 5 holds, so that

$$\mathrm{E}_{\theta_0} \overline{V}(t) \nabla L(\theta_0, t) = \nabla\alpha(\theta_0) + \frac{1}{t}\nabla b(\theta_0) + \mathrm{o}\left(\frac{1}{t}\right) \tag{26}$$

as $t \to \infty$. On the other hand, $(L(\theta, t)\colon\ t \geqslant 0)$ is typically a martingale process for each θ. Given that $(h^{-1}(L(\theta_0 + he_i, t) - L(\theta_0, t))\colon\ t \geqslant 0)$ is then a martingale for each unit vector e_i, one expects $(\nabla L(\theta_0, t)e_i\colon\ t \geqslant 0)$ to be a martingale as well. The martingale CLT (see, for example, p. 476 of Billingsley, 1995) then suggests that there exists a (deterministic) constant $\gamma \in (0, \infty)$ such that

$$t^{-1/2} \nabla L(\theta_0, t) e_i \Rightarrow \gamma\, \mathrm{N}(0, 1)$$

as $t \to \infty$. Slutsky's lemma then implies the weak convergence statement

$$t^{-1/2} \overline{V}(t) \nabla L(\theta_0, t) e_i \Rightarrow \alpha\gamma\, \mathrm{N}(0, 1)$$

as $t \to \infty$, so that we expect

$$t^{-1} \operatorname{var} \overline{V}(t) \nabla L(\theta_0, t) e_i \to \alpha^2\gamma^2 \operatorname{var} \mathrm{N}(0, 1) \tag{27}$$

as $t \to \infty$.

For a given computer budget c, how many independent replications m of length c/m should one simulate to minimize the mean square error of the resulting estimator? Note that the bias of each replication is then of order m/c (in view of (26)). The variance of a replication of length c/m is of order c/m (see (27)). So, the sample mean over m such replications has a variance of order c/m^2. The value of m that minimizes the mean square error of the corresponding gradient estimator is then of order $c^{3/4}$, yielding an estimator with a root mean square error of order $c^{-1/4}$.

However, if regenerative structure can be identified algorithmically, then a different means of estimating $\nabla\alpha(\theta_0)$ is available to simulationist. If the process V is regenerative under P_θ, the ratio formula

$$\alpha(\theta) = \frac{\mathrm{E}_\theta \int_0^{\tau_1} V(s)\, \mathrm{d}s}{\mathrm{E}_\theta \tau_1} \tag{28}$$

holds (provided that the process is initiated with a regeneration at $t = 0$, so that $T(0) = 0$). Assuming the existence of a likelihood ratio process $(L(\theta, t)\colon\ t \geqslant 0)$, we expect to be able to extend the identity (25) separately

to the numerator and denominator of (28) to

$$\alpha(\theta) = \frac{\mathrm{E}_{\theta_0} \int_0^{\tau_1} V(s)\, \mathrm{d}s\, L(\theta, \tau_1)}{\mathrm{E}_{\theta_0} \tau_1 L(\theta, \tau_1)}. \tag{29}$$

Assuming that the gradient operator can be interchanged with the expectation operator $\mathrm{E}_{\theta_0}(\cdot)$ (as above), we find that

$$\nabla\alpha(\theta_0) = \frac{\mathrm{E}_{\theta_0}\left[\int_0^{\tau_1} (V(s) - \alpha(\theta_0))\, \mathrm{d}s\, \nabla L(\theta, \tau_1)\right]}{\mathrm{E}_{\theta_0} \tau_1}. \tag{30}$$

Consequently, $\nabla\alpha(\theta_0)$ can be expressed as $\nabla\alpha(\theta_0) = k(\mathrm{E}\xi)$, where

$$\xi = \left(\int_0^{\tau_1} V(s)\, \mathrm{d}s,\, \tau_1,\, \int_0^{\tau_1} V(s)\, \mathrm{d}s\, \nabla L(\theta_0, \tau_1),\, \tau_1 \nabla L(\theta_0, \tau_1)\right)$$

and

$$k(x_1, x_2, x_3, x_4) = \frac{x_3 - (x_1/x_2)x_4}{x_2}.$$

Suppose that one simulates n i.i.d. cycles of v under the probability P_{θ_0}, thereby obtaining n independent copies $\xi_1, \xi_2, \ldots, \xi_n$ of the random vector ξ. The estimator $\nabla\alpha_n(\theta_0) \triangleq k(n^{-1} \sum_{i=1}^n \xi_i)$ can then be analyzed via "delta-method" type arguments (see, for example, Serfling, 1980) to establish that the estimator converges at rate $n^{-1/2}$ (and hence, in units of computer time, at rate $c^{-1/2}$) to the gradient $\nabla\alpha(\theta_0)$. This $c^{-1/2}$ convergence rate is to be contrasted with the $c^{-1/4}$ rate observed earlier, and shows clearly that regenerative structure can be usefully exploited in obtaining substantial efficiency improvements.

