Regenerative Steady-State Simulation of Discrete-Event Systems

SHANE G. HENDERSON
Cornell University
and
PETER W. GLYNN
Stanford University

The regenerative method possesses certain asymptotic properties that dominate those of other steady-state simulation output analysis methods, such as batch means. Therefore, applying the regenerative method to steady-state discrete-event system simulations is of great interest. In this paper, we survey the state of the art in this area. The main difficulty in applying the regenerative method in our context is perhaps in identifying regenerative cycle boundaries. We examine this issue through the use of the "smoothness index." Regenerative cycles are easily identified in systems with unit smoothness index, but this is typically not the case for systems with nonunit smoothness index. We show that "most" (in a certain precise sense) discrete-event simulations will have nonunit smoothness index, and extend the asymptotic theory of regenerative simulation estimators to this context.

Categories and Subject Descriptors: G.3 [**Probability and Statistics**]: *Markov processes*; I.6.6 [**Simulation and Modeling**]: Simulation Output Analysis

General Terms: Theory

INTRODUCTION

Discrete-event modeling and simulation is one of the most widely used techniques in operations research today. A widely accepted model of a discrete-event dynamical system (DEDS) is a generalized semi-Markov process (GSMP) [Whitt 1980; Glynn 1989b; Shedler 1993; Haas 1999]. Except in very special

The work of the first author was partially supported by New Zealand Public Good Science Fund Grant UOA803 and National Science Foundation Grant DMI 0085167. The work of the second author was partially supported by Army Research Office Contract No. DAAG55-97-1-0377-P0001 and National Science Foundation Grant DMS-9704732-001.
Authors' addresses: Shane G. Henderson, School of Operations Research and Industrial Engineering, Cornell University, Ithica, NY 14853; email: shane@orie.cornell.edu; Peter W. Glynn, Management Science and Engineering, Terman Engineering Center Room 429A, Stanford University, Stanford, CA 94305-4023; email: glynn@leland.stanford.edu.
Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.
© 2001 ACM 1049-3301/01/1000–0313 $5.00

cases, closed-form analytical solutions for performance measures of GSMPs are not known, so one turns to simulation for analysis of these systems.

Simulation output analysis methods allow one to make a statistically valid statement about the output from discrete-event simulations. The regenerative method for steady-state output analysis is one such method that holds great appeal, for several reasons. First, the problem of "initialization bias" does not arise in regenerative simulation [Bratley et al. 1987; Law and Kelton 2000]. Second, regenerative estimators are relatively simple to construct. Third, it is known that the regenerative method has the fastest *asymptotic* rate of convergence (see Section 1) of all time-average variance estimation methods. Of course, this final argument is subject to the usual proviso that it is an asymptotic result, and may not be true for a finite runlength, perhaps because of excessively long regenerative cycles. Further comments on this issue may be found in Section 5.

To motivate the content of this paper we need some notation. Let $\tilde{X} = (\tilde{X}(t) : t \geq 0)$ be a stochastic process on state space S that models a DEDS, and let $f : S \to \mathbb{R}$ be a real valued cost (or reward) function on the state space S. The steady-state estimation problem is the problem of estimating the "long run average cost"

$$\alpha = \lim_{t\to\infty} \frac{1}{t} \int_0^t f(\tilde{X}(s))\, ds,$$

when this limit exists. If the process $\tilde{X}$ is regenerative, then under mild conditions α exists and may be written as $\mathbb{E}\, Y_1 / \mathbb{E}\, \tau_1$, where Y_i and τ_i represent the cost accumulated over, and the length of, the ith regenerative cycle respectively. Therefore, one can estimate α via

$$\alpha_n = \frac{\sum_{i=1}^n Y_i}{\sum_{i=1}^n \tau_i}.$$

Associated with the estimator α_n is a "time-average variance constant" (TAVC) estimator v_n^2 that, together with α_n, may be used to construct confidence intervals for α.

It is known (see Section 2) that any "well-behaved" steady-state simulation is regenerative, and the resulting cycles are either independent, or one-dependent. The independent cycle case has been extensively studied, but the one-dependent cycle case has not. This suggests our focus on (i) and (ii) below. In principle, (i) and (ii) provide all the steady-state simulation methodology one will ever need for such well-behaved steady-state simulations. On the other hand, we ultimately need to apply these ideas in the context of DEDS. So, we need some basic theory to establish when classical (independent) regeneration pertains, as opposed to one-dependent regeneration (see (iii) below).

We view the primary contributions of this paper as follows.

(i) The paper surveys the state-of-the-art for the basic mathematical theory associated with the regenerative method for steady-state simulation. This is done, in part, to provide a "frame of reference" for the later results on the extension of the regenerative method to processes where the regenerative cycles are one-dependent.

(ii) The paper extends the basic theory and methodology known for classically regenerative processes to one-dependent processes. In particular, we obtain a joint central limit theorem (Theorem 8) for α_n and v_n that extends a result by Glynn and Iglehart [1987] for the independent case. Furthermore, we show that the bias in α_n is $O(n^{-1})$ (Theorem 9), where the notation $O(g(n))$ represents a deterministic sequence $(d_n : n \geq 1)$ with the property that $|d_n| \leq Dg(n)$ for some finite constant D. The estimator α_n involves a fixed number of cycles, but a random length of simulated time. One might instead consider the estimator $\alpha_{N(t)}$, where $N(t)$ is the number of completed regenerative cycles by simulated time t. Meketon and Heidelberger [1982] showed that in the independent case, the bias in $\alpha_{N(t)}$ is $O(t^{-1})$, while that of $\alpha_{N(t)+1}$ is $O(t^{-2})$. This result was further explained by Glynn and Heidelberger [1990], who showed that it results from a fortunate cancellation in asymptotic expressions for the bias. We extend the Meketon and Heidelberger result to the one-dependent case (Theorem 10) by showing that the bias in the estimator $\alpha_{N(t)+2}$ is $o(t^{-1})$, so that the bias goes to 0 at a rate that is faster than t^{-1}. (The notation $o(g(n))$ represents a deterministic sequence $(d_n : n \geq 1)$ with the property that $d_n/g(n) \to 0$ as $n \to \infty$.) To implement this estimator, one simply needs to complete the cycle in progress at time t, and the following cycle as well.

(iii) A major thrust of this paper is to put the implications of the above methodology and theory into the DEDS environment, making clear to the simulation community precisely the class of DEDS to which the classical regenerative method applies, versus those to which the one-dependent regenerative method applies. In this vein, our theory precisely partitions (for a large class of DEDS) the family of GSMPs to which the two methods potentially apply; see Theorem 17 and Proposition 18. This theoretical partitioning is perhaps the single most important result in the paper.

(iv) As a by-product of our theoretical analysis in (iii), we obtain new results on uniqueness of invariant (steady-state) measures for DEDS (Corollary 13), which sheds light on the related question of when the steady-state distribution is independent of the initial condition. We also obtain new results on when a GSMP is nonexplosive (Theorem 15).

(v) We survey and discuss implementation issues regarding use of these regenerative-type methods in the DEDS context.

This paper is primarily a theoretical contribution to the literature on the regenerative method as it applies to discrete-event systems. We do not believe that the paper proposes the final solution for dealing with models in which nonunit smoothness indices arise. Rather, the paper clearly identifies this as an issue that will need to be surmounted if the regenerative method is to be applied, in generality, to such simulations, and it discusses the modifications to conventional regenerative methodology that must be made in the presence of the one-dependent cycles that arise in such a context.

This paper is organized as follows.

In Section 1 we describe GSMP models of DEDS, and show how one can formally define the GSMP through a related general state-space Markov chain

(GSSMC). GSMPs with "single-states" easily admit a regenerative analysis. We outline this method for determining regenerative cycles, and provide several asymptotic results for estimators constructed in this fashion.

In Section 2, we argue that one may restrict attention to GSMPs that give rise to GSSMCs that are positive Harris recurrent, with essentially no loss of generality. This is the "well-behaved" result referred to earlier. We will also see that positive Harris recurrent chains may be classified according to their "smoothness index."

In Section 3, we consider chains with unit smoothness index. Two methods are reviewed for determining regeneration points, and we show that if the chain exhibits so-called "classical regenerative" behaviour, then it must have unit smoothness index. The significance of this result is that the chains that are discussed in the next section *cannot* exhibit classically regenerative behaviour.

Section 4 provides an analysis of chains with nonunit smoothness index. We explain that a weakened form of regeneration exists for such chains, so that the resulting cycles are one-dependent, and provide asymptotic theory for estimators based on this form of regeneration.

We also review the known methods for detecting regeneration times in chains with nonunit smoothness index m say. To implement these methods, one basically requires knowledge of the m-step transition probabilities of a GSSMC. The m-step transition probabilities are typically difficult to compute (although Henderson and Glynn [1999a] offer one promising direction for dealing with this problem). Therefore, characterizing the class of discrete-event systems with unit smoothness index is of great interest, and this is the subject of the remainder of the paper.

Perhaps the most important result in this paper is Theorem 17, which basically shows that in the absence of "event cancellation," the only systems with unit smoothness index are those with single states.

The consequences of this observation, and the other results in the paper, for regenerative simulation of discrete-event systems are discussed in Section 5.

Unless otherwise stated, proofs are given in Section 6.

1. GENERALIZED SEMI-MARKOV PROCESSES WITH SINGLE-STATES

We begin by defining a GSMP, following Whitt [1980] and Glynn [1989b] closely. In contrast to Glynn, and as in Whitt, our event clocks record residual lives and not time since last activation of the clock, thereby matching simulation software more closely. In contrast to Glynn and Whitt, and as in Haas [1999], we assume that state transitions are triggered by *sets* of events, and not necessarily single events, thus bypassing the difficulty of ensuring unique triggering events. Here we describe a *time-homogeneous* GSMP, henceforth referred to simply as a GSMP.

Let S be a (finite or countable) collection of *states*, and let E be a finite collection of *events*. Associated with each state $s \in S$, there is a set of *active events* $E(s) \subseteq E$ that can trigger a state transition out of s. If E^* denotes a set of events that simultaneously trigger a transition from state s, then the new GSMP state is chosen according to the probability mass function $p(\cdot\,; s, E^*)$ on S

independent of all else. We require $p(\cdot\,; s, E^*)$ to be a probability mass function on S for each $s \in S$, and each $E^* \subseteq E(s)$.

Associated with each event $e \in E$ is a *clock reading* c_e that indicates the amount of time remaining before the event e is scheduled to occur. When the clock reading c_e runs down to zero, the associated event e (together with any other active events in E that occur simultaneously) will trigger a state transition of the GSMP. We allow clocks to count down at different *rates*. When the GSMP is in state s, and $e \in E(s)$, the reading c_e decreases at rate r_{se}. Hence, if it were uninterrupted, event e would trigger a state transition in c_e/r_{se} time units.

All inactive clocks and rates are set to 0: $c_e = r_{se} = 0$ for all $e \notin E(s)$. Furthermore, we require that $r_{se} > 0$ for at least one $e \in E(s)$ for all $s \in S$, so that the GSMP does not "stall" in state s.

Define the *clock-reading vector* $c = (c_e : e \in E(s))$ so that c consists of the (ordered) list of clock readings for all active events when the GSMP is in state s. Now, when the GSMP is in state s, the time $\Delta = \Delta(s, c)$ until the next state transition is given by

$$\Delta = \min\{c_e/r_{se} : e \in E(s), r_{se} > 0\},$$

and the set of events which achieve this minimum is denoted E^*. Suppose that the GSMP then moves to state s'. The clock readings are updated as follows. Each event $e \in \mathcal{O}(s', s, E^*) \stackrel{\triangle}{=} E(s') \cap (E(s) - E^*)$ is an "old" event that remains active in the new state s'. Its clock reading c_e must be adjusted for the time that passed while in state s, so that the new reading is $c_e^* \stackrel{\triangle}{=} c_e - r_{se}\Delta$. For each new event $e' \in \mathcal{N}(s', s, E^*) \stackrel{\triangle}{=} E(s') - \mathcal{O}(s', s, E^*)$, the associated clock reading $c_{e'}$ is sampled, independent of all else, according to a distribution function $F(\cdot\,; s', e', s, E^*)$ that may depend on the old and new states, the triggering events and the new event. To ensure that the new clock reading is nonnegative, we require that $F(0; s', e', s, E^*) = 0$ for all s, s', e' and E^*. Each newly inactive event $e \in E(s) - E(s')$ has both its clock reading c_e and its rate r_{se} set to 0. If E^* consists of a single event e say, then we write $\mathcal{N}(s', s, e)$ for $\mathcal{N}(s', s, \{e\})$, $p(s'; s, e)$ for $p(s'; s, \{e\})$, and so forth.

The GSMP is a continuous-time stochastic process that evolves on state space S. To formally define the GSMP, we will first define its related GSSMC $X = (X_n : n \geq 0)$. The idea is that X_n represents the state and clock readings of the GSMP immediately after its nth state transition, and we will write $X_n = (S_n, C_n)$, where S_n is the state of the GSMP and C_n is the vector of active clock readings.

Let $\mathbb{R}_+ = [0, \infty)$, and $\mathbb{R}_+^k$ be the k-fold product space. Let R_s be the set of possible active clock readings in state s, so that

$$R_s = \mathbb{R}_+^{|E(s)|},$$

and let $\Sigma = \bigcup_{s \in S} \{s\} \times R_s$. The set Σ is the state-space of our GSSMC. Let $\mathcal{S}$ denote the usual product topology on Σ consisting of counting measure on S and Lebesgue measure on R_s. Let X_0 have distribution μ, where μ is a probability measure on $(\Sigma, \mathcal{S})$.

