

Stationarity Detection in the Initial Transient Problem

SØREN ASMUSSEN,
Aalborg University

PETER W. GLYNN
Stanford University
and

HERMANN THORISSON
University of Iceland

Let $\mathbf{X} = \{X(t)\}_{t \geq 0}$ be a stochastic process with a stationary version \mathbf{X}^* . It is investigated when it is possible to generate by simulation a version $\tilde{\mathbf{X}}$ of \mathbf{X} with lower initial bias than \mathbf{X} itself, in the sense that either $\tilde{\mathbf{X}}$ is strictly stationary (has the same distribution as \mathbf{X}^*) or the distribution of $\tilde{\mathbf{X}}$ is close to the distribution of \mathbf{X}^* . Particular attention is given to regenerative processes and Markov processes with a finite, countable, or general state space. The results are both positive and negative, and indicate that the tail of the distribution of the cycle length τ plays a critical role. The negative results essentially state that without some information on this tail, no a priori computable bias reduction is possible; in particular, this is the case for the class of all Markov processes with a countably infinite state space. On the contrary, the positive results give algorithms for simulating $\tilde{\mathbf{X}}$ for various classes of processes with some special structure on τ . In particular, one can generate $\tilde{\mathbf{X}}$ as strictly stationary for finite state Markov chains, Markov chains satisfying a Doeblin-type minorization, and regenerative processes with the cycle length τ bounded or having a stationary age distribution that can be generated by simulation.

Categories and Subject Descriptors: I.6.6 [Simulation and Modeling]: Simulation Output Analysis; I.6.8 [Simulation and Modeling]: Types of Simulation—Monte Carlo

General Terms: Algorithms

Additional Key Words and Phrases: Coupling, initial transient, regenerative processes, simulation, stationary processes

1. INTRODUCTION

When performing a steady-state simulation, simulation analysts are often concerned with the problem of dealing with the initial transient. The term

Authors' addresses: S. Asmussen, Institute of Electronic Systems, Aalborg University, Fr. Bajersv. 7, DK-9220 Aalborg, Denmark P. W. Glynn, Department of Operations Research, Stanford University, Stanford, CA 94305-4022. H. Thorisson, Science Institute, University of Iceland, Dunhaga 3, 107 Reykjavik, Iceland.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1992 ACM 1049-3301/92/0400-0130 \$01.50.

ACM Transactions on Modeling and Computer Simulation, Vol. 2, No. 2, April 1992, Pages 130-157.

“initial transient” refers to that initial segment of the simulation that is contaminated by bias introduced by starting the system in some state that is not typical of the long-run behavior of the system. The observations gathered during the initial transient are therefore not representative of the steady-state behavior of the system and are biased. Perhaps the most popular means of dealing with the initial transient is to discard the observations gathered during this period. In other words, the simulation analyst lets the simulation “warm up” before collecting any observations.

Of course, the key question is to determine how long the warm-up period must be for a given simulation. An essentially equivalent formulation of the problem is to identify that time at which the initial transient terminates and steady-state behavior begins. Since steady-state behavior is characterized by stationarity of the stochastic process, we can view the initial transient problem as involving the determination of that time at which the simulation is behaving like a stationary stochastic process. This paper is concerned with the question of existence and construction of such stationarity detection times (and suitable generalizations). The algorithms developed here are intended primarily to establish the boundaries of what is theoretically possible, rather than as proposals to the practical simulation analyst as to how to eliminate initial bias in real-world simulations—in fact, most of them have a large cost in terms of run lengths. Nevertheless, the paper seems to indicate that sampling from stationary distributions can be done much more frequently than usually considered possible, and thereby opens up the interesting and important question of designing efficient algorithms.

In the set-up of the paper, we consider a stochastic process $\mathbf{X} = \{X(t)\}$, which is available basically through simulated values. Typically, \mathbf{X} is regenerative or Markovian with a finite, countably infinite, or general state space. Note that these two settings are not intrinsically different because general Markov processes can typically be made regenerative (see Asmussen [4]), and conversely a regenerative process with i.i.d. cycles can be made Markovian by adding extra variables.

In the Markov case, the problem, roughly speaking, is to determine a random time at which the process possesses the stationary distribution. We assume, in constructing such a random time, that the constructions depend only on simulation of a finite time segment of the chain. In particular, we require that the algorithm that implements the construction be independent of the explicit transition function of the chain. We impose this requirement because most discrete-event simulations are implemented without any need to ever directly calculate an explicit expression for the transition function. In addition, the development of universal algorithms that are valid over broad classes of simulations is an important issue in writing software intended to be of wide applicability. Finally, this assumption is necessary in order to obtain nontrivial nonexistence results: if the algorithm is permitted to adapt itself to the particular transition function in question, then one can design deterministic algorithms for computing the stationary distribution numerically, so that one can then argue that the initial transient problem becomes irrelevant.

Similarly, in the regenerative setting the problem is to generate the stationary delay, assuming only that the zero-delayed process is available through simulated values.

The initial transient problem has challenged the simulation community for many years. A few of the papers that have addressed this question are [9, 12, 32, 33], and [25]; see also Fishman [11] and Law [21]. It is probably safe to say that no technique yet proposed satisfactorily solves this problem. It is also worth noting that when one estimates the steady-state via multiple replications of the process, the ability to determine the end of the initial transient is enhanced. The basic idea is that by averaging over multiple replicates, much of the variability in the system is damped out, so that the convergence to steady-state is easier to determine. Among the papers that take advantage of this idea are [20, 29, 30]. In the current paper, our interest focuses on stationarity detection rules that are based on a single run of the system.

We will see, in Section 2 of this paper, precisely why the initial transient problem has been so challenging. We will prove, in a mathematically precise sense, that without some restrictions on the class of simulations to be considered, there can exist no universally satisfactory means for detecting stationarity in a stochastic simulation. This negative result is probably expected, and suggests that any successful stationarity detection rule will need to take explicit advantage of some additional structure of the system being simulated.