11 Efficiency improvement via regeneration: Computing infinite horizon discounted reward

We now offer a second illustration of the principle that the presence of regenerative structure can be usefully exploited to obtain efficiency improvements. Consider the infinite horizon expected discounted reward $\alpha = \mathrm{E}D$, where

$$D = \int_0^\infty \mathrm{e}^{-rt} V(t)\, \mathrm{d}t$$

for some $r > 0$. From a computational standpoint, an algorithm based on simulating i.i.d. copies of the r.v. D cannot be operationalized, because it takes infinite time to generate the above r.v. As a consequence, one needs to consider computationally feasible alternatives.

One such approach is to exploit regeneration. In particular, suppose that V is regenerative with regeneration times $0 = T(0) < T(1) < \cdots$ (so that

V regenerates at $T = 0$) that split the sample path of V into i.i.d. cycles. Then,

$$\begin{aligned}\alpha &= \mathrm{E}\int_0^\infty \mathrm{e}^{-rt}V(t)\,\mathrm{d}t \\ &= \mathrm{E}\int_0^{\tau_1} \mathrm{e}^{-rt}V(t)\,\mathrm{d}t + \mathrm{E}\mathrm{e}^{-r\tau_1}\int_0^\infty \mathrm{e}^{-rt}V(\tau_1+t)\,\mathrm{d}t \\ &= \mathrm{E}\int_0^{\tau_1} \mathrm{e}^{-rt}V(t)\,\mathrm{d}t + \mathrm{E}\mathrm{e}^{-r\tau_1}\alpha\end{aligned}$$

so that $\alpha = k(\mathrm{E}\xi)$, where

$$\xi = \left(\int_0^{\tau_1} \mathrm{e}^{-rt}V(t)\,\mathrm{d}t, \mathrm{e}^{-r\tau_1}\right)$$

and $k(x_1, x_2) = x_1/(1 - x_2)$.

As in Section 10, the corresponding regenerative estimator for α is $\alpha_n = k(n^{-1}\sum_{i=1}^n \xi_i)$, where $\xi_1, \xi_2, \ldots, \xi_n$ are i.i.d. copies of ξ obtained by simulating n independent cycles of V. Note that the estimator for α can be computed in finite time and is computationally feasible. Furthermore, the delta method again applies, yielding the conclusion that α_n typically converges to α at rate $c^{-1/2}$ in the computer budget c. Thus, use of regeneration in this setting makes feasible and practical a computation that is problematic in its original form.

References

Anderson, T.W. (1971). *The Statistical Analysis of Time Series*. Wiley, New York.

Andradóttir, S., Calvin, J.M., Glynn, P.W. (1995). Accelerated regeneration for Markov chain simulations. *Probability in the Engineering and Informational Sciences* 9, 497–523.

Billingsley, P. (1995). *Probability and Measure*, 3rd edition. Wiley, New York.

Breiman, L. (1968). *Probability*. *Addison–Wesley Series in Statistics*. Addison–Wesley, Reading, MA.

Calvin, J. (1994). Return state independent quantities in regenerative simulation. *Operations Research* 42, 531–542.

Calvin, J.M., Glynn, P.W., Nakayama, M.K. (2006). The semi-regenerative method of simulation output analysis. *ACM Transactions on Modeling and Computer Simulation*, in press.

Cox, D., Smith, W. (1961). *Queues*. Methuen, London.

Cramér, H. (1946). *Mathematical Methods of Statistics*. Princeton University Press, Princeton, NJ.

Crane, M., Iglehart, D. (1974a). Simulating stable stochastic systems, I: General multi-server queues. *Journal of the Association for Computing Machinery* 21, 103–113.