We can now define the transition kernel P on Σ by setting

$$P((s,c),A) = p(s';s,E^*) \prod_{e'\in\mathcal{N}(s',s,E^*)} F(a_e;s',e',s,E^*) \prod_{e\in\mathcal{O}(s',s,E^*)} I(c_e^*\in[0,a_e]), \quad (1)$$

where

$$A = \{s'\}\times\{c'\in R_{s'} : 0\le c'_e\le a_e \quad \forall e\in E(s')\},$$

and $I(\cdot)$ is the indicator function that is 1 if its argument is true, and 0 otherwise. It is a standard measure-theoretic result (see Theorem 3.3 of Billingsley [1986] for example) that (1) uniquely specifies P. The definition of P, together with the initial distribution μ, uniquely defines the GSSMC X. The GSMP may now be defined.

Let ξ_n be the time of the nth transition in the GSMP, so that $\xi_0 = 0$, and for $n \ge 0$, $\xi_{n+1} = \xi_n + \Delta(S_n, C_n)$. We henceforth assume that the GSMP is *nonexplosive* in that $\xi_n \to \infty$ as $n\to\infty$ P_μ almost surely, where $P_\mu(\cdot) \stackrel{\triangle}{=} \int_\Sigma P(\cdot|X_0 = x)\mu(dx)$. This condition holds, for example, when the state-space is finite [Haas 1985, Proposition 2.1.18]. The GSMP $\tilde{X} = (\tilde{X}(t) : t\ge 0)$ may now be defined by

$$\tilde{X}(t) = S_{N(t)},$$

where

$$N(t) = \sup\{n : \xi_n \le t\}$$

is the number of state transitions to occur by time t. The nonexplosive assumption implies that $\tilde{X}(t)$ is defined for all $t \ge 0$.

See Shedler [1993] for examples of GSMPs.

Definition 1. Let $X = (X_n : n\ge 0)$ be a discrete-time stochastic process. For an increasing sequence of finite random times $T = (T(k) : k \ge -1)$ with $T(-1) = 0$ and $T(n)\to\infty$ with probability 1 as $n\to\infty$, define the collection of random elements (cycles) $W = (W_n : n\ge 0)$ constructed from T and X, where, for $n\ge 0$, $W_n = (W_n(k) : k\ge 0)$ is defined by

$$W_n(k) = \begin{cases} X_{T(n-1)+k} & \text{if } 0\le k < T(n)-T(n-1), \quad \text{and} \\ \Xi & \text{otherwise,} \end{cases}$$

and Ξ is some distinguished point not contained in the state space of X. We say that X is *classically regenerative* if there exists such a sequence T with the property that $W_0, W_1, \ldots$ are independent, and $W_1, W_2, \ldots$ are identically distributed.

Remark 1. The point Ξ acts as a "cemetery state," with all cycles of the process eventually absorbed in Ξ. Note that Ξ is simply a device to enable the cycles to be defined for all $k\ge 0$, and not just up to an almost surely finite random time.

Remark 2. An analogous definition applies when X is a continuous-time process.

One method for determining regeneration times in GSSMCs that arise from GSMPs is to identify *single-states*. This method for defining regenerations for GSMPs has been well studied; see Fossett [1979], Haas [1985], and Shedler [1987, 1993]. Shedler [1993] discusses several queueing network models in which single-states can be identified, and provides references to further applications.

Definition 2. Let $\tilde{X}$ be a GSMP as defined above. We say that $s \in S$ is a *single-state* if $|E(s)| = 1$, that is, there is only one active event when the GSMP is in state s.

When a GSMP leaves a single-state, all active clocks in the new state will have been set at the time of the transition, and are therefore independent of the previous history of the chain. A regeneration therefore occurs. Let $P_x(\cdot) \triangleq P(\cdot|X_0 = x)$.

PROPOSITION 1. *Let $\tilde{X}$ be a GSMP as defined above, and let X be the related GSSMC. Suppose that $\tilde{X}$ has a single state s^* such that*

$$P_x(S_n = s^* \textit{ infinitely often}) = 1, \quad \textit{for all } x \in \Sigma.$$

Then, X is classically regenerative with regeneration times $(T(n) : n \geq 0)$ given by

$$\begin{aligned} T(0) &= \inf\{k \geq 1 : S_{k-1} = s^*\}, \quad \textit{and for } n \geq 0, \\ T(n+1) &= \inf\{k > T(n) : S_{k-1} = s^*\}. \end{aligned}$$

Furthermore, $\tilde{X}$ is regenerative, with regeneration times given by $\tilde{T}(n) = \xi_{T(n)}$ $(n \geq 0)$.

For a proof of this result, see Theorem 5.13 of Shedler [1993], or for a similar result, see Proposition 4.3 of Haas and Shedler [1987].

Example 1. Consider the classical single-server queue in which the interarrival times of customers to a server form a renewal process. Customers are served in first in-first out order, and service times are independently identically distributed, and independent of the arrival process. This system may be modeled using a GSMP with state space $S = 0, 1, 2, \ldots,$ and event list $E = \{A, B\}$, where A corresponds to an arrival, and B corresponds to a service completion.

For this system, S_n represents the number of customers present in the system immediately after the nth customer arrival or service completion. Similarly, C_n gives the time until the next customer arrival, and the time until the next customer service (if $S_n > 0$) immediately after the nth customer arrival or service. The GSMP $(\tilde{X}(t) : t \geq 0)$ gives the number of customers $\tilde{X}(t)$ in the system at time t.

If the mean service time is finite and less than or equal to the mean interarrival time, then the state $s^* = 0$ is a single-state, since only A is active when the system is empty, and the system empties infinitely often. The (integer) regeneration times $(T(n) : n \geq 0)$ for the Markov chain $((S_i, C_i) : i \geq 0)$ correspond to the ordered indices n of arrival events where a customer arrives to an empty system: $S_{n-1} = 0$ and $S_n = 1$. The (real-valued) regeneration times $\tilde{T}(n)$ for the GSMP then correspond to the simulated time that a customer arrives to an empty system.

These regeneration times correspond to epochs when the GSMP *leaves* a single state, rather than when it *enters* a single state. This may be a source of confusion for the reader familiar with regeneration times for irreducible, finite state space continuous-time Markov chains defined by the times at which the chain enters a distinguished state. Such times are regenerations because of the conditional independence of the future and past evolution of the process, given the state of the process at a given time. This special property does not hold for the more general GSMP. But notice that the times at which an irreducible, finite state space continuous-time Markov chain leaves a distinguished state are also regeneration times, and these regeneration times are analogous to the regeneration times in a GSMP with a single state.

For $n \geq 0$, the nth regenerative cycle $(W_n(t) : t \geq 0)$ is given by the number of customers in the system at time $\tilde{T}(n-1) + t$ for $0 \leq t < \tilde{T}(n) - \tilde{T}(n-1)$, and by -1 say, otherwise. Note that here, we are using -1 as the cemetery state Ξ.

We now turn to limit theory for regenerative GSMPs. In particular, we discuss how to estimate long-run averages of the form

$$\lim_{t\to\infty} t^{-1} \int_0^t f(\tilde{X}(u))\, du,$$

where $f : S \to \mathbb{R}$ is a real-valued function on the state space of the GSMP. Our first result is a strong law for classically regenerative GSMPs that establishes that this limit exists, and provides an expression for the limit in terms of regenerative quantities.

For $x \in \Sigma$, the state space of the discrete-time chain $(X_n : n \geq 0)$, let $\mathbb{E}_x(\cdot) \stackrel{\triangle}{=} \mathbb{E}(\cdot | X_0 = x)$. If ν is a distribution on $(\Sigma, \mathcal{S})$, let $\mathbb{E}_\nu(\cdot) \stackrel{\triangle}{=} \int_\Sigma \mathbb{E}_x(\cdot)\nu(dx)$.

THEOREM 2. *Let $\tilde{X}$ be a classically regenerative GSMP with regeneration times $(\tilde{T}(n) : n \geq 0)$, and let $f : S \to \mathbb{R}$. Set $\varphi(dy) = P_x(X_{T(0)} \in dy)$ (note that φ is defined independently of x). For $j \geq 1$, set*

$$\begin{aligned} \tau_j &= \tilde{T}(j) - \tilde{T}(j-1), \quad \textit{and} \\ Y_j &= \int_{\tilde{T}(j-1)}^{\tilde{T}(j)} f(\tilde{X}(u))\, du. \end{aligned} \tag{2}$$

If $f \geq 0$ and $\mathbb{E}_\varphi \tau_1 < \infty$, then

$$\alpha(t) \stackrel{\triangle}{=} t^{-1} \int_0^t f(\tilde{X}(u))\, du \to \alpha \stackrel{\triangle}{=} \mathbb{E}_\varphi Y_1 / \mathbb{E}_\varphi \tau_1 \textit{ almost surely}$$

as $t \to \infty$.

PROOF. This is a special case of Theorem 3.1, of Asmussen [1987, p. 136]. □

Remark 3. The quantities Y_j and τ_j can be easily computed, since the sample paths of $\tilde{X}$ are piecewise constant. In particular,

$$\tau_j = \sum_{k=T(j-1)}^{T(j)-1} \Delta(S_k, C_k), \quad \text{and}$$
$$Y_j = \sum_{k=T(j-1)}^{T(j)-1} \Delta(S_k, C_k) f(S_k).$$

Theorem 2 establishes that under moderate conditions, $\alpha(t)$ converges to a constant as $t \to \infty$. We now wish to assess the variability of $\alpha(t)$ for finite t. Suppose that $\mathbb{E}_\varphi \tau_1 < \infty$ and τ_1 has a spread out distribution. (We say that a random variable with distribution function F is spread out if, for some $n \geq 1$, the n-fold convolution of F with itself has an absolutely continuous component (see p. 140 of Asmussen [1987] for example.) Then a stationary version $\tilde{X}^* = (\tilde{X}^*(t) : t \geq 0)$ of the regenerative process $\tilde{X}$ exists, and $\tilde{X}(t)$ converges in distribution to $\tilde{X}^*(0)$ as $t \to \infty$ (Corollary 1.4, p. 141, Asmussen [1987]). Under quite general conditions,

$$\operatorname{var}\left(\frac{1}{t}\int_0^t f(\tilde{X}(u))\,du\right) \sim \frac{2}{t}\int_0^\infty \operatorname{cov}(f(\tilde{X}^*(0)), f(\tilde{X}^*(u)))\,du \tag{3}$$
$$\triangleq \frac{\sigma^2}{t} \tag{4}$$

as $t \to \infty$. (We say that $x_t \sim y_t$ as $t \to \infty$ if $x_t/y_t \to 1$ as $t \to \infty$.) This holds, for example, under mixing assumptions on the process $(f(\tilde{X}(t)) : t \geq 0)$ (see the proof of Theorem 20.1, Billingsley [1968, p. 174]) and if the collection of random variables

$$\left(\frac{1}{t}\left(\int_0^t [f(\tilde{X}(u)) - \alpha]\,du\right)^2 : t \geq 0\right)$$

is uniformly integrable.

The quantity σ^2 is known as the *time-average variance constant* (TAVC), and is somewhat difficult to estimate. Estimators for the TAVC based on spectral density estimation methods typically have mean squared errors (MSEs) that converge at rate $t^{-\beta}(\beta < 1)$, where t is the simulation time horizon (Grenander and Rosenblatt [1984, p. 129]). For nonoverlapping and overlapping batch means estimators of σ^2, the results of Goldsman and Meketon [1986] and Song and Schmeiser [1995] imply that the mean squared error converges to 0 at most at rate $t^{-2/3}$. So for either of these classes of estimators, the MSE of the estimator of σ^2 converges to 0 at an asymptotic rate that is of the order $O(t^{-\beta})$ where $\beta < 1$.

Estimators of σ^2 derived using the regenerative method provide an attractive alternative because of their simplicity, and as we shall see, we can typically expect their MSE to be $O(t^{-1})$.

The regenerative point estimator of α, based on n regenerative cycles is given by

$$\alpha_n = \bar{Y}_n/\bar{\tau}_n,$$

where $\bar{Y}_n$ and $\bar{\tau}_n$ are the sample means of $(Y_1, Y_2, \ldots, Y_n)$ and $(\tau_1, \tau_2, \ldots, \tau_n)$ respectively.

THEOREM 3. *Let $X = (X_n : n \geq 0)$ be a classically regenerative process and let $f : S \to \mathbb{R}$. Suppose that $\mathbb{E}_\varphi(\tau_1 + |Y_1|) < \infty$, and define $Z_i = Y_i - \alpha\tau_i$ for $i \geq 1$. Suppose that $0 < \mathbb{E}_\varphi Z_1^2 < \infty$.*

(1) *The central limit theorem (CLT)*

$$n^{1/2}(\alpha_n - \alpha) \Rightarrow \sigma_{\text{cyc}} N(0, 1)$$

holds, where $\Rightarrow$ denotes weak convergence, $N(0, 1)$ is a standard normal random variable, and $\sigma_{\text{cyc}}^2 = \mathbb{E}_\varphi Z_1^2 / (\mathbb{E}_\varphi \tau_1)^2$. (Here, the suffix cyc is meant to be mnemonic for "cycle").

(2) *The estimator*

$$v_n^2 \triangleq \frac{n^{-1} \sum_{i=1}^n (Y_i - \alpha_n \tau_i)^2}{(\bar{\tau}_n)^2}$$

of σ_{cyc}^2 is weakly consistent: $v_n^2 \Rightarrow \sigma_{\text{cyc}}^2$ as $n \to \infty$.