In the rest of the paper we complement the above negative result with positive ones, which are perhaps more surprising. In Section 3 we show that it is possible to generate a r.v. Z having the stationary distribution of a finite Markov chain $\{X_n\}$, using only simulated values and randomization, and in Section 4 it is shown (using some recently developed ideas of [27]) that for certain classes of regenerative stochastic processes, one can identify a random time T such that the system is in exact stationarity at this instant. These constructions are possible even for certain systems in which the steady-state distribution is not analytically available and must be simulated. The approach taken here to developing stationarity detection rules strongly suggests that, in the regenerative setting, one must take advantage of a priori knowledge of the tail behavior of the regenerative cycle-length random variable.

Section 5 discusses settings in which approximate stationarity can be achieved. In Section 6 we provide further discussion, and Section 7 concludes the paper with some illustrative examples and applications. Unless otherwise stated, all proofs are deferred to the Appendix.

2. STATIONARITY DETECTION: DEFINITIONS AND BASIC THEORY

We restrict our formulation and discussion in this section to the Markov chain setting. However, the ideas described here can be easily extended to the general discrete-event simulation context. One need only observe that if one views the typical discrete-event simulation at transition epochs, one can

make the process Markovian by adding supplementary variables to the state description that include information on the time that remains before the “clocks” corresponding to each possible event will trigger a state transition of the system. It follows that a discrete-event system can be written as a functional of a discrete-time Markov chain, taking values in a complicated state space S in which both physical state and clock state is recorded. For additional information on this way of looking at discrete-event simulations, see Glynn [14].

Suppose that $\mathbf{X} = \{X_n\}_{n=0,1,\dots}$ is a Markov chain taking values in S . We can (without any loss of generality) view \mathbf{X} as being defined on the probability space

$$\Omega = (S \times (0, 1)) \times (S \times (0, 1)) \times \dots$$

A typical element in Ω then takes the form $\omega = \{(x_n, u_n)\}_{n=0,1,\dots}$. The sequence \mathbf{X} can be defined via the coordinate projections $X_n(\omega) = x_n$, and we further let \mathbf{U} be the sequence of random variables defined by $U_n(\omega) = u_n$.

Let K be a transition kernel defined on S , so that $K(x, B)$ represents the probability that the chain \mathbf{X} moves from x into $B \subseteq S$ in one step. For each initial distribution μ on S , we can then define a probability distribution $\mathbb{P}_{\mu, K}$ on Ω via the formula

$$\begin{aligned} & \mathbb{P}_{\mu, K}[X_0 \in A_0, \dots, X_n \in A_n, U_0 \in B_0, \dots, U_n \in B_n] \\ &= \int_{A_0} \mu(dx_0) \int_{A_1} K(x_0, dx_1) \cdots \int_{A_n} K(x_{n-1}, dx_n) \cdot \int_{B_0} dy_0 \cdots \int_{B_n} dy_n. \end{aligned}$$

Hence, under the distribution $\mathbb{P}_{\mu, K}$, \mathbf{X} is a Markov chain having initial distribution μ and transition kernel K . Also, \mathbf{U} is a sequence of i.i.d.-uniform $(0, 1)$ r.v.'s which is itself independent of \mathbf{X} . We need the uniform r.v.'s in order to define randomized algorithms for detecting stationarity. Much of our subsequent discussion will involve such randomized detection rules.

We let \mathcal{M} denote the subset of transition kernels on S such that for each $K \in \mathcal{M}$, there exists a unique stationary distribution π_K .

We say that T is a *random time* if T is a nonnegative integer-valued r.v. defined on Ω and let $\mathbf{X}(T + \cdot) = \{X(T + t)\}_{t \geq 0}$ be the post- T process. Roughly speaking, our goal is to construct a random time T such that $\mathbf{X}(T + \cdot)$ is in steady state (is strictly stationary).

Definition 2.1. Let $K \in \mathcal{M}$. The random time T is said to be a stationarity detection time for K if for each initial distribution μ

$$\mathbb{P}_{\mu, K}(\mathbf{X}(T + \cdot) \in \cdot) = \mathbb{P}_{\pi_K, K}(\mathbf{X} \in \cdot). \quad (2.1)$$

One way to construct stationarity detection times is by means of randomized stopping times (recall that a random time T is a *randomized stopping time* if for each n there exists a deterministic 0–1 valued function f_n such that $I(T = n) = f_n(X_0, U_0, \dots, X_n, U_n)$; sometimes also the term *nonanticipating*

ating is used for a randomized stopping time). Such randomized stopping times have the following nice property:

PROPOSITION 2.1. *Let $K \in \mathcal{K}$. If T is a randomized stopping time such that X_T has distribution π_K for each initial distribution μ , then T is a stationarity detection time for K .*

The proof of this follows immediately from the strong Markov property of \mathbf{X} .

Ideally, one would like stationarity detection times to detect stationarity immediately if the initial distribution is π_K . Our first result shows that no such stationarity detection time typically exists.

PROPOSITION 2.2. *Let $K \in \mathcal{K}$ and assume that π is not concentrated on any single point $x \in S$. Then there exists no stationarity detection time T for K such that*

$$\mathbb{P}_{\pi_K, K}(T = 0) = 1. \quad (2.2)$$

Thus, requirement (2.2) demands too much from random time T . If we drop (2.2), it turns out that stationarity detection times can often be constructed:

Example 2.1. Suppose that S is finite or countably infinite and that K is irreducible. If \mathbf{X} is positive recurrent under K , there exists a unique stationary distribution π_K , and (by applying inversion), we can find a deterministic function g such that $g(U_0)$ has distribution π_K . Then

$$T = \inf\{n = 0, 1, \dots : X_n = g(U_0)\}$$

is a randomized stopping time such that X_T has distribution π_K , and we may apply Proposition 2.1.