Crane, M., Iglehart, D. (1974b). Simulating stable stochastic systems, II: Markov chains. *Journal of the Association for Computing Machinery* 21, 114–123.

Crane, M., Iglehart, D. (1975). Simulating stable stochastic systems, III: Regenerative processes and discrete-event simulations. *Operations Research* 23, 33–45.

Efron, B., Tibshirani, R.J. (1993). *An Introduction to the Bootstrap*. Chapman and Hall, New York.

Ethier, S.N., Kurtz, T.G. (1986). *Markov Processes: Characterization and Convergence*. *Wiley Series in Probability and Statistics*. Wiley, New York.

Fishman, G. (1973). Statistical analysis for queueing simulations. *Management Science* 20, 363–369.

Fishman, G. (1974). Estimation in multiserver queueing simulations. *Operations Research* 22, 72–78.

Glynn, P.W. (1990). Likelihood ratio gradient estimation for stochastic systems. *Communications of the ACM* 33, 75–84.
Glynn, P.W. (1994). Some topics in regenerative steady-state simulation. *Acta Applicandae Mathematicae* 34 (1/2), 225–236.
Glynn, P.W., Haas, P.J. (2006). A law of large numbers and functional central limit theorem for generalized semi-Markov processes. *Stochastic Models*, in press.
Glynn, P.W., Heidelberger, P. (1990). Bias properties of budget constrained simulation. *Operations Research* 38 (5), 801–814.
Glynn, P.W., Heidelberger, P. (1992). Jackknifing under a budget constraint. *ORSA Journal on Computing* 4 (3), 226–234.
Glynn, P.W., Iglehart, D.L. (1986). Recursive moment formulas for the regenerative method of simulation. In: Janssen, J. (Ed.), *Semi-Markov Models*. Plenum Press, New York.
Glynn, P.W., Iglehart, D.L. (1987). A joint central limit theorem for the sample mean and regenerative variance estimator. *Annals of Operations Research* 8, 41–55.
Glynn, P.W., Iglehart, D.L. (1993). Conditions for the applicability of the regenerative method. *Management Science* 39, 1108–1111.
Glynn, P.W., Whitt, W. (1993). Limit theorems for cumulative processes. *Stochastic Processes and Their Applications* 47, 299–314.
Glynn, P.W., Whitt, W. (2002). Necessary conditions in limit theorems for cumulative processes. *Stochastic Processes and Their Applications* 98 (2), 199–209.
Henderson, S.G., Glynn, P.W. (1999). Can the regenerative method be applied to discrete-event simulation? In: Farrington, P.A., Nembhard, H.B., Evans, J., Sturrock, D. (Eds.), *Proceedings of the 1999 Winter Simulation Conference*. IEEE Press, Piscataway, NJ, pp. 367–373.
Henderson, S.G., Glynn, P.W. (2001). Regenerative steady-state simulation of discrete-event systems. *ACM Transactions on Modeling and Computer Simulation* 11 (4), 313–345.
Ibragimov, I.A., Has'minskii, R.Z. (1981). *Statistical Estimation – Asymptotic Theory*. *Applications of Mathematics*. Springer-Verlag, New York.
Kabak, I. (1968). Stopping rules for queueing simulations. *Operations Research* 16, 431–437.
Karlin, S., Taylor, H.M. (1981). *A Second Course in Stochastic Processes*. Academic Press, New York.
Lewis, P.A.W., Orav, E.J. (1989). *Simulation Methodology for Statisticians, Operations Analysts, and Engineers*. *Wadsworth & Brooks/Cole Statistics/Probability Series*. Wadsworth & Brooks/Cole, Pacific Grove, CA.
Meketon, M.S., Heidelberger, P. (1982). A renewal theoretic approach to bias reduction in regenerative simulations. *Management Science* 28 (2), 173–181.
Meyn, S., Tweedie, R. (1993). Criteria for stability of Markovian processes III: Foster–Lyapunov criteria for continuous time processes, with examples. *Advances in Applied Probability* 25, 518–548.
Miller, R.G. (1974). The jackknife – A review. *Biometrika* 61 (1), 1–15.
Serfling, R.J. (1980). *Approximation Theorems of Mathematical Statistics*. *Wiley Series in Probability and Statistics*. Wiley, New York.
Sigman, K. (1990). One-dependent regenerative processes and queues in continuous time. *Mathematics of Operations Research* 15, 175–189.