(3) *If, in addition, $\mathbb{E}_\varphi(Y_1^4 + \tau_1^4) < \infty$, then $n^{1/2}(\alpha_n - \alpha, v_n - \sigma_{\text{cyc}}) \Rightarrow N(0, \Lambda)$, where*

$$\Lambda = (\mathbb{E}_\varphi \tau_1)^{-2} \begin{bmatrix} \mathbb{E}_\varphi Z_1^2 & (2\sigma_{\text{cyc}}\, \mathbb{E}_\varphi \tau_1)^{-1}\, \mathbb{E}_\varphi A_1 Z_1 \\ (2\sigma_{\text{cyc}}\, \mathbb{E}_\varphi \tau_1)^{-1}\, \mathbb{E}_\varphi A_1 Z_1 & \left(4\sigma_{\text{cyc}}^2 (\mathbb{E}_\varphi \tau_1)^2\right)^{-1} \mathbb{E}_\varphi A_1^2, \end{bmatrix},$$

and $$A_i = Z_i^2 - \mathbb{E}_\varphi Z_1^2 - 2(\mathbb{E}_\varphi \tau_1)\sigma_{\text{cyc}}^2(\tau_i - \mathbb{E}_\varphi \tau_1) - 2(\mathbb{E}_\varphi Z_1 \tau_1 / \mathbb{E}_\varphi \tau_1) Z_i.$$

PROOF. The first two results are given in Glynn and Iglehart [1993]. The proof of the last result is similar to that of part (iv) of Theorem 8, and so is omitted. It may be regarded as the cycle-time version of Equation (3.3) in Glynn and Iglehart [1987]. □

Remark 4. The weak consistency in Part 2 of Theorem 3 can be strengthened to strong consistency (almost sure convergence) under the stronger moment hypothesis that $\mathbb{E}_\varphi(Y_1^2 + \tau_1^2) < \infty$.

Note that these central limit theorems are given in the time scale of regenerative cycles, whereas the TAVC σ^2 (see (4)) arose on the natural (simulated) time scale. When one adjusts the quantity σ_{cyc}^2 to take account of the time-change (see, for example, Wolff [1989, p. 124]), the new variance constant is $\sigma_{\text{cyc}}^2 \mathbb{E}_\varphi \tau_1$. Under certain moment and regularity conditions (Glynn [1989a, Theorem 5.5]), $\sigma_{\text{cyc}}^2 \mathbb{E}_\varphi \tau_1 = \sigma^2$.

We can typically expect the MSE of the estimator v_n^2 of σ_{cyc}^2 to be $O(n^{-1})$. To see why, observe that under the conditions of Theorem 3 Part 3, we can assert that $n^{1/2}(v_n^2 - \sigma_{\text{cyc}}^2) \Rightarrow \eta N(0, 1)$ as $n \to \infty$, for some appropriate constant η. Thus, $n(v_n^2 - \sigma_{\text{cyc}}^2)^2 \Rightarrow \eta^2 N(0, 1)^2$ as $n \to \infty$. Assuming the family of random variables $\{n(v_n^2 - \sigma_{\text{cyc}}^2)^2 : n \geq 1\}$ is uniformly integrable, it follows that $nE(v_n^2 - \sigma_{\text{cyc}}^2)^2 \to \eta^2$ as $n \to \infty$, that is, that the MSE of v_n^2 is $O(n^{-1})$.

Hence, we can typically expect the MSE of the regenerative variance estimator to decrease linearly with the number of regenerative cycles. This is in contrast to the sublinear rate of convergence exhibited by the TAVC estimators mentioned earlier.

In steady-state simulation, one must typically deal with the "initial transient" problem (see Bratley et al. [1987] and Law and Kelton [2000]). This occurs when the initial conditions of the simulation are not representative of steady-state conditions, so that point estimates are biased. The regenerative method sidesteps this particular difficulty because all calculations performed are based on cycle structure, irrespective of any initial transient period. However, bias is still exhibited through the fact that the estimator α_n is a *ratio* of sample means. Several methods are available to combat this difficulty. Iglehart [1975] discussed the relative merits of several estimators, concluding that a jackknife estimator should be used. Another method is suggested by Theorem 4. The proof of this result is similar to the proof of Theorem 7 in Glynn and Heidelberger [1990], and a special case of Theorem 9, and so is omitted.

THEOREM 4. *Let $X = (X_n : n \geq 0)$ be a classically regenerative GSMP, and let $(Y_i : i \geq 1)$, $(\tau_i : i \geq 1)$, and α_n be defined as above. If $|f|$ is bounded and $\mathbb{E}\,\tau_1^4 < \infty$, then*

$$\mathbb{E}\,\alpha_n = \alpha - \frac{1}{n}\frac{\mathbb{E}_\varphi Z_1 \tau_1}{(\mathbb{E}_\varphi \tau_1)^2} + o(n^{-1}).$$

Theorem 4 shows that the bias in the ratio estimator α_n decreases at rate n^{-1} as the number of cycles $n \to \infty$. Under appropriate uniform integrability conditions, the final result of Theorem 3 establishes that the MSE of α_n decreases at rate n^{-1}. Recall that $\text{MSE}(\alpha_n) = \text{Var}(\alpha_n) + \text{Bias}(\alpha_n)^2$, so that for large n, the dominant contribution to MSE is from the variance, which is $O(n^{-1})$. Hence, bias will only be a significant contributor to MSE in regenerative estimators if the number of cycles simulated is small. In such cases, one could use the estimator

$$\alpha_n' = \alpha_n + \frac{1}{n}\frac{n^{-1}\sum_{i=1}^{n}(Y_i - \alpha_n \tau_i)\tau_i}{(\bar{\tau}_n)^2},$$

instead of α_n. Glynn and Heidelberger [1990] give a similar bias-reducing estimator, and discuss when one might expect that bias is indeed reduced. We do not give a proof that bias is reduced using this new estimator because the proof is somewhat involved, and in any case, we believe that the following bias-reduction technique is typically more effective, and more easily implemented.

The estimator α_n is based on a fixed number of regenerative cycles n. One can also base an estimator of α on the (random) number of regenerative cycles completed by simulation time t. Define $N(t) = \max\{n \geq 0 : \tilde{T}(n) \leq t\}$ to be the number of identically distributed regenerative cycles completed by time t. Then one may estimate α using

$$\alpha(t) = \frac{\sum_{i=1}^{N(t)} Y_i}{\sum_{i=1}^{N(t)} \tau_i}.$$

Meketon and Heidelberger [1982] showed that this estimator typically exhibits bias of the order t^{-1}. They also showed that by completing the cycle in progress at time t, the bias properties of the estimator are improved. In particular, they gave conditions under which the bias in the estimator

$$\alpha'(t) = \frac{\sum_{i=1}^{N(t)+1} Y_i}{\sum_{i=1}^{N(t)+1} \tau_i}$$

is $O(t^{-2})$. Intuitively, the cycle in progress at time t is longer than a typical cycle because of length-biasing, and so completing the cycle has a large impact on the bias properties of the estimator. Glynn and Heidelberger [1990] looked at bias properties of such estimators, and carefully explain how this bias reduction arises. In Theorem 9 we prove that the bias in $\alpha'(t)$ is $o(t^{-1})$ in a more general setting.

2. HARRIS RECURRENCE OF GENERALIZED SEMI-MARKOV PROCESSES

In the previous section we saw how to define essentially any discrete-event simulation as a GSSMC. It is known that the problem of steady-state simulation of GSSMCs is well-posed (in a certain precise sense—see Theorem 5 below), if and only if the chain is positive Harris recurrent [Glynn 1982; Glynn 1994]. In view of this result, we may restrict our attention to the analysis of positive Harris recurrent Markov chains with essentially no loss of generality. Such chains enjoy a number of attractive properties, which we will exploit in this and later sections.

Recall that $P_x(\cdot) = P(\cdot|X_0 = x)$, that is, the probability on the path space of a chain $(X_n : n \geq 0)$ where $X_0 = x$. Let $\mathbb{E}_x$ denote the corresponding expectation operator.

Definition 3. The steady-state simulation problem is said to be *well-posed* for the Markov chain $X = (X_n : n \geq 0)$ on state space Σ if for every bounded real-valued measurable function g, there exists a number $c(g)$ such that for every $x \in \Sigma$

$$\frac{1}{n}\sum_{k=1}^{n} \mathbb{E}_x\, g(X_k) \to c(g)$$

as $n \to \infty$.

Note that $c(g)$ is required to be independent of the initial condition x. One would hope that the initial conditions play no role in the long-term behaviour of the system, so this definition seems reasonable.

Definition 4. Let $X = (X_n : n \geq 0)$ be a Markov chain on a complete, separable, metric space Σ. The chain X is said to be *Harris recurrent* if there exists a nonnegative function $\lambda : \Sigma \to [0, 1]$, a probability measure φ, an $\epsilon > 0$, and an $m \geq 1$ such that

(1) $P(X_m \in \cdot|X_0 = x) \geq \lambda(x)\varphi(\cdot), \forall x \in \Sigma$; and
(2) $P_x(\lambda(X_n) > \epsilon$ infinitely often$) = 1, \forall x \in \Sigma$.

Harris chains automatically possess a unique (up to a multiplicative constant) stationary measure π, and if $\pi(\Sigma) < \infty$, then π may be normalized to a probability, and X is then said to be *positive* Harris recurrent.

An intuitive discussion of Harris recurrence may be found in Henderson and Glynn [1999b]. We content ourselves with the following example.

Example 2. Let X be the GSSMC associated with a GSMP with a recurrent single state s^* (cf. Proposition 1). Then X is easily seen to be Harris recurrent. For if $x = (s, c)$, where s is the GSMP state and c is a vector of clock readings, then set $\lambda(x) = I(s = s^*)$, φ to be the regeneration distribution, $m = 1$ and $\epsilon \in (0, 1)$.

As mentioned earlier, the steady-state simulation problem for a given Markov chain is known to be well-posed if and only if the chain is positive Harris recurrent. The significance of this result is that we may, without loss of generality, restrict our attention to positive Harris recurrent Markov chains. For a proof of the following result, see Glynn [1994]. Theorem 17.1.7 of Meyn and Tweedie [1993] is similar, but assumes existence of a stationary probability distribution.

THEOREM 5. *The steady-state simulation problem is well-posed for the Markov chain X if and only if the chain X is positive Harris recurrent.*

When the number of states $|S|$ in a GSMP is finite, X is known to be a Harris chain under suitable conditions on the clock setting distributions. Before stating such a result, we need some definitions. Suppose that the GSMP $\tilde{X}$ and its related GSSMC X are defined as above.

Definition 5. We say that $s' \in S$ is *directly reachable* from $s \in S$ and write $s \to s'$ if $p(s'; s, e)r_{se} > 0$ for some $e \in E(s)$. We say that s' is *reachable* from s if there exist $s_1, s_2, \ldots, s_n \in S$ such that $s \to s_1 \to \cdots \to s_n \to s'$. The GSMP $\tilde{X}$ is said to be *irreducible* if s' is reachable from s for every $s, s' \in S$.

THEOREM 6. *[Haas 1999] Suppose that the GSMP $\tilde{X}$ is irreducible, the state space S is finite, and all clock rates r_{se} are positive. Suppose further that there exists $u \in (0, \infty)$ such that each clock setting distribution $F(\cdot; s', e', s, E^*)$ has a density function that is positive and continuous on $(0, u)$, and a finite first moment. Then the GSSMC X corresponding to $\tilde{X}$ is positive Harris recurrent.*

Conditions guaranteeing the applicability of the regenerative method were also given in König et al. [1967], Glynn [1989b], and Haas and Shedler [1987]. As discussed in Haas [1999], the sufficient conditions given in the above result are less restrictive than those in König et al. [1967] and Glynn [1989b]. They are more restrictive than those of Haas and Shedler [1987], but far easier to verify.

Clearly, continuous-time Markov chains on a countably infinite state space are a subclass of GSMPs. Since there is no general sharp result ensuring recurrence of such chains, we cannot expect to provide general sharp conditions under which countable-state GSMPs are recurrent. Any general recurrence theory for such GSMPs must necessarily involve imposition of Foster-Lyapunov hypotheses (see Meyn and Tweedie [1993]). In fact, even the finite-state space proof of Theorem 6 employs a Foster-Lyapunov type argument (to control the "continuous" clock readings of the related GSSMC).

Definition 6. Let X be a Harris recurrent Markov chain as in Definition 4. The *smoothness index* is the minimum value m such that Properties 1 and 2 of Definition 4 hold.

3. REGENERATIVE SIMULATION WITH UNIT SMOOTHNESS INDEX

The significance of Harris chains with unit smoothness index is that one can identify times $(T(k) : k \geq 0)$ such that X is classically regenerative, that is, the times $(T(k) : k \geq 0)$ yield independently identically distributed cycles. Therefore, the full power of Theorems 2 and 3 and their implications for steady-state estimation may be brought to bear on such chains once the regeneration times are identified.

Identifying the regeneration times is more complicated than in the special single-state case. There are basically two ways of doing this, both of which are based on Definition 4. Let $P(x, \cdot) \stackrel{\triangle}{=} P(X_1 \in \cdot | X_0 = x)$. We may write

$$P(x, \cdot) = \lambda(x)\varphi(\cdot) + (1 - \lambda(x))Q(x, \cdot), \tag{5}$$

where

$$Q(x, \cdot) = \frac{P(x, \cdot) - \lambda(x)\varphi(\cdot)}{1 - \lambda(x)}$$

if $\lambda(x) < 1$, and (arbitrarily) a point mass at x, if not. The decomposition (5) suggests that we might generate a transition from $X_0 = x$ by first generating a Bernoulli random variable Z with $P(Z = 1) = \lambda(x)$. If $Z = 1$, then X_1 is generated according to the distribution φ, and otherwise it is generated according to $Q(x, \cdot)$. The point here is that if $Z = 1$, then X_1 has distribution φ *independently* of $X_0 = x$, and so a regeneration occurs. Generating random variables from the distributions φ and $Q(x, \cdot)$ may prove somewhat difficult, and fortunately a second approach is possible that is based on acceptance/rejection ideas.