However, this construction obviously “cheats” by constructing the function g (and hence the stopping rule T) from explicit knowledge of π_K . This suggests that a more appropriate formulation for a stationarity detection time ought to somehow forbid the simulation analyst from using explicit knowledge of the stationary distribution to construct T .

We can accomplish this by requiring that T work uniformly well over a suitably large class \mathcal{N} of transition kernels K . Being defined only in terms of the simulated data \mathbf{X} and \mathbf{U} , T cannot explicitly modify itself to reflect knowledge of the various stationary distributions.

Definition 2.2. Let $\mathcal{N} \subseteq \mathcal{K}$. We say that a random time T is a \mathcal{N} -uniform stationarity detection time if T is a stationarity detection time for each $K \in \mathcal{N}$.

Perhaps surprisingly, it is often possible to construct such detection times.

Example 2.2. Suppose again that S is finite or countably infinite. Without loss of generality, we can take S to be $\{0, 1, \dots\}$. Let \mathcal{K}_1 be the class of irreducible positive recurrent transition matrices K defined on S . For $x \in S$, let

$$F(\mathbf{X}, x) = \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{k=0}^n I(X_k \leq x).$$

For a fixed but arbitrary $K \in \mathcal{M}_1$, the strong law of large numbers for X implies that $F(\mathbf{X}, \cdot)$ is almost surely equal to the distribution function of π_K , and thus,

$$T_1 = \inf\{n = 0, 1, \dots : X_n = F^{-1}(\mathbf{X}, U_0)\}$$

coincides a.s. with the T of Example 2.1. Thus T_1 is a stationarity detection time for the given K , and hence for all $K \in \mathcal{M}_1$.

Of course, here the difficulty is that T is constructed after observing the entire (infinite) sample trajectory of \mathbf{X} . Such a random time T cannot be implemented in a practical setting. This motivates restricting attention to stationarity detection times, which are *implementable* in the following sense.

Definition 2.3. A r.v. Z^* is *implementable*, if there exists an a.s. finite randomized stopping time β such that Z^* is a deterministic function of $\beta, (X_1, U_1), \dots, (X_\beta, U_\beta)$ alone.

We are now ready to state our main nonexistence result for strong stationarity times. It proves that well-behaved stationarity detection times T fail to exist even when one restricts attention to Markov chains with countably infinite state space. In fact, let \mathcal{M}_2 be the class of aperiodic irreducible positive recurrent transition matrices K .

THEOREM 2.1. *Assume that S is countably infinite. Then there exists no implementable \mathcal{M}_2 – uniform stationarity time.*

In fact, an even stronger result (Theorem 2.2 below) can be proved. Recall that the total variation distance between probability measures μ and ν on S is defined by

$$\|\mu - \nu\| = 2 \sup_{B \subseteq S} |\mu(B) - \nu(B)|.$$

Definition 2.4. Let $\mathcal{N} \subseteq \mathcal{M}$. (a) The family \mathcal{N} is said to be a weak *uniformity class* for the initial transient problem if, for each $\epsilon > 0$, there exists an implementable r.v. $Z^*(\epsilon)$ such that

$$\|\mathbb{P}_{\mu, K}(Z^*(\epsilon) \in \cdot) - \pi_K(\cdot)\| < \epsilon \quad (2.3)$$

for any initial distribution μ on S and any $K \in \mathcal{N}$. (b) The family \mathcal{N} is said to be a *uniformity class* if there exists an implementable r.v. Z^* such that $\mathbb{P}_{\mu, K}(Z^* \in \cdot) = \pi_K(\cdot)$ for any initial distribution μ on S and any $K \in \mathcal{N}$.

We call the r.v.'s $Z^*(\epsilon)$ and Z^* appearing in Definition 2.4 an ϵ -stationary r.v. and stationary r.v., respectively. If $Z^*(\epsilon)$ can be represented as $X_{T^{(\epsilon)}}$ for some random time $T^{(\epsilon)}$, we call $T^{(\epsilon)}$ an ϵ -stationarity time.

Note that (a) demands only that the marginal distribution of $Z^*(\epsilon)$ be approximately stationary. The extension from stationarity detection *times* to stationary *random variables* is motivated by the fact that given a stationary detection r.v. Z^* , one can simulate a strictly stationary version of the Markov chain by starting from $X_0 = Z^*$, and given a ϵ -stationary r.v. $Z^*(\epsilon)$, the version $\{\tilde{X}_n\}$ of the Markov chain started from $\tilde{X}_0 = Z^*(\epsilon)$ satisfies

$\|\mathbb{P}_{\mu, K}(\tilde{X}_n \in \cdot) - \pi_K(\cdot)\| < \epsilon$ because of the general inequality $\|\mathbb{P}_{\mu, K}(\tilde{X}_n \in \cdot) - \mathbb{P}_{\nu, K}(\tilde{X}_n \in \cdot)\| \leq \|\mu - \nu\|$.

THEOREM 2.2. *Assume that S is countably infinite. Then \mathcal{H}_2 is not a weak uniformity class.*

This follows from the following result:

PROPOSITION 2.3. *Let $S = \{0, 1, 2, \dots\}$, let $K^{(0)} = (k_{ij}^{(0)})_{i,j \in S}$ be a given transition matrix in \mathcal{H} , and let $\mathcal{N}(K^{(0)})$ be the set of transition matrices $K = (k_{ij})_{i,j \in S} \in \mathcal{H}$ such that there exists an integer $A = A(K^{(0)})$ such that $k_{ij} = k_{ij}^{(0)}$ for $i, j \leq A$. Then $\mathcal{N}(K^{(0)})$ is not a ϵ -uniformity class for $0 < \epsilon < 2$.*

These results suggest that the class \mathcal{N} of transition kernels needs to be carefully chosen in order to have any chance of being able to construct a uniformly well-behaved stationarity detection rule. The remainder of this paper is concerned with describing the type of information that needs to be present in \mathcal{N} so as to permit such constructions.

3. SIMULATION OF STATIONARY FINITE MARKOV CHAINS

Our main result on finite Markov processes is the following.

