SOME TOPICS IN REGENERATIVE STEADY STATE SIMULATION

PETER W. GLYNN

Stanford University
Stanford, CA 94305, USA

Revised May 1993

ABSTRACT

This paper offers a short introduction to the regenerative method of steady-state simulation output analysis. The paper also contains several new results. In particular, it is shown that regenerative methods necessarily apply to steady-state simulations that are "well-posed" in a certain precise sense. The paper also describes a bias-reduction algorithm that takes advantage of regenerative structure.

KEYWORDS

Steady-state simulation, regenerative processes, initial transient, Harris chains.

1. INTRODUCTION

In many performance settings, it is of great interest to be able to numerically calculate steady-state performance measures. Because of its great flexibility, simulation offers a powerful means of computing such quantities. In this paper, we hope to argue that the presence of regenerative structure in the stochastic system under consideration greatly enhances the ability of the simulator to apply output analysis and variance reduction techniques that offer substantial advantages over their non-regenerative analogs. We further take the point of view that while regenerative methods may not be universally applicable, developing a firm understanding of how to perform efficient simulation in the regenerative setting is a valuable first step in constructing more broadly-based methodologies.

This paper is organized as follows. In section 2, we show (in a precise sense) that in order that the steady-state simulation problem for a Markov process be well-posed, it is necessary that the process be Harris recurrent. We then point out that the steady-state of any such Harris recurrent process can be estimated by simulating an associated classically regenerative process. Section 3 describes the basic regenerative method for steady-state simulation output analysis, as well as a related bias-reducing estimation algorithm that also takes advantage of regenerative structure.

2. WELL POSED STEADY-STATE SIMULATION PROBLEMS

Let $X = (X(t) : t \geq 0)$ be a (time-homogeneous) Markov process in continuous time, taking values in a state space S. For a given real-valued function $f : S \to \Re$, the simulator is often interested in the long-run behavior of the sample time-average defined by

$$\alpha(t; f) \stackrel{\triangle}{=} \frac{1}{t} \int_0^t f(X(s))ds$$

We will now introduce the concept of a well-posed steady-state simulation problem. Let $P_x(\cdot)$ be the probability distribution of X on path-space under which $X(0) = x$, and let $E_x(\cdot)$ be its corresponding expectation operator.

DEFINITION 1. *The steady-state simulation problem is said to be well-posed for X if for every bounded real-valued measurable function f, there exists a number $\alpha(f)$ such that for each $x \in S$,*

$$E_x \alpha(t; f) \to \alpha(f)$$

as $t \to \infty$.

We note that the concept of well-posedness demands that the limiting behavior of the time-average $\alpha(t; f)$ must, at least in expectation, be asymptotically independent of the initial condition. Since one would hope that the simulator's chosen initial condition plays no role in determining the steady-state behavior of the system, this requirement for well-posedness seems quite reasonable.

There is an intimate relationship between the concept of well-posedness and that of Harris recurrence in continuous time.

DEFINITION 2. *We say that X is Harris recurrent (in continuous time) if there exists a non-trivial σ-finite measure η such that whenever $\eta(A) > 0$, then for each $x \in S$,*

$$\int_0^\infty I(X(t) \in A)dt = +\infty$$

P_x a.s. If, in addition, X has an invariant probability measure π, then X is said to be a positive recurrent Harris process.

THEOREM 1. *The steady-state simulation problem is well-posed for X if and only if X is a positive recurrent Harris process.*

PROOF. Suppose first that the steady-state simulation problem is well-posed for X. It is evident (by specializing to indicator functions f) that for each measurable subset A,

$$\frac{1}{t}\int_0^t P_x[X(s) \in A]ds \to \alpha(I_A) \tag{2.1}$$

as $t \to +\infty$. The Vitali-Hahn-Saks theorem (see Renyi (1970)) then implies that the set function $\alpha(I_A)$ is a probability measure, say $\pi(\cdot)$. Since

$$\begin{aligned}
&\frac{1}{t}\int_h^{t+h} P_x[X(s) \in A]ds \\
&= \int_S \frac{1}{t}\int_0^t P_x[X(s) \in dy]P_y[X(h) \in A]ds \\
&\to \int_S \pi(dy)P_y[X(h) \in A]
\end{aligned}$$

as $t \to \infty$, it follows from (2.1) that π is an invariant probability measure. Furthermore, it is clear that since (2.1) holds for every x, π must in fact be the unique invariant probability associated with X, and

$$\frac{1}{t}\int_0^t E_x f(X(s))ds \to \pi f \stackrel{\triangle}{=} \int_S f(y)\pi(dy) \tag{2.2}$$

for every bounded measurable f and $x \in S$.