Suppose that $X_1 = y$ has already been generated from $X_0 = x$, and we want to determine whether a regeneration occurred at time 1. We can now generate a Bernoulli random variable Z which indicates whether X_1 "came from φ" or not. Intuitively speaking, from acceptance/rejection ideas, Z should have success probability

$$w(x, y) = \frac{\lambda(x)\varphi(dy)}{P(x, dy)},$$

that is, $w(x, \cdot)$ should be a density of $\lambda(x)\varphi(\cdot)$ with respect to $P(x, \cdot)$. This expectation is correct [Glynn and L'Ecuyer 1995]. These two methods for generating regenerations are easily seen to be statistically equivalent.

These ideas were applied in the setting of positive recurrent discrete-time Markov chains on a discrete state-space in Andradóttir et al. [1995]. For such chains, it is well known that returns to a fixed state constitute regeneration times. Andradóttir et al. showed that with an appropriate choice of the function λ and distribution φ, regeneration times defined as above form a "super-sequence" of those defined by returns to a fixed state, and variance reduction in the estimation of the time-average variance is therefore guaranteed. A numerical example in Henderson and Glynn [1999b] shows that the reductions can be quite substantial.

There is a partial converse to the result that chains with unit smoothness index are classically regenerative. First we need a definition.

Definition 7. Let $X = (X_n : n \geq 0)$ be a GSSMC on state space Σ, equipped with σ-field $\mathcal{S}$. Let $\phi : \mathcal{S} \to \mathbb{R}_+$ be a nontrivial measure on $\mathcal{S}$. We say that X is *ϕ-irreducible* [Meyn and Tweedie 1993] if for every $A \in \mathcal{S}$ such that $\phi(A) > 0$, and for every $x \in \Sigma$, there is an $n = n(A, x) \geq 1$ such that $P^n(x, A) > 0$.

PROPOSITION 7. *Suppose that the GSSMC X is classically regenerative and ϕ-irreducible. Then the decomposition (5) holds, where φ is the distribution of the chain at regeneration times.*

For a proof, see Nummelin [1984, Theorem 4.3, p. 66].

In the next section, we will show how to obtain regeneration times for chains with nonunit smoothness index. Unfortunately, the resulting regenerative cycles are 1-dependent (non-adjacent cycles are independent, while adjacent cycles may be dependent). Proposition 7 implies that for such chains, *there is no way to obtain independent and identically distributed cycles*, since otherwise (5) would hold and the chain would have unit smoothness index.

4. REGENERATIVE SIMULATION WITH NONUNIT SMOOTHNESS INDEX

When the smoothness index $m > 1$, it is still possible to define regeneration times $(T(k) : k \geq 0)$ for X, but the regenerative cycles are now 1-dependent. (We explain below how this 1-dependence arises.) This weakened form of regeneration still allows one to analyze steady-state simulations, but the details are more complicated than in the unit smoothness case. Theorems 2 and 3 generalize to yield the following result.

THEOREM 8. *Let $X = (X_n : n \geq 0)$ be a regenerative process with 1-dependent cycles and regeneration times $(T(k) : k \geq 0)$. Let Y_i, τ_i be defined as in (2) for $i \geq 1$, let $\bar{Y}_n$ and $\bar{\tau}_n$ be sample means as defined earlier, and let $\alpha_n = \bar{Y}_n/\bar{\tau}_n$. Suppose that $\mathbb{E}_\varphi(\tau_1 + |Y_1|) < \infty$.*

(i) *The strong law*

$$\alpha_n \to \alpha = \frac{\mathbb{E}_\varphi Y_1}{\mathbb{E}_\varphi \tau_1}$$

holds almost surely.

Define $Z_i = Y_i - \alpha\tau_i$ for $i \geq 1$, and suppose that in addition to the above conditions, $0 < \mathbb{E}_\varphi Z_1^2 < \infty$.

(ii) *The CLT*

$$n^{1/2}(\alpha_n - \alpha) \Rightarrow \sigma_{\text{cyc}} N(0, 1)$$

holds, where $\sigma_{\text{cyc}}^2 = \left(\mathbb{E}_\varphi Z_1^2 + 2\,\mathbb{E}_\varphi Z_1 Z_2\right)/(\mathbb{E}_\varphi \tau_1)^2$.

(iii) *The estimator*

$$v_n^2 \triangleq \frac{n^{-1}\sum_{i=1}^{n-1}[(Z_i(n))^2 + 2Z_i(n)Z_{i+1}(n)]}{(\bar{\tau}_n)^2}$$

of σ_{cyc}^2 *is weakly consistent, that is,* $v_n^2 \Rightarrow \sigma_{\text{cyc}}^2$ *as* $n \to \infty$*, where* $Z_i(n) = Y_i - \alpha_n \tau_i$.

(iv) *If, in addition,* $\mathbb{E}_\varphi(Y_1^4 + \tau_1^4) < \infty$*, then* $n^{1/2}(\alpha_n - \alpha, v_n - \sigma_{\text{cyc}}) \Rightarrow N(0, \Lambda)$*, where*

$$\begin{aligned}
\Lambda_{11} &= \frac{\mathbb{E}_\varphi\left[Z_1^2 + 2Z_1Z_2\right]}{(\mathbb{E}_\varphi\tau_1)^2},\\
\Lambda_{12} &= \frac{\mathbb{E}_\varphi[Z_1D_1 + Z_1D_2 + D_1Z_2 + D_1Z_3]}{2\sigma_{\text{cyc}}(\mathbb{E}_\varphi\tau_1)^3} = \Lambda_{21},\\
\Lambda_{22} &= \frac{\mathbb{E}_\varphi\left[D_1^2 + 2D_1D_2 + 2D_1D_3\right]}{4\sigma_{\text{cyc}}^2(\mathbb{E}_\varphi\tau_1)^4},\\
\beta &= 2\,\mathbb{E}_\varphi[\tau_1Z_1 + \tau_1Z_2 + Z_2\tau_1]/\,\mathbb{E}_\varphi\tau_1, \quad \textit{and}\\
D_i &= Z_i^2 - \mathbb{E}_\varphi Z_1^2 + 2Z_iZ_{i+1} - 2\,\mathbb{E}_\varphi Z_1Z_2 - 2(\mathbb{E}_\varphi\tau_1)\sigma_{\text{cyc}}^2(\tau_i - \mathbb{E}_\varphi\tau_1) - \beta Z_i.
\end{aligned}$$

Glynn [1982] proves (i), and versions of results (ii) and (iii) on the natural (simulated) time scale under stronger moment conditions than we require. The proofs of (ii) and (iii) follow as in Glynn and Iglehart [1993] and are omitted. The proof of result (iv) may be found in Section 6.

Our next result discusses the bias properties of the estimator α_n, generalizing Theorem 4.

THEOREM 9. *Suppose that X is a GSMP with nonunit smoothness index. Let* $(Y_i : i \geq 1)$, $(\tau_i : i \geq 1)$*, and* α_n *be defined as in Section 1. If* $|f|$ *is bounded and* $\mathbb{E}_\varphi\tau_1^4 < \infty$*, then*

$$\mathbb{E}\,\alpha_n = \alpha - bn^{-1} + o(n^{-1}),$$

where

$$b = \frac{\mathbb{E}_\varphi Z_1\tau_1 + \mathbb{E}_\varphi Z_1\tau_2 + \mathbb{E}_\varphi Z_2\tau_1}{(\mathbb{E}_\varphi\tau_1)^2}.$$

As discussed in Section 1, bias will only be significant for simulation runs containing few cycles. In such situations, one could use

$$\alpha_n' = \alpha_n + \frac{1}{n}\frac{n^{-1}\sum_{i=1}^{n-1} Z_i(n)\tau_i + Z_i(n)\tau_{i+1} + Z_{i+1}(n)\tau_i}{(\bar{\tau}_n)^2}$$

to estimate α, where $Z_i(n) = Y_i - \alpha_n\tau_i$. As in the independent cycle case, we believe that a different bias-reduction technique is typically more effective than the use of α_n', and so we do not offer (involved) conditions for when α_n' provably reduces bias. We refer the reader to Glynn and Heidelberger [1990] for further discussion of this point.

As in the independent cycle case, one can use an estimator $\alpha(t)$ based on the number of regenerative cycles completed by simulation time t. In particular, if $N(t)$ is defined as in Section 1 as the number of identically distributed regenerative cycles completed by simulation time t, then we may estimate α by

$$\alpha(t) = \frac{\sum_{i=1}^{N(t)} Y_i}{\sum_{i=1}^{N(t)} \tau_i}.$$

Since independence is a special case of 1-dependence, it follows that the bias of this estimator is of the order t^{-1}. We now extend a result by Meketon and Heidelberger [1982] to show that this bias can be reduced to $o(t^{-1})$. Because of the 1-dependent structure, one must not only complete the cycle in progress at time t, but also the following cycle. Define

$$\alpha'(t) = \frac{\sum_{i=1}^{N(t)+2} Y_i}{\sum_{i=1}^{N(t)+2} \tau_i}.$$

THEOREM 10. *Suppose that X is a GSMP with unit or nonunit smoothness index, and let Y_i and τ_i be defined as in Section 1. If $\mathbb{E}_\varphi(Y_1^2 + \tau_1^2) < \infty$, then the bias of $\alpha'(t)$ is $o(t^{-1})$.*

Remark 5. An examination of the proof of Theorem 10 shows that under stronger moment conditions on the cycle length distribution, the bias can be shown to be $O(t^{-3/2})$. This follows by using renewal theory arguments to bound the right-hand side of (26) by a constant that is independent of t.

The regeneration times $(T(k) : k \geq 0)$ may be determined using methods analogous to those for chains with unit smoothness index [Glynn and L'Ecuyer 1993]. Defining $P^m(x, \cdot) = P(X_m \in \cdot | X_0 = x)$, we have

$$P^m(x, \cdot) = \lambda(x)\varphi(\cdot) + (1 - \lambda(x))Q(x, \cdot)$$

with a suitable definition of the (m-step) transition kernel Q. Given X_0, one can generate the following m transitions as follows. First generate a Bernoulli random variable Z with $P(Z = 1) = \lambda(x)$. If $Z = 1$, then X_m is generated from φ, otherwise it is generated from Q. The intermediate values $X_1, \ldots, X_{m-1}$ are then generated from the appropriate conditional distributions. Because of the difficulty in generating random variables from these distributions, this approach is rarely implementable.

A second approach parallels the acceptance/rejection method for the unit smoothness index case. First $X_0 = x, \ldots, X_m = y$ are generated in any convenient manner. Then one generates a Bernoulli random variable Z having success probability $w(x, y) = \lambda(x)\varphi(dy)/P^m(x, dy)$. If $Z = 1$, then a regeneration is recorded at time m, and X_m is distributed according to φ independent of the value of X_0. However, there will almost certainly be correlation between X_m and $(X_1, \ldots, X_{m-1})$. Therefore, the resulting cycles are no longer independent, but are 1-dependent. One might then simulate a further m transitions and repeat this process, or perhaps wait until the chain enters some favourable region before repeating the test for regeneration (see Andradóttir et al. [1995] for more details).

Both methods outlined above require knowledge of the transition kernel P^m to identify cycle boundaries. This quantity is unlikely to be readily available to the simulationist except in problems with very special structure. This is the essence of the difficulty in applying regenerative simulation to simulation of chains with nonunit smoothness index.

Glynn [1994] suggests a slightly different approach to regenerative simulation in the nonunit smoothness index context. After a regeneration is detected at time n say, a new value of X_n is independently sampled from φ. The resulting process X^* say, has independent cycles, and identical marginals and steady-state as X. However, one still needs to identify the cycle boundaries, and so simulation of X^* still requires explicit knowledge of P^m. The principal problem has not been avoided.

These negative comments should be set against a recent positive result of Henderson and Glynn [1999a] that shows that it *is* possible to explicitly compute $P^m(x, dy)$, for a restricted set of values x and y. Henderson and Glynn explain how this observation can be used to identify regeneration times using the second method described above. The approach is potentially useful for small systems, but for models of even moderate complexity, it appears that the regenerative cycles constructed using this method will be excessively long.

In view of these implementation difficulties, characterizing GSMPs with unit smoothness index becomes of great interest. To the extent possible, we will treat the case $|S| = \infty$ (a countably infinite state space) as well as the more tractable case $|S| < \infty$.

The next result gives conditions under which the GSSMC associated with a GSMP is ϕ-irreducible for a certain measure ϕ. The proof of this result is constructive, and borrows ideas from an analogous result for finite state spaces due to Haas [1999]. Let ϕ_∞ be defined by

$$\phi_\infty(A) = \prod_{e \in E(s)} a_e, \tag{6}$$

where

$$A = \{s\} \times \{c \in R_s : 0 \le c_e \le a_e \quad \forall e \in E(s)\},$$

that is, ϕ_∞ is basically counting measure on the states together with Lebesgue measure on the clocks. (Equation (6) defines the measure on a class of sets that is a π-system. Theorem 3.3 of Billingsley [1986] allows us to conclude that this is sufficient to define ϕ_∞ uniquely.) Similarly, for $u > 0$, define

$$\phi_u(A) = \prod_{e \in E(s)} \min(a_e, u),$$

so that ϕ_u is the same measure confined to clock readings bounded by u.

THEOREM 11. *Let $\tilde{X}$ be a GSMP with related GSSMC X. Suppose that*

(1) *the GSMP $\tilde{X}$ is irreducible (Definition 5),*

(2) *for some $\epsilon > 0$ all clock setting distributions have density components that are positive almost everywhere (a.e.) on $[0, \epsilon)$, and*

(3) *all active clock speeds are 1.*

Then X is ϕ_u-irreducible, for any $u \in (0, \epsilon)$, and therefore ϕ_ϵ-irreducible.

The assumption that all active clock speeds equal 1 may appear somewhat restrictive. However, we expect that a similar result also holds when active clock speeds are restricted to lie in some interval $[r_*, r^*]$, where $r_* > 0$.

Theorem 11 immediately gives the following corollary.