THEOREM 3.1. *The class $\mathcal{H}^{(s)}$ of irreducible Markov chains with a fixed number s of states is a uniformity class.*

Thus, there exist algorithms generating an r.v. having the stationary distribution π of a finite Markov chain using only simulated values and randomization. We proceed to describe one such algorithm, thereby providing a proof of Theorem 3.1.

The first step is to translate the problem into one on continuous-time Markov process $\{Y(t)\}_{t \geq 0}$ by uniformization (Poissonification). Indeed, it is a standard fact that if V_1, V_2, \dots are i.i.d. exponential (say with unit rate), then the process $\{Y(t)\}$ defined by

$$Y(t) = X_0, \quad 0 \leq t < V_1, \quad Y(t) = X_n, \quad V_1 + \dots + V_{n-1} \leq t < V_1 + \dots + V_n$$

is a s -state irreducible Markov process with the same stationary distribution as $\{X_n\}$. Note that $\{Y(t)\}$ does not necessarily jump at $V_1 + \dots + V_n$, but only if $X_{n+1} \neq X_n$. In terms of the transition probabilities $(k_{ij})_{i,j \in E}$ for $\{X_n\}$, $\{Y(t)\}$ has intensity $\lambda_{ij} = k_{ij}$ for jumping from i to j when $i \neq j$.

The next step is to observe that the construction of a stationary detection r.v. for $\{Y(t)\}$ is easy when $s = 2$ (e.g., $S = \{1, 2\}$), where

$$\pi_1 = \frac{\lambda_{21}}{\lambda_{12} + \lambda_{21}}, \quad \pi_2 = \frac{\lambda_{12}}{\lambda_{12} + \lambda_{21}}.$$

Indeed, let T_i be the first holding time of state i , $i = 1, 2$. Then T_1, T_2 are exponential with intensities λ_{12} , resp. λ_{21} , and an easy calculation shows that

$$\mathbb{P}(T_1 > T_2) = \frac{\lambda_{21}}{\lambda_{12} + \lambda_{21}} = \pi_1. \quad (3.4)$$

Thus, we have the following algorithm:

Algorithm A. (Generating a Stationary Detection r.v. Z for a Markov Chain $\{X_n\}$ with two states 1, 2)

1. Let $j \leftarrow 1, t \leftarrow 0$
2. Generate V , an exponential random variable with unit intensity, and X_1 starting from $X_0 = 1$. Let $j \leftarrow X_1, t \leftarrow t + V$
3. If $j = 2$, let $T_1 \leftarrow t$;
else return to 2
4. Repeat steps 1, 2, 3 with states 1 and 2 interchanged to generate T_2
5. If $T_1 > T_2$, let $Z \leftarrow 1$;
else let $Z \leftarrow 2$

The last step (and the most intricate one) is to treat the case $s > 2$ recursively. We need to introduce the F -valued process $\{Y^{(F)}(t)\}_{t \geq 0}$, defined as the process $\{Y(t)\}$ on F , where $F \subseteq S$ (see, e.g., [4, pp. 13–14]). This means that in the path of $\{Y(t)\}$ we delete all segments where $Y(t) \notin F$ and glue together the remaining segments. Algorithmically, this can be implemented as follows:

Algorithm B. (Generating the first holding time $T^{(F)}(i)$ and the next state $Y^{(F)}(i)$ of $\{Y^{(F)}(t)\}$ starting from $Y^{(F)}(0) = X_0 = i \in F$)

1. Let $j \leftarrow i, t \leftarrow 0$
2. Generate V , an exponential random variable with unit intensity. If $j \in F$, let $t \leftarrow t + V$.
3. Generate X_1 starting from $X_0 = j$. Let $j \leftarrow X_1$.
4. If $j \neq i$ and $j \in F$, let $T^{(F)}(i) \leftarrow t, Y^{(F)}(i) \leftarrow j$;
else return to 2

It is well known that the stationary distribution $\pi^{(F)}$ of $\{Y^{(F)}(t)\}$ is obtained by conditioning $\pi = \pi^{(S)}$ to F :

$$\pi_i^{(F)} = \frac{\pi_i}{\pi_F} \quad \text{where} \quad \pi_F = \sum_{j \in F} \pi_j. \quad (3.5)$$

We also have the principle of local balance ([19, p. 8]): when $\{Y^{(F+G)}(t)\}$ is in stationarity, the rate of flow of mass from F to G is the same as the flow of mass from G to F (here F, G are disjoint subsets of S). In terms of the intensities $\lambda_{ij}^{(F+G)}$ for $\{Y^{(F+G)}(t)\}$ (which exist because $\{Y^{(F+G)}(t)\}$ is a Markov process—the analytical form is unimportant here) thus

$$\sum_{i \in F} \pi_i^{(F+G)} \lambda_{iG}^{(F+G)} = \sum_{j \in G} \pi_j^{(F+G)} \lambda_{jF}^{(F+G)}, \quad (3.6)$$

where $\lambda_{iG}^{(F+G)} = \sum_{j \in G} \lambda_{ij}^{(F+G)}$. We can rewrite (3.6) as

$$\pi_F^{(F+G)} \sum_{i \in F} \pi_i^{(F)} \lambda_{iG}^{(F+G)} = \pi_G^{(F+G)} \sum_{j \in G} \pi_j^{(G)} \lambda_{jF}^{(F+G)}. \quad (3.7)$$