If Z is a bounded invariant r.v., then it is evident that $E_x[Z|X(s) : 0 \le s \le t] = g(X(t))$ a.s., where $g(x) = E_x Z$. Hence, $g(x) = E_x g(X(t))$ and (2.2) implies that

$$g(x) = \frac{1}{t}\int_0^t E_x g(X(s))ds \to \pi g$$

for $x \in S$. So, $g(X(t)) = \pi g$. Furthermore, martingale theory establishes that

$$E_x[Z|X(s) : 0 \le s \le t] \to E_x[Z|X(s) : s \ge 0] = Z$$

a.s. as $t \to \infty$, so that $Z = \pi g$ a.s. It follows that the invariant σ-field of X is trivial.

Now, for any $t > 0$, the shift θ_t is a measure-preserving transformation under P_π, where $P_\pi(\cdot)$ is the distribution on path-space under which $X(0)$ has distribution π. The ergodic theorem therefore yields

$$\alpha(t; f) \to \pi f \qquad P_\pi \quad \text{a.s.}$$

as $t \to \infty$, for bounded measurable f. If $Z = I(\alpha(t; f) \to \pi f)$, it is evident that Z is invariant and $E_x Z = 1$ for π a.e. $x \in S$. Our above argument therefore shows that $g(x) = E_x Z$ is a constant, and hence $g(x) = 1$ for all $x \in S$. In other words, $\alpha(t; f) \to \pi f \qquad P_x$ a.s. for every $x \in S$. By specializing f to indicator functions, we conclude that X is a positive recurrent Harris process.

The converse is straightforward, and follows from the law of large numbers for positive recurrent Harris processes (see, for example, Revuz(1984)).

REMARK 1. Glynn (1982) establishes the analog of Theorem 1 for discrete-time Markov chains. In particular, it is shown there that the steady-state simulation problem is well-posed for a discrete-time Markov chain $X = (X_n : n \geq 0)$ if and only if X is a positive recurrent Harris chain.

REMARK 2. Assume that S is a Polish space. One could reasonably weaken the concept of well-posedness by demanding only that

$$E_x \alpha(t; f) \to \alpha(f)$$

as $t \to \infty$ for each $x \in S$ and each bounded continuous function $f : S \to \Re$. With this weakened assumption on X, X need no longer be Harris recurrent. To illustrate, we consider the discrete-time example $X_{n+1} = \frac{1}{2} X_n + (\chi_{n+1} - 1)$, where the χ_i's are i.i.d. Bernoulli $(\frac{1}{2})$ r.v. 's; similar examples may be constructed in continuous time. It is easily shown that $X_n \Rightarrow U$, where U is uniform on $[-1, 1]$, regardless of X_0. On the other hand, if $f(x)$ is 1 or 0 depending on whether or not $x \in \mathcal{Q}$, it is evident that $E_x \alpha(t; f) \to 1$ if $x \in \mathcal{Q}$ and $E_x \alpha(t; f) \to 0$ otherwise. Consequently, the steady-state simulation problem is not well-posed for X, and X is therefore not a positive recurrent Harris chain (see Remark 1). In fact, X can not be a null recurrent Harris chain either, since it is easily shown that the uniform distribution on $[-1, 1]$ is in fact an invariant probability measure for X.