COROLLARY 12. *If Condition 2 of Theorem 11 is strengthened to all clock setting distributions have density components that are positive almost everywhere on* $[0, \infty)$*, then* X *is* ϕ_∞*-irreducible.*

Irreducibility implies what we believe is the first uniqueness result for invariant measures of GSMPs with countably infinite state space. Meyn and Tweedie [1993] define *positive* chains to be those chains that are ϕ-irreducible and that possess an invariant probability measure. By Proposition 10.1.1, page 231, a positive chain is recurrent, and by Theorem 10.4.4, page 242, recurrent chains possess a unique invariant measure. We have established the following corollary.

COROLLARY 13. *Under the conditions of Theorem 11, if* X *has an invariant probability measure, then it is unique.*

The following definition will prove useful in characterising the value of the smoothness index.

Definition 8. Let X be the GSSMC associated with a GSMP. Suppose that X has a stationary probability distribution π. We say that π satisfies the *clock smoothness condition* if π is absolutely continuous with respect to $\phi_\infty(\pi \ll \phi_\infty)$.

We now wish to establish simple sufficient conditions for the clock smoothness condition to hold. Our proof is based on the following rather general result, which is interesting in its own right.

THEOREM 14. *Let* $X = (X_n : n \geq 0)$ *be a positive Harris recurrent Markov chain on state space* Σ *equipped with* σ*−field* $\mathcal{S}$*. Let* π *be the stationary probability distribution and* P *be the transition kernel of* X*. Let* ψ *be a given measure on* $(\Sigma, \mathcal{S})$*. Suppose that there is a stopping time* T *(with respect to the natural filtration on* X*) such that*

(i) $T < \infty$ P_x *almost surely, and*
(ii) $P_x(X_T \in \cdot) \ll \psi$*.*

Finally, suppose that P *has the property that if a probability measure* $\nu \ll \psi$*, then* $\nu P \ll \psi$*, where* νP *is the probability measure defined by* $\nu P(A) = \int_\Sigma \nu(dx)P(x, A)$ *for* $A \in \mathcal{S}$*, i.e., the distribution of* X_1 *when* X_0 *has distribution* ν*. Then* $\pi \ll \psi$*.*

This theorem basically states that if the transition kernel satisfies a certain "smoothing property," then once the distribution of the chain is "smooth," it will remain "smooth," and so the stationary distribution of the chain will be "smooth." We will also need the following result, which gives conditions under which a GSMP is nonexplosive.

THEOREM 15. *Let X be the GSSMC associated with a GSMP $\tilde{X}$. Suppose that all clock setting distributions are non-null, in the sense that no clock setting distributions correspond to a point mass at 0. If X is Harris recurrent, then the GSMP is nonexplosive.*

We are now in a position to give sufficient conditions for the clock smoothness condition to hold. Our conditions are far from being the "tightest" possible, but they demonstrate the type of result one might expect to hold.

THEOREM 16. *Let X be the GSSMC associated with a GSMP where all clock-setting distributions have densities (with respect to Lebesgue measure). If X is positive Harris recurrent with stationary probability distribution π, then π satisfies the clock smoothness condition.*

We are now able to quantify the smoothness index.

THEOREM 17. *Let $X = (X_n : n \geq 0)$ be a positive Harris recurrent GSSMC associated with a GSMP, and let π be the stationary distribution of X. Assume that all clocks remain active until they trigger a state transition, and that the triggering event set consists of a single event almost surely. Suppose that π satisfies the clock smoothness condition. Set*

$$m^* = \min_{s \in S : \pi_s > 0} |E(s)|,$$

where $\pi_s = P_\pi(S_0 = s)$. If m is the smoothness index for X, then $m \geq m^*$.

Remark 6. The fundamental idea behind this theorem is that $P^m(x, \cdot)$ is singular with respect to π for $m < m^*$.

Remark 7. The class of GSMPs satisfying the conditions of Theorem 17 is very large. For example, insensitive GSMPs [Burman 1981] have this form. These GSMPs are such that π_s depends on the clock random variable distributions only through their moments, so that they are, in some sense, insensitive to the clock random variable distributions.

Theorem 17 basically establishes that the smoothness index m must be at least as large as m^*, the minimum number of active clocks at any time. The following proposition provides a partial converse, establishing conditions under which the GSSMC has an m^*-minorization.

PROPOSITION 18. *Suppose that*

(1) *X is the GSSMC associated with an irreducible GSMP,*
(2) *all clock setting distributions have density components that are positive and bounded away from 0 (by a common lower bound) on $[0, \epsilon]$, for some $\epsilon > 0$,*
(3) *all active clock speeds are 1, and*
(4) *all clocks remain active until they trigger a state transition.*

Then there exists a set A with $\phi_\epsilon(A) > 0$, a $\lambda > 0$, and a probability measure φ on $(\Sigma, \mathcal{S})$ such that

$$P^{m^*}(x, \cdot) \geq \lambda \varphi(\cdot)$$

for all $x \in A$, where $m^* = \min_{s \in S} |E(s)|$.

5. CONCLUSIONS

Theorem 5 establishes that any "well-posed" simulation is necessarily positive Harris recurrent, and therefore possesses regenerative structure. The question then is whether the regenerative structure can be identified. When the chain possesses a single state, this is easy, because the exit times from single states are regeneration times. But such chains appear to form a very special class of all discrete-event systems, so that a more general method for detecting regenerations is warranted.

We described methods based on minorizations of the n-step transition kernel, and defined the smoothness index m to be the minimum value n such that an appropriate minorization can be constructed. The special case of a unit smoothness index ($m = 1$) is of great interest, because it is quite straightforward to compute the 1-step transition kernel; see (1). In Theorem 17 and Proposition 18 we basically established that in the absence of event cancellation, the only chains that have unit smoothness index are those with single states. Therefore, if these minorization methods are to be used to detect regenerative structure in chains without single states, then we have to deal with the $m > 1$ case.

The problem is that to detect regenerative structure in the nonunit smoothness index ($m > 1$) case, the methods that we have outlined require explicit knowledge of the m-step transition kernel. At least currently, it is unlikely that such information will be available. Henderson and Glynn [1999a] describe an exception that may prove useful in small systems, but the cycles that can be constructed using their method are likely to be excessively long in models of moderate complexity.

Indeed, it is well known that excessively long regenerative cycles are a practical barrier to implementation of the regenerative method. Long cycles can result from the construction based on minorizations if the function λ appearing in Definition 4 is such that $\pi\lambda \stackrel{\triangle}{=} \int_\Sigma \lambda(x)\pi(dx)$ is very small, where π is the stationary probability distribution corresponding to the positive Harris recurrent Markov chain X. However, some hope comes from the observation that for a positive Harris recurrent aperiodic Markov chain, $P^n(x, \cdot) \Rightarrow \pi(\cdot)$ as $n \to \infty$ for all x (Meyn and Tweedie [1993, Theorem 13.3.3]). Hence, for n sufficiently large, it may be the case that $P^n(x, \cdot)$ can be more strongly minorized, leading to a larger function λ and, in turn, shorter regenerative cycles.

Another interesting question is whether it is possible to identify regenerative structure in the GSMP $\tilde{X}$ *without having to compute* $P^m(x, \cdot)$. Such a method would allow the very appealing statistical properties of the regenerative method to be brought to bear on the general problem of discrete-event simulation.

We continue to work on these and other interesting possibilities with the ultimate goal of obtaining a practical method for applying the regenerative method to general steady-state discrete-event simulation.

6. PROOFS

Let $o_p(g(n))$ represent the nth term of a stochastic sequence ζ_n say, with the property that $\zeta_n/g(n) \Rightarrow 0$ as $n \to \infty$.

We will need the following lemma in proving Theorem 8.

LEMMA 19. *Let* $(\bar{V}_n : n \geq 1)$ *be a sequence of random variables.*

(1) *If* $\bar{V}_n \Rightarrow v$ *for some finite constant* v*, then* $\bar{V}_n = v + o_p(1)$.
(2) *If* $n^{1/2}(\bar{V}_n - v)$ *converges in distribution, then* $\bar{V}_n = v + o_p(n^{-1/2+\epsilon})$ *for all* $\epsilon > 0$.
(3) *Suppose that* $\bar{V}_n \to v$ *almost surely as* $n \to \infty$ *and* $n^{1/2}(\bar{V}_n - v)$ *converges in distribution. If* g *is a real-valued function that is continuously differentiable in a neighbourhood,* Γ *say, of* v*, then*

$$g(\bar{V}_n) = g(v) + g'(v)(\bar{V}_n - v) + o_p(n^{-1/2}).$$

PROOF. The proofs of Parts 1 and 2 are straightforward and omitted. For Part 3, note that

$$g(\bar{V}_n) = g(\bar{V}_n)I(\bar{V}_n \in \Gamma) + g(\bar{V}_n)I(\bar{V}_n \notin \Gamma). \tag{7}$$

The second term in (7) is $o_p(n^{-1/2})$ because $I(\bar{V}_n \notin \Gamma)$ is 0 for n sufficiently large with probability 1. By Taylor's theorem, there is some ζ_n between v and $\bar{V}_n$ such that

$$\begin{aligned} g(\bar{V}_n)I(\bar{V}_n \in \Gamma) &= [g(v) + g'(v)(\bar{V}_n - v) + (g'(\zeta_n) - g'(v))(\bar{V}_n - v)]I(\bar{V}_n \in \Gamma) \\ &= [g(v) + g'(v)(\bar{V}_n - v) + o_p(n^{-1/2})]I(\bar{V}_n \in \Gamma) \\ &= [g(v) + g'(v)(\bar{V}_n - v) + o_p(n^{-1/2})][1 - I(\bar{V}_n \notin \Gamma)] \qquad (8) \end{aligned}$$

where (8) follows because $g'(\zeta_n) \to g'(v)$ with probability 1. The result then follows because $I(\bar{V}_n \notin \Gamma) = o_p(n^{-1/2})$. □

PROOF. (Theorem 8) We only need to prove part (iv). The proof basically follows from certain Taylor expansions. For notational convenience we write $\mathbb{E}(\cdot)$ for $\mathbb{E}_\varphi(\cdot)$ and $\bar{\tau}, \bar{Y}$ for $\mathbb{E}\,\tau_1$ and $\mathbb{E}\,Y_1$. Observe that

$$\begin{aligned} v_n^2 &= \left(n\bar{\tau}_n^2\right)^{-1} \sum_{i=1}^{n-1} [Z_i(n)^2 + 2Z_i(n)Z_{i+1}(n)] \\ &= \left(n\bar{\tau}_n^2\right)^{-1} \sum_{i=1}^{n-1} [(Z_i - (\alpha_n - \alpha)\tau_i)^2 + 2(Z_i - (\alpha_n - \alpha)\tau_i)(Z_{i+1} - (\alpha_n - \alpha)\tau_{i+1})] \\ &= \left(n\bar{\tau}_n^2\right)^{-1} \sum_{i=1}^{n-1} \left[Z_i^2 + 2Z_iZ_{i+1} - 2(\alpha_n - \alpha)(\tau_i Z_i + \tau_i Z_{i+1} + \tau_{i+1}Z_i) \right. \\ &\quad \left. + (\alpha_n - \alpha)^2\left(\tau_i^2 + 2\tau_i\tau_{i+1}\right)\right]. \qquad (9) \end{aligned}$$

From Part 2 of Lemma 19, $\alpha_n - \alpha = o_p(n^{-1/2+\epsilon})$ for all $\epsilon > 0$. Hence, $(\alpha_n - \alpha)^2 = o_p(n^{-1+2\epsilon})$ for all $\epsilon > 0$. Furthermore, the strong law of large numbers ensures that

$$\left(n\bar{\tau}_n^2\right)^{-1} \sum_{i=1}^{n-1} \tau_i^2 + 2\tau_i\tau_{i+1} \to \frac{\mathbb{E}\,\tau_1^2 + 2\,\mathbb{E}\,\tau_1\tau_2}{\bar{\tau}^2} \text{ almost surely.}$$

Hence,

$$\left(n\bar{\tau}_n^2\right)^{-1}\sum_{i=1}^{n-1}(\alpha_n-\alpha)^2\left(\tau_i^2+2\tau_i\tau_{i+1}\right)=o_p\left(n^{-1/2}\right). \tag{10}$$

Now, for $i \geq 1$ let $R_i = \tau_i Z_i + \tau_i Z_{i+1} + \tau_{i+1} Z_i$. From Part 1 of Lemma 19

$$\begin{aligned} 2(\alpha_n-\alpha)\frac{1}{\bar{\tau}_n^2}\frac{1}{n}\sum_{i=1}^{n-1}R_i &= 2(\alpha_n-\alpha)\left(\frac{1}{\bar{\tau}^2}+o_p(1)\right) \\ &\quad\times\left(\frac{1}{n}\sum_{i=1}^{n-1}(R_i-\mathbb{E}\,R_1)+\mathbb{E}\,R_1+o_p(1)\right) \\ &= 2(\alpha_n-\alpha)\left(\frac{1}{\bar{\tau}^2}+o_p(1)\right)(\mathbb{E}\,R_1+o_p(1)) \\ &= \frac{2\,\mathbb{E}\,R_1}{\bar{\tau}^2}(\alpha_n-\alpha)+o_p\left(n^{-1/2}\right) \end{aligned} \tag{11}$$

where the second equality follows since $n^{-1}\sum_{i=1}^{n-1}(R_i-\mathbb{E}\,R_i)\to 0$ almost surely and (11) follows since $(\alpha_n-\alpha)o_p(1)=o_p(n^{-1/2})$.