Now assume that we can generate exponential r.v.'s $T_F^{(F+G)}, T_G^{(F+G)}$ having parameters $\sum_{i \in F} \pi_i^{(F)} \lambda_{iG}^{(F+G)}$, resp. $\sum_{j \in G} \pi_j^{(G)} \lambda_{jF}^{(F+G)}$, and stationarity detection

r.v.'s $Z^{(F)}, Z^{(G)}$ having distributions $\pi^{(F)}$, resp. $\pi^{(G)}$. Then as in (3.4),

$$\begin{aligned} \mathbb{P}(T_F^{(F+G)} > T_G^{(F+G)}) &= \frac{\sum_{i \in F} \pi_i^{(F)} \lambda_{iG}^{(F+G)}}{\sum_{i \in F} \pi_i^{(F)} \lambda_{iG}^{(F+G)} + \sum_{j \in G} \pi_j^{(G)} \lambda_{jF}^{(F+G)}} \\ &= \frac{1}{1 + \sum_{j \in G} \pi_j^{(G)} \lambda_{jF}^{(F+G)} / \sum_{i \in F} \pi_i^{(F)} \lambda_{iG}^{(F+G)}} \\ &= \frac{1}{1 + \pi_G^{(F+G)} / \pi_F^{(F+G)}} = \pi_F^{(F+G)}, \end{aligned}$$

where we used (3.7) for the third equality. That is, we can use the ordering of $T_F^{(F+G)}, T_G^{(F+G)}$ to decide whether $Z^{(F+G)}$ should be in F or G , and next $Z^{(F)}$ or $Z^{(G)}$ to decide which value in F, G is appropriate. Thus

$$\mathbb{P}(Z^{(F+G)} = i) = \mathbb{P}(T_F^{(F+G)} > T_G^{(F+G)}) \mathbb{P}(Z^{(F)} = i) = \pi_i^{(F+G)} \pi_i^{(F)} = \pi_i^{(F+G)}$$

for $i \in F$ as desired, and similarly $\mathbb{P}(Z^{(F+G)} = j) = \pi_j^{(F+G)}$ for $j \in G$.

To construct $T_F^{(F+G)}$, note that the distribution of $T_F^{(F+G)}$ is that of the minimum of exponential r.v.'s W_i with intensity $\pi_i^{(F)} \lambda_{iG}^{(F+G)}$ for the i th state ($i \in F$). If $\lambda_{iG}^{(F+G)} > 0$, we can sample a Poisson stream with intensity $\lambda_{iG}^{(F+G)}$ by repeatedly starting $\{Y^{(F+G)}(t)\}$ in state i and accumulate the time until a transition to G occurs. We can then thin this stream by generating a sequence of i.i.d. copies of $Z^{(F)}$, say $Z_1^{(F)}, Z_2^{(F)}, \dots$. The n th point is retained if $Z_n^{(F)} = i$, thereby obtaining a Poisson stream with intensity $\pi_i^{(F)} \lambda_{iG}^{(F+G)}$. The r.v. W_i is then the first epoch of the thinned Poisson process. In practice, this construction requires a small modification, since $\lambda_{iG}^{(F+G)}$ may be zero for some i . However, by irreducibility $\lambda_{iG}^{(F+G)} > 0$ for at least one $i \in F$, and this ensures that the following algorithm is valid (roughly, the idea is to visit the states i in F cyclically; t_i indicates the amount of time in which the i th Poisson stream has been simulated so far, T_F is the current earliest occurrence of a transition to G , and $s_i = I(s_i > T_F)$ is a binary variable indicating whether it is necessary to simulate any further; by irreducibility, T_F will assume a finite value at some stage):

Algorithm C. (Generating $Z^{(F+G)}$ if it is Known How to Generate $Z^{(F)}, Z^{(G)}$)

1. Number the states in F in some way, say $F = \{l_0, \dots, l_{p-1}\}$
2. Let $T_F \leftarrow \infty, t_i \leftarrow 0, s_i \leftarrow 0, i = 0, \dots, p-1$
3. Let $i \leftarrow p-1$
4. If all $s_k = 1$, go to 10;
else let $i \leftarrow (i+1) \bmod p$
5. If $s_i = 1$, return to 4
6. Generate $Y^{(F+G)}(l_i), T^{(F+G)}(l_i)$ using Algorithm B and let $j \leftarrow Y^{(F+G)}(l_i)$,
 $t_i \leftarrow t_i + T^{(F+G)}(l_i)$.
If $t_i > T_F$, let $s_i = 1$ and return to 4
7. If $j \in G$, generate $Z^{(F)}$ and let $k \leftarrow Z^{(F)}$;
else return to 4
8. If $j \neq k$, return to 4;
else if $t_i < T_F$, let $T_F \leftarrow t_i$
9. Return to 4
10. Repeat steps 1, ..., 9 with F and G interchanged to generate T_G
11. If $T_F > T_G$, generate $Z^{(F)}$ and let $Z^{(F+G)} \leftarrow Z^{(F)}$;
else generate $Z^{(G)}$ and let $Z^{(F+G)} \leftarrow Z^{(G)}$

To generate $Z = Z^{(S)}$, start by noting that obviously $Z^{(F)} = i$ if $F = \{i\}$ is a one-point set. Using Algorithm C, one can then generate $Z^{(F)}$ for a partitioning of S into two-point sets; then $Z^{(F)}$ for a coarser partitioning into three- or four-point sets, and so on.

In Section 6, we give some (rather pessimistic) estimates indicating how the number of steps needed to generate Z may depend on the number s of states.

Remark 3.1. As a corollary of Theorem 3.1, it follows that *there exists a stationarity detection time T for the class $\mathcal{M}^{(s)}$ of irreducible Markov chains with a fixed number s of states.* Indeed, by a minor (though presumably less efficient) variant of the construction, we may assume that all necessary values of \mathbf{X} for $Z^{(S)}$ are generated by observing a single sample path, say up to time T_1 where T_1 is a randomized stopping time. Then

$$T = \inf\{n > T_1: X_n = Z^{(S)}\}$$

is a stationarity detection time, i.e. $\mathbb{P}(X_T = i) = \pi_i$.