Positive recurrent Harris processes (PRHP's) enjoy a weakened form of regenerative structure. For a given non-decreasing sequence of random times $T(0) < T(1) < \ldots$, let

$$W_n(t) = \begin{cases} X(T(n-1) + t) & ; 0 \leq t < \tau_n \\ \triangle & ; t \geq \tau_n \end{cases}$$

where $\tau_n = T(n) - T(n-1)$ and $\triangle \notin S$, and $T(-1)$ is (by convention) set equal to zero.

THEOREM 2. *If X is a PRHP, there exists a sequence of random times $(T(n) : n \geq 0)$ such that:*

i) $P_x[(W_1, W_2, \ldots) \in \cdot] = P[(W_1, W_2, \ldots) \in \cdot]$ *is independent of* x*;*

ii) $(W_n : n \geq 0)$ *is a sequence of 1-dependent random elements under* P_x*,* $x \in S$*;*

iii) $(W_n : n \geq 1)$ *is a stationary sequence of random elements under* P_x*,* $x \in S$*.*

Furthermore, if π *is the invariant probability measure of* X*, then for each non-negative function* f*,*

$$\pi f = \lambda E Y_1$$

where $\lambda^{-1} = E\tau_1$ *and* $Y_1 = \int_0^{\tau_1} f(X(T(0) + s))ds$*.*

For a proof, see Azema, Duflo, and Revuz (1967) and Sigman (1990).

The most obvious approach to estimating the steady-state performance measure $\alpha \stackrel{\triangle}{=} \pi f$ is to simulate the process X over a long time horizon t, and to then form the estimator

$$\alpha(t) \stackrel{\triangle}{=} \alpha(t; f).$$

The following central limit theorem (CLT) for $\alpha(t)$ describes its asymptotic error as an estimator for α. Set $Z_i = Y_i - \alpha\tau_i$, and put $N(t) = \max\{n \geq -1 : T(n) \leq t\}$. Based on the fact that

$$\begin{aligned} \alpha(t) - \alpha &= \frac{1}{t} \sum_{i=1}^{N(t)} Z_i \\ &\approx \frac{1}{t} \sum_{i=1}^{\lfloor \lambda t \rfloor} Z_i, \end{aligned} \tag{2.3}$$

a CLT for $\alpha(t)$ can easily be derived by appealing to standard results for l-dependent sequences.

PROPOSITION 1. *Suppose that* f *is non-negative and* $E[Y_1^2 + \tau_1^2] < \infty$*. Then,*

$$t^{\frac{1}{2}}(\alpha(t) - \alpha) \Rightarrow \gamma N(0, 1)$$

as $t \to \infty$*, where* $\gamma^2 = \lambda(EZ_1^2 + 2EZ_1Z_2)$ *and* $\Rightarrow$ *denotes weak convergence.*

This follows easily from the l-dependence of the W_i's ; see Glynn(1982) and Sigman (1990) for details.

We now describe an alternative approach to estimating α. Rather than simulate X itself, we simulate a process X^* that is formed by "piecing together" i.i.d. replications of the "cycle" W_1. In other words, let $W_1^*, W_2^*, \ldots$ be a sequence of i.i.d. copies of the r.v. W_1, and let $\tau_1^*, \tau_2^*, \ldots$ be their corresponding cycle lengths. Set

$$T^*(n) = \tau_1^* + \ldots + \tau_n^*$$

$$N^*(t) = \max\{n \geq 0 : T^*(n) \leq t\}.$$

Then, we define X^* through the formula $X^*(t) = W^*_{N^*(t)+1}(t - T^*(N^*(t)))$.

The process X^* is a non-delayed classically regenerative process having i.i.d. cycles consisting of the random elements $W_1^*, W_2^*, \ldots$. (We adopt the convention that a process having i.i.d. cycles is a (non-delayed) classically regenerative process.) Set

$$\alpha^*(t) = \frac{1}{t}\int_0^t f(X^*(s))ds$$

Standard arguments establish that if f is non-negative, $\alpha^*(t) \to \alpha$ a.s. as $t \to \infty$, so that the steady-state of X^* is identical to that of X. Thus, we have constructed a classically regenerative process that has precisely the same long-run behavior as that of the original l-dependent process.