If we define $\beta = 2\,\mathbb{E}\,R_1/\bar{\tau}$ then, combining (9), (10) and (11), we find that

$$v_n^2=\left(n\bar{\tau}_n^2\right)^{-1}\sum_{i=1}^{n-1}\left[Z_i^2+2Z_iZ_{i+1}\right]-\frac{\beta}{\bar{\tau}}(\alpha_n-\alpha)+o_p\left(n^{-1/2}\right). \tag{12}$$

Now, $\alpha_n-\alpha=\bar{Y}_n/\bar{\tau}_n-\bar{Y}/\bar{\tau}$. By Part 3 of Lemma 19 with $g(x)=1/x$, $1/\bar{\tau}_n = 1/\bar{\tau}-(\bar{\tau}_n-\bar{\tau})/\bar{\tau}^2+o_p\left(n^{-1/2}\right)$. Hence,

$$\begin{aligned} \alpha_n-\alpha &= \bar{Y}_n/\bar{\tau}-\bar{Y}_n(\bar{\tau}_n-\bar{\tau})/\bar{\tau}^2+o_p\left(n^{-1/2}\right)-\alpha \\ &= \bar{Y}_n/\bar{\tau}-\bar{Y}(\bar{\tau}_n-\bar{\tau})/\bar{\tau}^2-\alpha+o_p\left(n^{-1/2}\right) && (13)\\ &= \bar{Y}_n/\bar{\tau}-\alpha\bar{\tau}_n/\bar{\tau}+o_p\left(n^{-1/2}\right) && (14)\\ &= \bar{Z}_n/\bar{\tau}+o_p\left(n^{-1/2}\right). && (15) \end{aligned}$$

The equality (13) follows since $\bar{Y}_n=\bar{Y}+o(1)$ and $n^{1/2}(\bar{\tau}_n-\bar{\tau})$ converges in distribution. Equations (14) and (15) hold because $\alpha=\bar{Y}/\bar{\tau}$ and $\bar{Z}_n=\bar{Y}_n-\alpha\bar{\tau}_n$.

Furthermore, by Part 3 of Lemma 19 with $g(x)=x^{-2}$,

$$\bar{\tau}_n^{-2}=\bar{\tau}^{-2}-2\bar{\tau}^{-3}(\bar{\tau}_n-\bar{\tau})+o_p\left(n^{-1/2}\right),$$

so that $(n\bar{\tau}_n^2)^{-1}\sum_{i=1}^{n-1}[Z_i^2+2Z_iZ_{i+1}]$ is given by

$$\begin{aligned} &\frac{1}{n\bar{\tau}^2}\sum_{i=1}^{n-1}\left[Z_i^2+2Z_iZ_{i+1}\right]-\frac{2}{\bar{\tau}^3}(\bar{\tau}_n-\bar{\tau})\frac{1}{n}\sum_{i=1}^{n}\left[Z_i^2+2Z_iZ_{i+1}\right]+o_p\left(n^{-1/2}\right) \\ &=\frac{1}{n\bar{\tau}^2}\sum_{i=1}^{n-1}\left[Z_i^2+2Z_iZ_{i+1}\right]-\frac{2}{\bar{\tau}^3}(\bar{\tau}_n-\bar{\tau})\mathbb{E}\left(Z_1^2+2Z_1Z_2\right)+o_p\left(n^{-1/2}\right) && (16)\\ &=\frac{1}{n}\sum_{i=1}^{n}\frac{Z_i^2+2Z_iZ_{i+1}-2\bar{\tau}\sigma_{\mathrm{cyc}}^2(\tau_i-\bar{\tau})}{\bar{\tau}^2}+o_p\left(n^{-1/2}\right), && (17) \end{aligned}$$

where (16) follows from Part 1 of Lemma 19 and the fact that $(\bar{\tau}_n - \bar{\tau})o_p(1) = o_p(n^{-1/2})$, and (17) holds since $\sigma^2_{\text{cyc}} = \mathbb{E}(Z_1^2 + 2Z_1Z_2)/\bar{\tau}^2$.

Combining (12), (15) and (17), we obtain that

$$\begin{aligned} v_n^2 - \sigma^2_{\text{cyc}} &= (n\bar{\tau}^2)^{-1}\sum_{i=1}^n \left[Z_i^2 + 2Z_iZ_{i+1} - 2\bar{\tau}\sigma^2_{\text{cyc}}(\tau_i - \bar{\tau}) - \beta Z_i\right] - \sigma^2_{\text{cyc}} + o_p\left(n^{-1/2}\right) \\ &= n^{-1}\sum_{i=1}^n \bar{\tau}^{-2}\left[Z_i^2 - \mathbb{E}\,Z_1^2 + 2Z_iZ_{i+1} - 2\,\mathbb{E}\,Z_1Z_2 - 2\bar{\tau}\sigma^2_{\text{cyc}}(\tau_i - \bar{\tau}) - \beta Z_i\right] \\ &\quad + o_p\left(n^{-1/2}\right) \\ &= n^{-1}\sum_{i=1}^n \frac{D_i}{\bar{\tau}^2} + o_p\left(n^{-1/2}\right), \end{aligned} \tag{18}$$

where $D_i = Z_i^2 - \mathbb{E}\,Z_1^2 + 2Z_iZ_{i+1} - 2\,\mathbb{E}\,Z_1Z_2 - 2\bar{\tau}\sigma^2_{\text{cyc}}(\tau_i - \bar{\tau}) - \beta Z_i$.

Noting that $v_n = \sqrt{v_n^2}$, Part 3 of Lemma 19 with $g(x) = \sqrt{x}$ gives

$$v_n - \sigma_{\text{cyc}} = (2\sigma_{\text{cyc}})^{-1}\left(v_n^2 - \sigma^2_{\text{cyc}}\right) + o_p\left(n^{-1/2}\right). \tag{19}$$

From (15), (18) and (19), we then have that

$$\begin{pmatrix} \alpha_n - \alpha \\ v_n - \sigma_{\text{cyc}} \end{pmatrix} = \begin{pmatrix} \bar{\tau}^{-1}\bar{Z}_n \\ (2\sigma_{\text{cyc}}\bar{\tau}^2)^{-1}\bar{D}_n \end{pmatrix} + o_p\left(n^{-1/2}\right), \tag{20}$$

where $\bar{D}_n$ is the sample mean of $D_1, \ldots, D_n$. Now, $(Z_i : i \geq 1)$ is a 1-dependent sequence of identically distributed random variables and $(D_i : i \geq 1)$ is a 2-dependent sequence of identically distributed random variables. Hence, if L_i is the column vector $(Z_i/\bar{\tau}, D_i/(2\sigma_{\text{cyc}}\bar{\tau}^2))'$ and x' denotes the transpose of the vector x, then $(L_i : i \geq 1)$ is a 2-dependent sequence of identically distributed random vectors whose components have finite second moment and zero mean. Application of a central limit theorem for such sequences (see Billingsley [1968, p. 177]) gives the result. The expressions for the covariance matrix may be obtained by noting that $\Lambda = E(L_1L_1' + L_1L_2' + L_2L_1' + L_1L_3' + L_3L_1')$ and using the fact that $EZ_1D_3 = 0$ to simplify the resulting expressions. □

Proof. (Theorem 9) We follow the proof of Theorem 7 of Glynn and Heidelberger [1990] closely. Since $|f|$ is bounded (by A say), $\alpha_n, \alpha \leq A$ almost surely. Let $X_i = (Y_i, \tau_i)'$, and $\bar{X}_n = (\bar{Y}_n, \bar{\tau}_n)'$, where $\bar{Y}_n$ and $\bar{\tau}_n$ are the sample means of $(Y_1, \ldots, Y_n)$ and $(\tau_1, \ldots, \tau_n)$ respectively. Define $g(y, z) = y/z$, so that $\alpha_n = g(\bar{X}_n)$ and $\alpha = g(\mu)$, where $\mu = (\mathbb{E}\,Y, \mathbb{E}\,\tau)'$. (For notational convenience, we write Y, τ for Y_1, τ_1). Observe that g is C^∞ on $[0, \infty) \times (0, \infty)$. Define Γ_n to be the indicator of the event $\{\|\bar{X}_n - \mu\| \leq \epsilon\}$, where $\|x\| = \max_i |x_i|$ and ϵ is positive but strictly smaller than the minimum of $\mathbb{E}\,\tau_1$ and $\mathbb{E}\,|Y_1|$. Then

$$\mathbb{E}\,\alpha_n = \alpha + \mathbb{E}((\alpha_n - \alpha)\Gamma_n) + \mathbb{E}((\alpha_n - \alpha)(1 - \Gamma_n)).$$

Now

$$\begin{aligned}
&\mathbb{E}((\alpha_n - \alpha)(1 - \Gamma_n)) \\
&\le 2AP(\|\bar{X}_n - \mu\| > \epsilon) \\
&\le 2A\{P(|\bar{Y}_n - \mathbb{E}\,Y| > \epsilon) + P(|\bar{\tau}_n - \mathbb{E}\,\tau| > \epsilon)\} \\
&= 2A\{P(\sqrt{n}|\bar{Y}_n - \mathbb{E}\,Y| > \epsilon\sqrt{n}) + P(\sqrt{n}|\bar{\tau}_n - \mathbb{E}\,\tau| > \epsilon\sqrt{n})\} \\
&\le \frac{2A}{n^{1+\delta}\epsilon^{2+2\delta}}\{\mathbb{E}(\sqrt{n}|\bar{Y}_n - \mathbb{E}\,Y|)^{2+2\delta} + \mathbb{E}(\sqrt{n}|\bar{\tau}_n - \mathbb{E}\,\tau|)^{2+2\delta}\} \qquad (21)
\end{aligned}$$

for any $\delta > 0$. Now, $\mathbb{E}[\sqrt{n}|\bar{Y}_n - \mathbb{E}\,Y|]^4 = n^2\,\mathbb{E}(\bar{Y}_n - \mathbb{E}\,Y)^4$, which, by direct calculation, is bounded in n. Similarly, $n^2\,\mathbb{E}(\bar{\tau}_n - \mathbb{E}\,\tau)^4$ is bounded in n. Hence, for $u \in [0, 4)$, the sequences

$$([\sqrt{n}|\bar{Y}_n - \mathbb{E}\,Y|]^u : n \ge 1) \quad \text{and} \quad ([\sqrt{n}|\bar{\tau}_n - \mathbb{E}\,\tau|]^u : n \ge 1) \qquad (22)$$

are uniformly integrable. Therefore, the expectations in (21) are bounded in n for small δ and $\mathbb{E}(\alpha_n - \alpha)(1 - \Gamma_n) = o(n^{-1})$.

When $\Gamma_n = 1$, a Taylor expansion shows that $(\alpha_n - \alpha)\Gamma_n$ is given by

$$\nabla g(\mu)^T(\bar{X}_n - \mu)\Gamma_n + \frac{1}{2}\sum_{i=1}^{2}\sum_{j=1}^{2}\nabla^2 g(\zeta_n)_{ij}(\bar{X}_n(i) - \mu_i)(\bar{X}_n(j) - \mu_j)\Gamma_n, \qquad (23)$$

where ζ_n lies on the line segment joining $\bar{X}_n$ and μ. Now,

$$\mathbb{E}((\bar{X}_n - \mu)\Gamma_n) = -\,\mathbb{E}((\bar{X}_n - \mu)(1 - \Gamma_n)),$$

and by Hölder's inequality,

$$\begin{aligned}
\mathbb{E}((\bar{Y}_n - \mathbb{E}\,Y)(1 - \Gamma_n)) &\le (\mathbb{E}(\bar{Y}_n - \mathbb{E}\,Y)^4)^{1/4}(\mathbb{E}[1 - \Gamma_n])^{3/4} \\
&\le 2A(\mathbb{E}\,\tau^4)^{1/4}(\mathbb{E}[1 - \Gamma_n])^{3/4} \\
&\le Bn^{-9/8}[\mathbb{E}(\sqrt{n}|\bar{Y}_n - \mathbb{E}\,Y|)^3 + \mathbb{E}(\sqrt{n}|\bar{\tau}_n - \mathbb{E}\,\tau|)^3]^{3/4} \\
&= o(n^{-1}),
\end{aligned}$$

for some constant B, where the second inequality follows from the boundedness of $|f|$, and the final inequality follows as in (21). A similar calculation for $\mathbb{E}(\bar{\tau}_n - \mathbb{E}\,\tau)(1 - \Gamma_n)$ shows that the expectation of the first term in (23) is $o(n^{-1})$.

Now, as $n \to \infty$,

$$n\Gamma_n\nabla^2 g(\zeta_n)_{ij}(\bar{X}_n(i) - \mu_i)\,(\bar{X}_n(j) - \mu_j) \Rightarrow G_{ij}N_iN_j, \qquad (24)$$

where $G = \nabla^2 g(\mu)$, and N_i, N_j are correlated normal random variables with mean zero and $\text{cov}(N_i, N_j) = \text{cov}(X_1(i), X_1(j)) + \text{cov}(X_1(i), X_2(j)) + \text{cov}(X_1(j), X_2(i))$. Note that

$$\mathbb{E}\,\frac{1}{2}\sum_{i=1}^{2}\sum_{j=1}^{2}G_{ij}N_iN_j = -b,$$

where b was defined in the statement of the theorem, and so to complete the proof, it remains to show that the left hand side of (24) is uniformly integrable (as a sequence of random variables in n).