Remark 3.2. From a statistical point of view, a relevant concept to ask for would be an unbiased sequential randomized estimator (u.s.r.e) of $\boldsymbol{\pi}$, i.e., a random probability vector $\hat{\boldsymbol{\pi}}$ which is measurable w.r.t. \mathcal{F}_σ for some randomized stopping time σ and has the property $\mathbb{E}\hat{\pi}_i = \pi_i$ for all i . It seems worthwhile to note that the *existence of an u.s.r.e. of $\boldsymbol{\pi}$ within the statistical model of all irreducible s -state Markov chains is equivalent to the existence of a stationarity detection rule.* Indeed, obviously, a stationarity detection rule like $Z^{(S)}$ leads immediately to an u.s.r.e by letting $\hat{\pi}_i = 1$ when $Z^{(S)} = i$, all other $\hat{\pi}_j = 0$. Conversely, if $\hat{\boldsymbol{\pi}}$ is an u.s.r.e., we can choose an additional uniform random number V and obtain a stationarity detection rule Z as $\inf\{i = 1, \dots, s: \hat{\pi}_1 + \dots + \hat{\pi}_i > V\}$.

4. SIMULATION OF STATIONARY REGENERATIVE PROCESSES

Let S be a state space endowed with a metric under which S is separable (e.g., \mathbb{R}^d). If $\mathbf{X} = \{X(t)\}_{t \geq 0}$ is a right-continuous stochastic process taking values in S , we say that \mathbf{X} is a (nondelayed) *wide-sense regenerative process* if there exist random times $0 = T(0) < T(1) < \dots$ such that

- (i) $\mathbf{X}(T(n) + \cdot) \stackrel{\mathcal{L}}{=} \mathbf{X}(\cdot)$ for $n \geq 1$;
- (ii) $T(n)$ is independent of $\mathbf{X}(T(n) + \cdot)$ for $n \geq 1$.

Note that we are not requiring the process evolution prior to time $T(n)$ to be independent of that subsequent to $T(n)$. Instead, the post- $T(n)$ process $\mathbf{X}(T(n) + \cdot)$ is required to be independent only of the time $T(n)$ itself. This extension of classical regeneration (known as wide-sense regeneration) turns out to be useful in the study of Harris recurrent Markov chains; see pp. 150–158 of [4]. (Note that a discrete time sequence $\{X_n\}_{n=0,1,\dots}$ can be

analyzed by studying the associated continuous time process $\{X(t)\}_{t \geq 0}$, where $X(t) = X_{\lfloor t \rfloor}$, and $\lfloor t \rfloor$ denotes the greatest integer less than or equal to T .)

We say that $\mathbf{X}^* = \{X^*(t)\}_{t \geq 0}$ is a *stationary version* of \mathbf{X} if \mathbf{X}^* is a strictly stationary process possessing an associated random time $T^*(0)$ such that

- (i) $\mathbf{X}^*(T^*(0) + \cdot) \stackrel{\mathcal{L}}{=} \mathbf{X}(\cdot)$
- (ii) $T^*(0)$ is independent of $\mathbf{X}^*(T^*(0) + \cdot)$.

In the remainder of this section, we describe two settings in which \mathbf{X}^* can be simulated, given the ability to simulate \mathbf{X} .

Let $\tau_n = T(n) - T(n-1)$ for $n \geq 1$, write $\tau = \tau_1$ for the generic cycle, and set $m = \mathbb{E}\tau$. The following description of the stationary version \mathbf{X}^* is given in [27]:

PROPOSITION 4.1. *Assume $m < \infty$. Suppose that $\mathbf{X}' = \{X'(t)\}_{t \geq 0}$ is an S -valued stochastic process with associated random time $T'(0)$ satisfying*

$$\mathbb{P}(T'(0) \in dx) = \frac{x}{m} \mathbb{P}(\tau \in dx) \quad (4.8)$$

and

$$\mathbb{P}(\mathbf{X}' \in \cdot | T'(0) = x) = \mathbb{P}(\mathbf{X} \in \cdot | \tau_1 = x) \quad (4.9)$$

for each $x \geq 0$. If U is a uniformly distributed r.v. on $[0, 1]$ and independent of \mathbf{X}' , then \mathbf{X}^* is a stationary version of \mathbf{X} where

$$X^*(t) = X'(UT'(0) + t), \quad t \geq 0.$$

There is an intuitive explanation for why this construction should give a stationary version. Imagine that the process \mathbf{X} has been running for a time interval of length t and that we pick a point η uniformly in the interval $[0, t]$. Then, the post- η process $\mathbf{X}(\eta + \cdot)$ converges to a stationary version \mathbf{X}^* of \mathbf{X} when $t \rightarrow \infty$. The possibility of the point η ending in a given cycle interval of length x should be proportional to x , and the relative number of such intervals is $\mathbb{P}(\tau \in dx)$, which (together with the “normalization” $\int x \mathbb{P}(\tau \in dx) = m$) gives us (4.8). Given the length of the picked cycle, it should behave as an ordinary cycle, i.e., (4.9) should hold. Finally, the picked point should lie uniformly within its “length-biased” interval (corresponding to the U in Proposition 4.1), independently of everything else.

Suppose that τ has a density. In that case, (4.8) states that the ratio of the density of $T'(0)$ to that of τ is x/m . Hence, if the r.v. τ is a.s. bounded above by the deterministic finite constant a , say, the ratio will be bounded above by a/m . It is well known (see, for example, [21]) that the boundedness of the ratio permits one to generate the r.v. $T'(0)$ via acceptance-rejection (given an algorithm to generate ordinary regenerative cycle lengths with the distribution of τ). A similar analysis is valid without assuming the existence of densities, in particular when τ is a discrete r.v.

Combining this acceptance-rejection idea and Proposition 3.1 leads to the following algorithm for generating the stationary version \mathbf{X}^* corresponding to \mathbf{X} :

Algorithm D. (Generating a Stationary Version \mathbf{X}^* When the Cycle Lengths Are Bounded, $\tau \leq a < \infty$ a.s.)

1. Let $n \leftarrow 1$
2. Generate \mathbf{X} over $[T(n-1), T(n))$
3. Generate V_n , a uniform r.v. on $[0, 1]$
4. If a $V_n < \tau_n$, go to 5;
else let $n \leftarrow n + 1$ and return to 2
5. Generate U , a uniform r.v. on $[0, 1]$, and let $X^*(t) \leftarrow X(T(n-1) + U\tau_n + t)$, $t \geq 0$.