PROPOSITION 2. *Suppose that f is non-negative and $E[Y_1^2 + \tau_1^2] < \infty$. Then,*

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

as $t \to \infty$, where $\sigma^2 = \lambda E Z_1^2$.

This easily established result shows that if the computational effort required to simulate t time units of X^* is identical to that of X, then $\alpha^*(t)$ is more efficient than $\alpha(t)$ if and only if $EZ_1Z_2 \geq 0$. Since positive correlations are typical of most real-world stochastic systems, it is evident that $\alpha^*(t)$ is often the better estimator. Our remaining discussion will therefore focus on the analysis of classically regenerative processes.

3. OUTPUT ANALYSIS FOR REGENERATIVE STEADY-STATE SIMULATION

Let $X = (X(t) : t \geq 0)$ be a non-delayed classically regenerative process with associated regeneration times $0 = T(0) < T(1) < \ldots$ If f is a non-negative function, then

$$\alpha(t) = \frac{1}{t}\int_0^t f(X(s))ds \to \pi f$$

a.s. as $t \to \infty$. In order to construct confidence intervals for α based on the estimator $\alpha(t)$, it is necessary to estimate the time-average variance constant σ^2 appearing in Proposition 2. Let

$$s(t) = \sqrt{\frac{1}{t}\sum_{i=1}^{N(t)}(Y_i - \alpha(t)\tau_i)^2}$$

where $N(t) = \max\{n \geq 0 : T(n) \leq t\}$. If $E[Y_1^2 + \tau_1^2] < \infty$, it is evident that $s(t) \to \sigma$ a.s. as $t \to \infty$, and consequently

$$t^{\frac{1}{2}}(\alpha(t) - \alpha)/s(t) \Rightarrow N(0,1)$$

as $t \to \infty$ provided that $\sigma^2 > 0$. Hence, it follows that

$$[\alpha(t) - z\frac{s(t)}{t^{\frac{1}{2}}}, \quad \alpha(t) + z\frac{s(t)}{t^{\frac{1}{2}}}]$$

is an asymptotic 100(1-δ)% confidence interval for α, when z is selected so that

$$P[N(0,1) \leq z] = 1 - \frac{\delta}{2}.$$

We note that the estimator $\alpha(t)$ is defined completely independently of the regenerative structure of X. The regenerative structure does, however, play a critical role in defining the estimator $s(t)$. It is the role that regeneration plays in constructing $s(t)$ that is the core of the so-called "regenerative method" of steady-state simulation.

Alternative estimators for σ^2 can be based on the fact that σ^2 can typically be re-expressed in the form

$$\sigma^2 = 2\int_0^\infty \mathrm{cov}(\tilde{X}(0), \tilde{X}(t))dt$$

where $(\tilde{X}(t) : t \geq 0)$ is the stationary version of the regenerative process X (see Glynn (1990) for the technical details.) The constant σ^2 is therefore just a multiple of the spectral density of $\tilde{X}$ evaluated at zero; see Anderson (1971) for definitions and results pertaining to stationary time series. The constant σ^2 may consequently be estimated by appealing to the substantial literature on spectral density estimation that has been developed by time series analysts. Such estimators typically converge to σ^2 at a rate $t^{-\beta}(\beta < \frac{1}{2})$ in the run-length t.

By contrast, the estimator $s(t)$ associated with the regenerative method of simulation converges at rate $t^{-1/2}$ in the run-length t. In fact, $s(t)$ and $\alpha(t)$ jointly satisfy a CLT with convergence rate $t^{-1/2}$. Recall that $Z_i = Y_i - \alpha\tau_i$. Put