When $\Gamma_n = 1$, $\|\nabla^2 g(\bar{X}_n)\|$ is bounded. Furthermore,

$$\begin{aligned}
&\mathbb{E}\,|n\Gamma_n(\bar{X}_n(i) - \mu_i)(\bar{X}_n(j) - \mu_j)|^{1+\delta} \\
&\quad\le \mathbb{E}\,|\sqrt{n}(\bar{X}_n(i) - \mu_i)|^{1+\delta}|\sqrt{n}(\bar{X}_n(j) - \mu_j)|^{1+\delta} \\
&\quad\le [\mathbb{E}\,|\sqrt{n}(\bar{X}_n(i) - \mu_i)|^{2+2\delta}\,\mathbb{E}\,|\sqrt{n}(\bar{X}_n(j) - \mu_j)|^{2+2\delta}]^{1/2}
\end{aligned}$$

by the Cauchy-Schwarz inequality, and these expectations are bounded in n, as before. This proves uniform integrability and completes the proof. □

PROOF. (Theorem 10) We may write

$$\begin{aligned}
\alpha'(t) - \alpha &= \frac{\sum_{i=1}^{N(t)+2} Y_i - \alpha\tau_i}{\sum_{i=1}^{N(t)+2}\tau_i} \\
&= \frac{S_{N(t)+2}}{t(1 + R_t/t)},
\end{aligned}$$

where

$$S_n = \sum_{i=1}^{n} Z_i \quad \text{and} \quad R_t = \sum_{i=1}^{N(t)+2} \tau_i - t \ge 0.$$

Now, for $x \ge 0$, $(1+x)^{-1} = 1-(1+\zeta)^{-2}x$ for some $\zeta \in [0, x]$ by Taylor's theorem. Thus

$$\alpha'(t) - \alpha = t^{-1} S_{N(t)+2}\left(1 - \frac{1}{(1+\zeta_t)^2}\frac{R_t}{t}\right)$$

for some $\zeta_t \in [0, R_t/t]$. Hence

$$|\,\mathbb{E}\,\alpha'(t) - \alpha| = \left|t^{-1}\,\mathbb{E}\,S_{N(t)+2} - \mathbb{E}\,\frac{R_t}{t^2(1+\zeta_t)^2} S_{N(t)+2}\right|.$$

An extension of Wald's lemma to 1-dependent sequences [Janson 1983] shows that $\mathbb{E}\,S_{N(t)+2} = 0$, since $N(t) + 1$ is a stopping time relative to the filtration $\mathcal{F} = (\mathcal{F}_n : n \ge 1)$, where $\mathcal{F}_n = \sigma(Y_1, \tau_1, \ldots, Y_n, \tau_n)$. Hence

$$\begin{aligned}
|\,\mathbb{E}\,\alpha'(t) - \alpha| &= \left|\mathbb{E}\,\frac{R_t}{t^2(1+\zeta_t)^2} S_{N(t)+2}\right| \\
&\le \frac{1}{t^2}\,\mathbb{E}^{1/2}\left(\frac{R_t}{(1+\zeta_t)^2}\right)^2 \mathbb{E}^{1/2}\, S_{N(t)+2}^2 \\
&\le \frac{1}{t^2}\,\mathbb{E}^{1/2}\, R_t^2\,\mathbb{E}^{1/2}\, S_{N(t)+2}^2 \qquad (25)
\end{aligned}$$

where the first inequality follows by the Cauchy-Schwarz inequality, and the second from the fact that $\zeta_t \ge 0$. We will show that $\mathbb{E}\,R_t^2 = o(t)$, and $\mathbb{E}\,S_{N(t)+2}^2 = O(t)$. This will then establish that the bias in $\alpha'(t)$ is $o(t^{-1})$.

Observe that

$$\mathbb{E}\,R_t^2 \le \mathbb{E}(\tau_{N(t)+1} + \tau_{N(t)+2})^2 \le 2\,\mathbb{E}\left(\tau_{N(t)+1}^2 + \tau_{N(t)+2}^2\right). \qquad (26)$$

Adapting a technique from the proof of Corollary 1.1 of Janson [1983], we have that

$$\mathbb{E}\,\tau_{N(t)+1}^2 \le \mathbb{E} \sum_{i=1}^{N(t)+2} \tau_i^2 = \mathbb{E}(N(t)+2)\,\mathbb{E}\,\tau_1^2.$$

For all $\epsilon > 0$, let $A = A(\epsilon)$ be such that $\mathbb{E}(\tau_1^2; \tau_1 > A) \le \epsilon$. Define $\tau_i' = \tau_i I(\tau_i > A)$ and $\tau_i'' = \tau_i I(\tau_i \le A)$. Then

$$\begin{aligned}\mathbb{E}\,\tau_{N(t)+1}^2 &= \mathbb{E}\left(\tau'_{N(t)+1}\right)^2 + \mathbb{E}\left(\tau''_{N(t)+1}\right)^2 \\ &\le \mathbb{E}(N(t)+2)\,\mathbb{E}(\tau_1')^2 + A^2 \\ &\le \epsilon\,\mathbb{E}\,N(t) + B(\epsilon)\end{aligned}$$

for some constant $B(\epsilon)$. Thus,

$$\limsup_{t\to\infty} t^{-1}\,\mathbb{E}\,\tau_{N(t)+1}^2 \le \epsilon \limsup t^{-1}\,\mathbb{E}\,N(t) = \epsilon / \mathbb{E}\,\tau_1$$

where the equality follows from Theorem 3.1 of Janson [1983]. Since ϵ was arbitrary, we have established that $\mathbb{E}\,\tau_{N(t)+1}^2 = o(t)$. Exactly the same approach may be applied to $\mathbb{E}\,\tau_{N(t)+2}^2$, so that from (26), we obtain $\mathbb{E}\,R_t^2 = o(t)$.

We turn now to the final term in (25), adapting techniques from Janson [1983] and Chow et al. [1965]. Define

$$W_n = \mathbb{E}\left(S_{n+1}^2 - \sum_{i=1}^{n+1} Z_i^2 - 2\sum_{i=1}^{n} Z_i Z_{i+1} \mid \mathcal{F}_n\right),$$

and observe that $W = (W_n : n \ge 1)$ is a martingale with respect to $\mathcal{F}$. Define $a \wedge b = \min\{a, b\}$. If T is any stopping time with respect to $\mathcal{F}$, then $\mathbb{E}\,W_{T\wedge n} = \mathbb{E}\,W_1 = 0$. Hence,

$$\begin{aligned}\mathbb{E}\,S_{1+(T\wedge n)}^2 &= \mathbb{E} \sum_{i=1}^{1+(T\wedge n)} Z_i^2 + 2\sum_{i=1}^{T\wedge n} Z_i Z_{i+1} \\ &\le \mathbb{E} \sum_{i=1}^{1+(T\wedge n)} Z_i^2 + \sum_{i=1}^{T\wedge n}\left(Z_i^2 + Z_{i+1}^2\right) \\ &\le 3\,\mathbb{E} \sum_{i=1}^{1+(T\wedge n)} Z_i^2.\end{aligned}$$

Recalling that $N(t)+1$ is a stopping time, we obtain

$$\begin{aligned}\mathbb{E}\,S_{N(t)+2}^2 &\le \liminf_{n\to\infty} \mathbb{E}\,S_{1+[(N(t)+1)\wedge n]}^2 &(27)\\ &\le 3 \liminf_{n\to\infty} \mathbb{E} \sum_{i=1}^{1+[(N(t)+1)\wedge n]} Z_i^2 \\ &= 3\,\mathbb{E} \sum_{i=1}^{N(t)+2} Z_i^2 &(28)\\ &= 3\,\mathbb{E}\,Z_1^2\,\mathbb{E}(N(t)+2),\end{aligned}$$

where (27) follows from Fatou's lemma, and (28) follows from monotone convergence. Theorem 3.1 of Janson [1983] gives $\mathbb{E}\, N(t) + 2 = O(t)$, so that the same is true of $\mathbb{E}\, S^2_{N(t)+2}$ and the proof is complete. □

Proof. (Theorem 11) Let $u \in (0, \epsilon)$, and let

$$A = \{s^*\} \times [a_1, b_1] \times \cdots \times \left[a_{|E(s^*)|}, b_{|E(s^*)|}\right],$$

where $0 < a_k < b_k < u$ for all k. Let $x \in \Sigma$ be the initial state of the GSSMC. We will describe how to set the clocks on successive transitions from x so that after some number of state transitions, the final state of the GSSMC will lie in the set A, proving ϕ_u-irreducibility.

Let $v \in (0, \epsilon - u)$. Our first step is to reduce all active clock readings to lie in $(0, v)$ if this is not already the case. To do so, repeatedly set every new clock reading to lie in $(v/2, v)$. After at most $|E|$ state transitions, the maximum clock reading must have decreased by at least $v/2$. Repeating this process ensures that all active clock readings will eventually lie in the interval $(0, v)$ at a time that we arbitrarily call the bounding instant. Let the active clock readings at this time be $0 < c_1 < \cdots < c_k < v$.

Suppose that after the bounding instant, every new clock reading is set *after* c_k. Then, after at most k transitions, all of the clocks set in the previous step will either be inactive, or have been reset. Let s_0 be the state of the GSMP after the last of the clocks is reset or made inactive. By the GSMP irreducibility assumption, there exists $e_0, s_1, e_1, \ldots, s_n, e_n$ such that $p(s_1; s_0, e_0)p(s_2; s_1, e_1) \cdots p(s^*; s_n, e_n) > 0$. We will construct this path by carefully setting the new clocks, and also ensuring that the final clock readings lie in the appropriate intervals of the set A.

Let t denote the amount of simulated time that has elapsed after the bounding instant. The idea is that the final event e_n will occur at time $t = \epsilon - u$, and at this time, all clocks will have their appropriate values, so that the GSSMC lies in the set A. Suppose the ith event ($0 \leq i < n$) in the above path has just occurred, and we are setting the clock for event e.

(1) If the event e will not be set again before the GSMP enters the set A, it is not the final triggering event e_n, and $e \in E(s^*)$ and its clock reading is required to lie in $[a_j, b_j]$, then set the clock to lie in the interval $[a_j + \epsilon - u - t, b_j + \epsilon - u - t]$. This will ensure that its reading will lie in the appropriate interval when the GSMP enters the set A at time $\epsilon - u$.

(2) If the event e will not be set again before the GSMP enters A, and it is the final triggering event e_n, then set it equal to $\epsilon - u - t$ (or at least in a very small interval containing this time point).

(3) If the event e will not be set again before the GSMP enters A, it is not the final triggering event e_n and $e \notin E(s^*)$, then set it in the interval $(\epsilon - u, \epsilon)$. It will then not affect the dynamics of the remaining state transitions and will be rendered inactive by the time the GSMP enters A.

(4) If the event e appears in $\{e_{i+1}, \ldots, e_{n-1}\}$, then set it in the interval $[0, \epsilon - u - t]$ in such a manner that the order of events required in the above path is retained. This event will trigger a state transition before the GSMP enters

the set A, and therefore must be carefully set to ensure the order of events is appropriate.

At the time that the nth event e_n occurs at time $t = \epsilon - u$, any newly active clocks are set to lie in the appropriate interval, and the GSMP then lies in the set A. □

Proof. (Theorem 14) Suppose that $\psi(B) = 0$. We will show that $\pi(B) = 0$. Note that

$$\nu_j(\cdot) \stackrel{\triangle}{=} P_x(X_T \in \cdot, T = j) \leq P_x(X_T \in \cdot) \ll \psi(\cdot),$$

so that $\nu_j \ll \psi$. Furthermore, we have that $\nu_j P \ll \psi$ and by induction, $\nu_j P^k \ll \psi$ for all $k \geq 0$. Now

$$\begin{aligned}
&P_x(X_n \in B) \\
&= P_x(X_n \in B, T \leq n) + P_x(X_n \in B, T > n) \\
&\leq \sum_{j=0}^{n} P_x(X_n \in B, T = j) + P_x(T > n) \\
&= \sum_{j=0}^{n} \int_\Sigma P_x(X_n \in B, T = j, X_j \in dy) + P_x(T > n) \\
&= \sum_{j=0}^{n} \int_\Sigma P_x(X_n \in B | X_j \in dy) P_x(T = j, X_j \in dy) + P_x(T > n) \\
&= \sum_{j=0}^{n} \int_\Sigma P^{n-j}(y, B)\nu_j(dy) + P_x(T > n) \\
&= \sum_{j=0}^{n} \nu_j P^{n-j}(B) + P_x(T > n) \\
&= 0 + P_x(T > n) \qquad (29)
\end{aligned}$$

where (29) follows from the Markov property and the fact that T is a stopping time. Thus

$$\begin{aligned}
\pi(B) &= \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} P_x(X_i \in B) \\
&\leq \lim_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} P_x(T > i) \\
&= 0
\end{aligned}$$

since $T < \infty$ P_x almost surely. Hence $\pi \ll \psi$. □

Proof. (Theorem 15) For $\epsilon > 0$ and $s \in S$, define the set $A_s(\epsilon) = \{s\} \times [\epsilon, \infty)^{|E(s)|}$, the set of Markov chain states in which the GSMP is in state s, and all active clock readings are at least ϵ. Define $A(\epsilon) = \cup_{s\in S} A_s(\epsilon)$ to be the set where all active clock readings are at least ϵ. Let π denote an invariant measure for X, where π is not necessarily finite.

We will show that for some $\epsilon > 0$, $\pi(A(\epsilon)) > 0$, which then implies that $\pi(A_s(\epsilon)) > 0$ for some $s \in S$. If so, then π is equivalent to a maximal irreducibility (probability) measure μ of X (see Meyn and Tweedie [1993, Theorem 10.4.9]). Hence $\mu(A_s(\epsilon)) > 0$, and from Theorem 9.1.4 of Meyn and Tweedie, X_n visits $A_s(\epsilon)$ infinitely often P_x almost surely, for μ almost all x. This then ensures that $\tilde{X}$ is nonexplosive, since $X_n \in A_s(\epsilon)$ implies that the time spent by $\tilde{X}$ in state s on the nth transition Δ_n is given by

$$\Delta_n = \min\{C_n(e)/r_{se} : r_{se} > 0\} \geq \epsilon / \max\{r_{se} : e \in E(s)\}.$$

The set of events is finite, so we have a lower bound on Δ_n, and hence the GSMP is nonexplosive P_x almost surely, for μ almost all x. Hence, it remains to show that $\pi(A(\epsilon)) > 0$, or equivalently $\mu(A(\epsilon)) > 0$ for some $\epsilon > 0$. But by Proposition 4.2.2 part (iii) of Meyn and Tweedie, this will follow if $X_n \in A(\epsilon)$ infinitely often P_x almost surely, for μ almost all x.