Note that $T^*(0) = (1 - U)\tau_n$ and that the probability of acceptance at step 4 of the above algorithm is m/a ; the expected number of times that the test at step 4 is executed is therefore a/m .

Algorithm A implies that the class of wide-sense regenerative processes with cycle lengths bounded above by a fixed constant form a uniformity class. Because of the intimate relation between regenerative processes and recurrent Markov chains, it is also clear how to translate this into a result about Markov chains.

Unfortunately, it is only rarely the case that the cycle lengths of a regenerative process are bounded (but see Example 7.2 for an interesting exception). However, [27] provides us with the tools necessary to develop a second interesting class of wide-sense regenerative processes for which generation of the stationary version \mathbf{X}^* is possible.

PROPOSITION 4.2. *Assume $m < \infty$. Then \mathbf{X}^* is a stationary version of \mathbf{X} if there exists an associated random time β such that*

$$\mathbb{P}(\beta \in dx) = \frac{1}{m} \mathbb{P}(\tau > x) \quad (4.10)$$

and

$$\mathbb{P}(\mathbf{X}^*(t + \cdot) \in \cdot | \beta = t) = \mathbb{P}(\mathbf{X}(t + \cdot) \in \cdot | \tau_1 > t) \quad (4.11)$$

for each $x, t \geq 0$.

Note that (4.10) states that β has the stationary age distribution (or the stationary excess life distribution) for the wide-sense regenerative process \mathbf{X} ; see p. 116 of [4]. Thus, (4.11) basically asserts that the stationary version \mathbf{X}^* can be obtained by conditioning the original process \mathbf{X} in such a way that the cycle currently in progress has the appropriate stationary age distribution.

The key to applying Proposition 3.2 to simulate \mathbf{X}^* is the ability to generate the r.v. β from the distribution specified by (4.10). We say that the stationary age distribution is *simulatable* if such variable generation is possible. Proposition 3.2 immediately establishes the validity of the following

algorithm:

Algorithm E. (Generating a stationary version \mathbf{X}^* when the stationary age distribution is simulatable)

1. Generate a r.v. β with distribution (4.10)
2. Let $n \leftarrow 1$
3. Generate \mathbf{X} over $[T(n-1), T(n))$
4. If $\tau_n > \beta$, go to 4;
else let $n \leftarrow n + 1$ and return to 2
5. Let $X^*(t) \leftarrow X(T(n-1) + \beta + t)$, $t \geq 0$.

Note that $T^*(0) = \tau_n - \beta$. If $\beta = x$, the probability of acceptance in step 4 given $\beta = x$ is $\mathbb{P}(\tau > x)$, so that the expected number of times N the test at step 4 is performed is

$$\int_0^\infty \frac{1}{\mathbb{P}(\tau > x)} \mathbb{P}(\beta \in dx) = \int_0^\infty \frac{1}{\mathbb{P}(\tau > x)} \cdot \frac{1}{m} \mathbb{P}(\tau > x) dx = \int_0^\infty \frac{1}{m} dx = \infty$$

(unless the support of τ is bounded, say a is the supremum, then we get a/m precisely as in Algorithm D). Thus typically Algorithm E has an infinite expected sample size. This indicates that applications that require repeated use of Algorithm E (see, for example, Section 6) will in practice have enormous sample sizes, whereas the problem is less serious if the algorithm is only used once, for example when starting a long simulation run, and is a strictly stationary way to eliminate bias. (It is tempting to circumvent the problem by generating β in step 3 instead, but this idea does not lead to the correct distribution of \mathbf{X}^* .)

Note that the r.v. N does not represent an extreme instance of an r.v. with a heavy-tailed distribution. Suppose for example that τ has a geometric distribution (see Examples 7.4, 7.5) or, more generally, that $\mathbb{P}(\tau > x) \leq e^{-\alpha x}$ for some $\alpha > 0$. By Jensen's inequality,

$$\mathbb{E}[N^p | \beta = x] \leq \mathbb{E}^p[N | \beta = x] = \mathbb{P}(\tau > x)^{-p}$$

for $p < 1$, and hence

$$\begin{aligned} \mathbb{E}N^p &= \int_0^\infty \mathbb{E}[N^p | \beta = x] \mathbb{P}(\beta \in dx) \leq \frac{1}{m} \int_0^\infty \left(\mathbb{P}(\tau > x) \right)^{1-p} dx \\ &\leq \frac{1}{m} \int_0^\infty e^{-(1-p)\alpha x} dx < \infty. \end{aligned}$$

One may also note that once $\beta = x$ has been picked in step 1, the conditional expected sample size $1/\mathbb{P}(\tau > x)$ is finite.

One might initially expect that the only case when the stationary age distribution is simulatable is that where the stationary distribution of \mathbf{X} is known in closed form, in which case one can simulate the stationary version \mathbf{X}^* from this distribution explicitly. This, however, is not the case; see Examples 7.3 and 7.4 for nontrivial applications of Algorithm E.

Algorithm B implies that the class of wide-sense regenerative processes with a given simulatable stationary age distribution is a uniformity class.

It turns out that Algorithms D and E can, in fact, be extended beyond the setting of wide-sense regenerative processes. Variants of these algorithms can be developed for synchronous processes; see Thorisson [28] for details. A synchronous process is one in which the “cycles” form a stationary sequence, but where the dependency structure between cycles can essentially be arbitrary. These processes arise naturally as Palm versions of stationary point processes; see [24]. Because it is not clear to the authors how this would show up in a simulation context, we don’t discuss this further here.

Recent research indicates that the class of regenerative processes such that $\mathbb{P}(\tau > t) \leq g(t)$ for a known integrable function g is a uniformity class, which would be a rather satisfying answer in terms of how few special properties are actually required to generate a stationary version. Our study of this case involves an application of Keane and O’Brien’s ongoing research [18], and will be published elsewhere.