$$A_i = Z_i^2 - \sigma^2\tau_i$$

$$\beta = 2\lambda E Z_1\tau_1.$$

If $h(x) = x^{\frac{1}{2}}$, it is evident that

$$\begin{aligned}
s(t) - \sigma &= h(s^2(t)) - h(\sigma^2) \\
&\approx h'(\sigma^2)(s^2(t) - \sigma^2) \\
&\approx \frac{1}{2\sigma} \cdot \frac{1}{t}\sum_{i=1}^{N(t)}[(Z_i - (\alpha(t) - \alpha)\tau_i)^2 - \sigma^2\tau_i \\
&= \frac{1}{2\sigma} \cdot \frac{1}{t}\sum_{i=1}^{N(t)}[A_i - (\alpha(t) - \alpha)Z_i\tau_i + (\alpha(t) - \alpha)^2\tau_i^2] \\
&= \frac{1}{2\sigma}\left(\frac{1}{t}\sum_{i=1}^{N(t)} A_i - \frac{\beta}{t}\sum_{i=1}^{N(t)} Z_i + O_p(\frac{1}{t})\right) \\
&\approx \frac{1}{2\sigma}\left(\frac{1}{t}\sum_{i=1}^{\lfloor\lambda t\rfloor} A_i - \frac{\beta}{t}\sum_{i=1}^{\lfloor\lambda t\rfloor} Z_i + o_p(\frac{1}{\sqrt{t}})\right).
\end{aligned}$$

In combination with (2.3), this basically establishes the following joint CLT.

THEOREM 3. *Suppose that f is a non-negative function for which $E[Y_1^4 + \tau_1^4] < \infty$. Then, if $\sigma^2 > 0$,*

$$t^{\frac{1}{2}}(\alpha(t) - \alpha, s(t) - \sigma) \Rightarrow N(\vec{0}, C),$$

where the covariance matrix C has elements defined by

$$\begin{aligned} C_{11} &= \lambda E Z_1^2 \\ C_{12} = C_{21} &= \lambda E[A_1 Z_1 - \beta Z_1^2]/2\sigma \\ C_{22} &= \lambda E[(A_1 - \beta Z_1)^2]/4\sigma^2. \end{aligned}$$

For a complete proof, see Glynn and Iglehart (1987)

We note that the estimator $\alpha(t)$ is defined in a manner that is completely independent of the regenerative structure of X, and hence we ought to expect that the constant C_{11} is independent of the particular choice of regenerative structure exploited. Specifically, in the countable state Markov process setting, the constant C_{11} is independent of the choice of the particular state used to "split" the sample path into regenerative cycles; see, for example, Chung (1967).

By contrast, since the estimator $s(t)$ is highly sensitive to the way in which one defines the "cycle boundaries," the constant C_{22} typically depends on the regenerative state used. Surprisingly, it turns out that the covariance term C_{12} is, like C_{11}, independent of the choice of regenerative state exploited; see Calvin (1989). We should point out that while simulation "folklore" suggests that one obtains maximal efficiency in the regenerative method by choosing to use the regenerative state with the smallest mean cycle length, the variance term C_{22} is not necessarily minimized by this particular choice. The question of how to choose a regenerative state in an optimal fashion remains open.

The rate $t^{-\frac{1}{2}}$ obtained for the estimator $s(t)$ takes advantage of the fact that since the cycle boundaries are known, one does not have to estimate covariance effects between adjacent regenerative cycles. The slower rate $t^{-\beta}(\beta < \frac{1}{2})$ associated with spectral density estimators appears largely because such estimators can make no such simplification in general.

We note that this enhanced convergence rate is not, however, costless. Specifically, any computational procedure based on using regeneration needs to explicitly identify regenerative cycle boundaries along the sample path simulated. In some settings, this can be done cheaply (e.g. Markov processes that return to points infinitely often). However, when it is necessary to use the randomized regenerations associated with the general theory of Harris recurrence, the computational and coding effort required to implement the regenerative method may be prohibitive. Nevertheless, there is a rich class of applications for which regenerative techniques appear quite feasible.

Thus far, we have focused on discussion of the estimation questions related to the time-average variance constant σ^2. A second statistical issue that is of some interest concerns the bias of the point estimator α. Specifically, it turns out that in the presence of classically regenerative structure that one can "correct" for the bias of $\alpha(t)$.

THEOREM 4. *Suppose that X is a non-delayed classically regenerative process for which τ_1 is non-arithmetic. If f is a non-negative function and $E[Y_1^2 + \tau_1^2] < \infty$, then*

$$E\alpha(t) = \alpha - \frac{\lambda}{t} E\Big[\int_0^{\tau_1} s\Big(f\big(X(s)\big) - \alpha\Big) ds\Big] + \circ(\frac{1}{t})$$

as $t \to \infty$.