Suppose that $X_n \notin A(\epsilon)$ eventually P_y almost surely, for some y and all $\epsilon > 0$. Then, if $c_n = \min\{C_n(e) : e \in E(s)\}$ is the minimum active clock reading on transition n, we must have that $c_n \to 0$ as $n \to \infty$ P_y almost surely. Hence, by Theorem 17.3.2 of Meyn and Tweedie [1993], π must be concentrated on the set where the minimum clock reading is 0:

$$\pi(D) = 0, \quad \text{where } D = \{x = (s, c), \min\{c_e : e \in E(s)\} > 0\}. \tag{30}$$

But for any $x = (s, c)$, our assumption of nonnull clock setting distributions implies that $P(x, D) > 0$ for all x. And then

$$\pi(D) = \int \pi(dx) P(x, D) > 0,$$

contradicting (30). Hence $X_n \in A(\epsilon)$ infinitely often P_y almost surely for all y, and the proof is complete. □

PROOF. (Theorem 16) Suppose that $X_0 = x = (s_0, t_1, \ldots, t_k)$, where $k = |E(s_0)|$ and $t_1 \leq t_2 \leq \cdots \leq t_k$. Define the stopping time $T = \inf\{n : \xi_n > t_k\}$, where, as before, ξ_n is the time of the nth transition in the GSMP, so that T is a time at which all of the original clocks have been reset or are inactive. We will show that the conditions of Theorem 14 are satisfied with the chain X, the measure ϕ_∞, and the stopping time T. It will immediately follow that the stationary distribution π is absolutely continuous with respect to ϕ_∞, establishing the result. First note that $T < \infty$ P_x almost surely for π almost all x, by Theorem 15.

Next, we need to show that $P_x(X_T \in \cdot) \ll \phi_\infty$. We will show this by first establishing the result for a closely related chain, and then showing that the result is, in some sense, "inherited" by X. Let $Y = ((S_n^Y, C_n^Y) : n \geq 0)$ be the GSSMC associated with a GSMP that is identical to X, except that its clock setting distributions are all exponential with mean 1. The chain Y is not necessarily positive Harris recurrent, but we do not need this property.

The chain Y is non-explosive, because the time between events is stochastically bounded below by an exponential random variable with rate $|E|$. Therefore, we may define the stopping time $T(Y)$ for the chain Y analogously to T above.

Let $\tilde{P}_x(\cdot) \stackrel{\triangle}{=} P(\cdot|Y_0 = x)$, and let $\widetilde{\mathbb{E}_x}$ be the associated expectation. Let $p_s = \tilde{P}_x(S^Y_{T(Y)} = s)$ be the probability that the GSMP associated with Y is in state s immediately after transition $T(Y)$. Let the set $A \stackrel{\triangle}{=} \{s\} \times [0, a(1)] \times \cdots [0, a(|E(s)|)]$. The memoryless property of the exponential distribution allows us to conclude that

$$\tilde{P}_x\big(Y_{T(Y)} \in A\big) = p_s \prod_{i=1}^{|E(s)|} \big(1 - e^{-a(i)}\big),$$

so that $\tilde{P}_x(Y_{T(Y)} \in \cdot) \ll \phi_\infty$.

Our assumption that all clock setting distributions have densities, implies that the transition probabilities $P(v, \cdot)$ of X are absolutely continuous with respect to the transition probabilities $\tilde{P}(v, \cdot)$ of Y for each $v \in S$. In particular, there exists $r(\cdot, \cdot)$ such that $P(v, dw) = r(v, w)\tilde{P}(v, dw)$. Hence, we can write

$$\begin{aligned} P_x(X_T \in \cdot) &= \mathbb{E}_x(I(X_T \in \cdot)) \\ &= \widetilde{\mathbb{E}_x}\left(I\big(Y_{T(Y)} \in \cdot\big) \prod_{i=0}^{T(Y)-1} r(Y_i, Y_{i+1})\right) \\ &\ll \tilde{P}_x\big(Y_{T(Y)} \in \cdot\big), \end{aligned}$$

and the absolute continuity of $\tilde{P}_x(Y_{T(Y)} \in \cdot)$ with respect to ϕ_∞ implies that $P_x(X_T \in \cdot) \ll \phi_\infty$.

It remains to show that if a probability measure ν is such that $\nu \ll \phi_\infty$, then $\nu P \ll \phi_\infty$. This property is straightforward, though cumbersome, to show, and so we omit the proof. □

Proof. (Theorem 17) Suppose that X has smoothness index $m < m^*$. Let $x \in \Sigma$ be such that $\lambda(x) > 0$. By definition of m^*, the number of active events/clocks in the state x must be at least m^*. At each state transition, at most one of the m^* original clocks is deactivated or reset (after triggering the transition). Therefore, after $m < m^*$ transitions, at least $m^* - m$ of the original clocks are still active. Hence, $P^m(x, \cdot)$ is concentrated on a set in which $\{c_e = b\}$, for some event e and some $b > 0$. But then, by the clock smoothness assumption, $P^m(x, \cdot)$ is concentrated on a set of π measure 0, and since $P^m(x, \cdot) \geq \lambda(x)\varphi(\cdot)$, so is φ. But φ is absolutely continuous with respect to π (as shown shortly). Hence π must also be concentrated on the set $\{c_e = b\}$, which is a contradiction. Therefore $m \geq m^*$.

To see that φ is absolutely continuous with respect to π, note that the first condition of the minorization implies that

$$\begin{aligned} \pi(\cdot) &= \int_\Sigma P^m(x, \cdot)\pi(dx) \\ &\geq \int_\Sigma \lambda(x)\varphi(\cdot)\pi(dx) \\ &= (\pi\lambda)\varphi(\cdot), \end{aligned}$$

where $\pi\lambda = \int_\Sigma \lambda(x)\pi(dx)$. But $\pi\lambda > 0$, since otherwise the second minorization condition would be violated. □

PROOF. (Proposition 18) The proof is constructive. Let s^* be such that $|E(s^*)| = m^*$ and suppose that $E(s^*) = \{e_1, e_2, \ldots, e_{m^*}\}$. Define $h = \epsilon/(2m^*)$, and set

$$A = \{s^*\} \times (0, h) \times (h, 2h) \times \cdots \times ((m^* - 1)h, m^*h).$$

Suppose that $X_0 \in A$, so that at time 0, the clock reading for event e_i is contained in the interval $((i-1)h, ih)$. Let c_{m^*} be the clock reading for e_{m^*}. Clearly, $\phi_\epsilon(A) = h^{m^*} > 0$.

Assume that all new clock readings are set in the interval $[c'_{m^*}, c'_{m^*} + \epsilon/2)$ where c'_{m^*} is the clock reading for event e_{m^*} at the time the new clock is being set. Then the sequence of activating events will be $e_1, e_2, \ldots, e_{m^*}$, and when event e_m^* triggers a state change, all active clock readings will be bounded above by $\epsilon/2$. Let $s_1, s_2, \ldots, s_{m^*}$ be such that

$$p \stackrel{\triangle}{=} p(s_1; s^*, e_1)p(s_2; s_1, e_2) \ldots p(s_{m^*}; s_{m^*-1}, e_{m^*}) > 0$$

and let δ be a lower bound on the value of the clock setting density components on $[0, \epsilon]$. The probability that a new clock reading will lie in the given interval is at least $q \stackrel{\triangle}{=} \epsilon\delta/2$. At each state transition, at most $|E|$ new clocks are set, so that after m^* transitions, at most $n \stackrel{\triangle}{=} |E|m^*$ new clocks are set. Defining φ to be the uniform probability distribution on

$$B \stackrel{\triangle}{=} \{s_{m^*}\} \times (0, \epsilon/2)^k,$$

where $k = |E(s_{m^*})|$, we see that for $x \in A$,

$$P^{m^*}(x, B) \geq pq^n\varphi(B).$$

Similarly, we can show that $P^{m^*}(x, C) \geq pq^n\varphi(C)$, for any $C \in \mathcal{S}$ and any $x \in A$. Taking $\lambda = pq^n$ yields the result. □

ACKNOWLEDGMENTS

We are grateful to Sean Meyn and Richard Tweedie for useful discussions. We would also like to thank the editors and referees, particularly Jim Wilson, for very helpful comments and suggestions that improved the paper.

REFERENCES

ANDRADÓTTIR, S., CALVIN, J. M., AND GLYNN, P. W. 1995. Accelerated regeneration for Markov chain simulations. *Probability in the Engineering and Informational Sciences 9*, 497–523.

ASMUSSEN, S. 1987. *Applied Probability and Queues*. Wiley, Chichester.

BILLINGSLEY, P. 1968. *Convergence of Probability Measures*. Wiley, New York.

BILLINGSLEY, P. 1986. *Probability and Measure* (2nd ed.). Wiley, New York.

BRATLEY, P., FOX, B. L., AND SCHRAGE, E. L. 1987. *A Guide to Simulation* (2nd ed.). Springer-Verlag, New York.

BURMAN, D. 1981. Insensitivity in queueing systems. *Advances in Applied Probability 13*, 846–859.

CHOW, Y. S., ROBBINS, H., AND TEICHER, H. 1965. Moments of randomly stopped sums. *Ann. Math. Stat. 36*, 789–799.

FOSSETT, L. D. 1979. Simulating generalized semi-Markov processes. Tech. rep., Dept. of Operations Research, Stanford University, Stanford CA.

GLYNN, P. W. 1982. *Simulation Output Analysis for General State Space Markov Chains*. Ph.D. dissertation, Dept. of Operations Research, Stanford University, Stanford, CA.

GLYNN, P. W. 1989a. The covariance function of a regenerative process. Tech. rep., Department of Operations Research, Stanford University, Stanford CA.

GLYNN, P. W. 1989b. A GSMP formalism for discrete event systems. *Proc. IEEE 77*, 14–23.

GLYNN, P. W. 1994. Some topics in regenerative steady-state simulation. *Acta Applic. Mathem. 34*, 225–236.

GLYNN, P. W. AND HEIDELBERGER, P. 1990. Bias properties of budget constrained simulations. *Oper. Res. 38*, 801–814.

GLYNN, P. W. AND IGLEHART, D. L. 1987. A joint central limit theorem for the sample mean and the regenerative variance estimator. *Ann. Oper. Res. 8*, 41–55.

GLYNN, P. W. AND IGLEHART, D. L. 1993. Conditions for the applicability of the regenerative method. *Manage. Sci. 39*, 1108–1111.

GLYNN, P. W. AND L'ECUYER, P. 1995. Likelihood ratio gradient estimation for stochastic recursions. *Advances Appl. Prob. 27*, 1019–1053.

GOLDSMAN, D. AND MEKETON, M. S. 1986. A comparison of several variance estimators. Tech. rep., School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta GA.

GRENANDER, U. AND ROSENBLATT, M. 1984. *Statistical Analysis of Stationary Time Series* (2nd ed.). Chelsea, New York NY.

HAAS, P. J. 1985. *Recurrence and Regeneration in Non-Markovian Simulations*. Ph.D. dissertation, Dept. of Operations Research, Stanford University, Stanford, CA.

HAAS, P. J. 1999. On simulation output analysis for generalized semi-Markov processes. *Commun. Stat.: Stoch. Models 15*, 53–80.

HAAS, P. J. AND SHEDLER, G. S. 1987. Regenerative generalized semi-Markov processes. *Commun. Stat.: Stoch. Models 3*, 409–438.

HENDERSON, S. G. AND GLYNN, P. W. 1999a. Can the regenerative method be applied to discrete-event simulation? In P. A. FARRINGTON, H. BLACK NEMBHARD, D. T. STURROCK, AND G. W. EVANS, Eds., *Proceedings of the 1999 Winter Simulation Conference* (Piscataway NJ, 1999), pp. 367–373. IEEE.

HENDERSON, S. G. AND GLYNN, P. W. 1999b. Derandomizing variance estimators. *Oper. Res. 47*, 907–916.

IGLEHART, D. L. 1975. Simulating stable stochastic systems, V: comparison of ratio estimators. *Naval Res. Log. Quart. 22*, 553–565.

JANSON, S. 1983. Renewal theory for m-dependent variables. *Ann. Prob. 11*, 558–568.

KÖNIG, D., MATTHES, K., AND NAWROTZKI, K. 1967. *Verallgemeinerungen der Erlangschen und Engsetschen Formeln*. Akadamie-Verlag, Berlin.

LAW, A. M. AND KELTON, W. D. 2000. *Simulation Modeling and Analysis* (3rd ed.). McGraw-Hill, New York.

MEKETON, M. S. AND HEIDELBERGER, P. 1982. A renewal theoretic approach to bias reduction in regenerative simulations. *Manage. Sci. 28*, 173–181.

MEYN, S. P. AND TWEEDIE, R. L. 1993. *Markov Chains and Stochastic Stability*. Springer-Verlag, London.

NUMMELIN, E. 1984. *General Irreducible Markov Chains and Non-Negative Operators*. Cambridge University Press.

SHEDLER, G. S. 1987. *Regeneration and Networks of Queues*. Springer-Verlag, New York.

SHEDLER, G. S. 1993. *Regenerative Stochastic Simulation*. Academic Press, Boston.

SONG, W. T. AND SCHMEISER, B. W. 1995. Optimal mean-squared-error batch sizes. *Manage. Sci. 41*, 110–123.

WHITT, W. 1980. Continuity of generalized semi-Markov processes. *Math. Oper. Res. 5*, 494–501.

WOLFF, R. W. 1989. *Stochastic Modeling and the Theory of Queues*. Prentice Hall, Englewood Cliffs NJ.

Received January 2000; revised November 2000 and April 2001; accepted May 2001