This result first appeared in Glynn (1987), and takes advantage of the fact that $a(t) = Ef(X(t))$ satisfies the renewal equation

$$a = b + F * a$$

where $F(dt) = P[\tau_1 \in dt]$ and $b(t) = E[f(X(t)); \tau_1 > t]$. It follows from standard arguments that

$$a = U * b$$

where $U = \sum_{n=0}^{\infty} F^{(n)}$. The result basically follows from the following theorem.

THEOREM 5. *Suppose that $U = \sum_{n=0}^{\infty} F^{(n)}$, where F is the distribution function of a positive r.v. having finite mean λ^{-1}. If b is a non-negative function, and $a = U * b$, then*

$$\frac{1}{t}\int_0^t a(s)ds \to \lambda\kappa, \tag{3.1}$$

where $\kappa = \int_0^{\infty} b(s)ds$. If, in addition, F is non-arithmetic and $\int_0^{\infty} s[b(s) + \overline{F}(s)]ds < \infty$ (where $\overline{F}(t) = 1 - F(t)$), then

$$\frac{1}{t}\int_0^t a(s)ds = \lambda\kappa - \frac{\lambda}{t}\int_0^{\infty} s[b(s) - \lambda\kappa\overline{F}(s)]ds + \circ(\frac{1}{t})$$

as $t \to +\infty$.

PROOF. To prove (3.1), we first assume that $\kappa < +\infty$. Then, for each $\epsilon > 0$, there exists T such that for $t \geq T$,

$$\int_0^t b(s)ds \geq \kappa - \epsilon.$$

We now observe that

$$\begin{aligned}
\int_0^t a(s)ds &= \int_0^t (U * b)(s)ds \\
&= \int_0^t \int_{[0,s]} U(du)b(s-u)ds \\
&= \int_{[0,t]} U(du) \int_u^t b(s-u)ds \\
&= \int_{[0,t]} U(du) \int_0^{t-u} b(s)ds \qquad (3.2)
\end{aligned}$$

Then, for $t \geq T$,

$$\int_{[0,t-T]} U(du) \int_0^T b(s)ds \leq \int_0^t a(s)ds \leq U(t) \int_0^t b(s)ds \leq U(t)\kappa.$$

Dividing by t and applying the elementary renewal theorem to the extreme members of this inequality yields

$$\lambda(\kappa - \epsilon) \leq \underline{\lim} \frac{1}{t} \int_0^t a(s)ds \leq \overline{\lim} \int_0^t a(s)ds \leq \lambda\kappa.$$

Since ϵ was arbitrary, this proves the result when $\kappa < +\infty$. For $\kappa = +\infty$, we can appropriately "truncate" $b(\cdot)$, take advantage of the result just established, and then remove the truncation by passing to the limit.

To establish the second part of the result, we note that (3.2) gives

$$\begin{aligned}
&\int_0^t a(s)ds - \lambda\kappa t \\
&= \int_0^t a(s)ds - U(t) \int_0^\infty b(s)ds + [U(t) - \lambda t] \cdot \kappa \\
&= \int_{[0,t]} U(du) \int_0^{t-u} b(s)ds - U(t) \int_0^\infty b(s)ds + [U(t) - \lambda t]\kappa \\
&= -\int_{[0,t]} U(du) \int_{t-u}^\infty b(s)ds + [U(t) - \lambda t]\kappa \\
&= -(U * c)(t) + [U(t) - \lambda t]\kappa,
\end{aligned}$$

where $c(t) = \int_t^\infty b(s)ds$. Since $c(\cdot)$ is a non-increasing integrable function under our assumptions, it follows that $c(\cdot)$ is directly Riemann integrable. Hence, the key renewal theorem applies and

$$\begin{aligned}
(U * c)(t) &\to \lambda \int_0^\infty c(s)ds \\
&= \lambda \int_0^\infty s\,b(s)ds
\end{aligned}$$

as $t \to \infty$. Also, it is well known that when F is non-arithmetic,

$$U(t) - \lambda t \to \lambda^2 \int_0^\infty s(1 - F(s))ds$$

as $t \to \infty$. Combining the two results yields the conclusion that

$$\int_0^t a(s)ds - \lambda\kappa t \to \lambda \int_0^\infty s[\lambda\kappa\overline{F}(s) - b(s)]ds$$

as $t \to \infty$, proving the theorem. $\|$

By appropriately substituting $b(t) = E[f(X(t)); \tau_1 > t]$ into Theorem 5, and noting that $\int_0^\infty b(t)dt = EY_1$, Theorem 4 follows.

Since the $o(\frac{1}{t})$ term appearing in Theorem 4 is determined by the rate of convergence in the renewal theorem, it is evident that it often converges to zero exponentially rapidly as $t \to \infty$ (assuming, for example, that τ_1 has exponential moments).

Based on Theorem 4, it seems reasonable to consider the estimator

$$\tilde{\alpha}(t) = \alpha(t) + \frac{1}{t^2} \sum_{i=1}^{N(t)} \int_0^{\tau_i} s\Big(f\big(X_i(T(i-1)+s)\big) - \alpha(t)\Big) ds.$$

While the correction added to $\alpha(t)$ to form $\tilde{\alpha}(t)$ is asymptotically negligible relative to $\alpha(t)$, it does have a positive impact on the bias properties of the new estimator. Specifically, Glynn (1987) shows that

$$E\tilde{\alpha}(t) = \alpha + \circ(\frac{1}{t})$$

as $t \to \infty$ (under certain regularity conditions). Typically, the $\circ(\frac{1}{t})$ term converges to zero at rate $\frac{1}{t^2}$. (It fails to converge at exponential rate because of the error introduced by having to estimate the bias term.)

Another means of reducing bias (or even eliminating it) is to simulate a stationary version of the regenerative process. In Asmussen, Glynn, and Thorisson (1992), it is shown that generation of a stationary version is possible for certain classes of regenerative processes (even processes for which no explicit representation of the stationary distribution is available).

ACKNOWLEDGEMENTS

The author wishes to thank the two referees for their helpful suggestions. In addition, the author wishes to acknowledge the support of the National Science Foundation under grant DDM-9101580, and the Army Research Office under Contract No. DAAL03-91-G-0319.

REFERENCES

[1] Anderson, T. W. (1971). *The Statistical Analysis of Time Series*. John Wiley & Sons, NY.

[2] Asmussen, S., P. W. Glynn, and H. Thorisson (1992). Stationarity Detection in the Initial Transient Problem. *ACM Transactions on Modeling and Computer Simulation* 2, 130-157.

[3] Azema, J., M. Duflo and D. Revuz (1967). Mesure invariate sur les classes le'currents. *Z. Wahrschlichkeitsth verw. Gebiete* 8, 157-181.

[4] Calvin, J. (1989). Stochastic Optimization Algorithms and Moment Formulas for Markov chains. Ph.D. Dissertation, Dept. of Operations Research, Stanford University.

[5] Glynn, P.W. (1982). Simulation Output Analysis for General State Space Markov Chains. Ph.D. Dissertation, Dept. of Operations Research, Stanford University.

[6] Glynn, P.W. and D.L. Iglehart (1987). A joint central limit theorem for the sample mean and regenerative variance estimator. *Annals of Operations Research* 8, 41-55.

[7] Glynn, P.W. (1987). A low bias steady-state estimator for equilibrium processes. Technical Report, Dept. of Operations Research, Stanford University.

[8] Glynn, P.W. (1990). The covariance function of a regenerative process. Technical Report, Dept. of Operations Research, Stanford University.

[9] Renyi, A. (1970). *Foundations of Probability.* Holden-Day, San Francisco.

[10] Revuz, D. (1984). *Markov Chains.* New Holland Publishing Company, Amsterdam.

[11] Sigman, K. (1990). One-dependent regenerative processes and queues in continuous time. *Mathematics of Operations Research* 15, 175-189.