5. WEAK UNIFORMITY CLASSES

We start by showing that, for a certain class of Markov chains, one can calculate a priori estimates on the rate at which the system converges to steady-state, and thereby construct ϵ -stationarity detection times, which are deterministic. In the finite state space setting, estimating the convergence rate essentially amounts to calculating a bound on the eigenvalue of K having the second largest modulus; since this is the parameter that determines the rate at which the n th power of an irreducible transition matrix converges to its limit. In any case, given an upper bound on the rate at which the system converges to steady-state, we can choose a time T so that the total variation distance to the steady-state distribution is arbitrarily small. The deterministic time T can then be used in (2.3) to obtain an appropriate uniform bound on the total variation distance.

We say that a transition kernel K satisfies a (λ, φ, m) minorization if

$$K^m(x, \cdot) \geq \lambda\varphi(\cdot), \quad x \in S. \quad (5.1)$$

Here $0 < \lambda \leq 1$, φ is a probability distribution on S , and $m \geq 1$ an integer ($K^m(x, B)$ denotes the probability that the chain \mathbf{X} moves from x to B in m transitions). Some discussion of condition (5.1) is given in Remark 5.1 below.

The following result is well-known and straightforward to show via coupling (see, e.g., [22]):

PROPOSITION 5.1. *If K satisfies a (λ, φ, m) minorization, then $K \in \mathcal{M}$ and*

$$\sup_{x \in S} \|K^n(x, \cdot) - \pi_K(\cdot)\| \leq (1 - \lambda)^{\lfloor n/m \rfloor}. \quad (5.2)$$

It follows that by choosing n sufficiently large, we can make $(1 - \lambda)^{\lfloor n/m \rfloor}$ arbitrarily small. This immediately proves that one can construct deterministic ϵ -stationarity detection times for the above class of systems.

THEOREM 5.1. *Fix $\lambda > 0$ and $m \geq 1$, and let \mathcal{M}_3 be the family of transition kernels defined on S such that K satisfies a (λ, φ, m) minorization for some φ . Then \mathcal{M}_3 is a weak uniformity class, and $T(\epsilon)$ is an ϵ -stationarity detection*

time when $T(\epsilon)$ is chosen as the least integer n having the property $(1 - \lambda)^{\lfloor n/m \rfloor} < \epsilon$.

We will see later (Example 7.3) that \mathcal{M}_3 is a uniformity class. In order to achieve strict stationarity (rather than ϵ -stationarity), we will use randomized stopping times (rather than the deterministic times of this section).

Remark 5.1. Condition (5.1) is closely related to the *Doebelin condition* studied, for example, in Chapter 5 of [10] (it is easy to show that condition D_0 stated on p. 221 of [10] implies (5.1)). In the discrete state space setting, (5.1) is equivalent to requiring that the m -step transition matrix have a column in which the elements are uniformly bounded away from zero (in the case of a finite S , (5.1) holds for some (λ, φ, m) if and only if the Markov chain is aperiodic and irreducible). Typically, we would expect (5.1) to hold when S is compact and $K(x, A)$ satisfies some continuity requirements [in fact, in (5.1) we may, preliminarily, for a fixed x take $m = m(x)$ where there is positive probability of coupling to the stationary version in m steps starting from x (cf., Lemma 2.2 of [5]); by continuity, the same m will then serve in a neighborhood of x , and by compactness, a finite m will do for all x].

It turns out, that the class of chains for which deterministic detection times work uniformly well over all possible initial distributions ν is precisely described by the set of kernels K satisfying (5.1). This follows by the following partial converse to Theorem 5.1: *suppose that for $0 < \epsilon < 1/2$ there exists a deterministic time $T(\epsilon) = n$ such that $\|\mathbb{P}_{x,K}(X_n \in \cdot) - \pi_K(\cdot)\| < \epsilon$ for each $x \in S$. Then there exists (λ, φ, m) such that K satisfies a (λ, φ, m) minorization.* The proof is easy and therefore omitted.

One difficulty with applying Algorithm D of Section 3 is that it requires that the cycle lengths be bounded. Very few regenerative processes have this property, although the class of (s, S) inventory systems is a notable exception (see Example 7.2 for further details). It is worth noting that, in general, boundedness of the regeneration times does not imply the existence of a deterministic stationarity detection time T ; see [15] for a discussion of the class of chains for which such deterministic times exist.

On the other hand, one might hope that the application of an appropriately derived truncation technique to the cycle length distribution would enable one to use Algorithm E to construct ϵ -stationarity detection times. The development of such a methodology is given in the Appendix, where we show the validity of the following algorithm:

Algorithm F. (Generating an ϵ -stationarity detection time $T(\epsilon)$ when an upper bound γ on $\mathbb{E}\tau^{p+1}$ is known)

1. Calculate $\alpha = \alpha(\epsilon) = (4\gamma/\epsilon^2)^{1/p}$
2. Let $n \leftarrow 1$
3. Generate \mathbf{X} over $[T(n-1), T(n))$
4. If $\tau_n \leq \alpha$, go to 5;
else, let $n \leftarrow n + 1$ and return to 3
5. Generate V_n , a uniform r.v. on $[0, 1]$
6. If $\alpha V_n < \tau_n$, go to 7;
else, let $n \leftarrow n + 1$ and return to 3
7. Generate U , a uniform r.v. on $[0, 1]$, and let $T(\epsilon) \leftarrow T(n-1) + U\tau_n$

Thus, the class of regenerative processes with $\mathbb{E}\tau^{p+1} \leq c$ is a weak uniformity class

6. FURTHER DISCUSSION

An underlying theme of Sections 2 through 5 is that the tail of the distribution of the cycle length τ plays a critical role in whether one can construct ϵ -stationarity times. In particular, we have positive results for

- finite-state Markov chains (stationarity detection r.v.'s are constructed in Algorithms A, B, C);
- wide-sense regenerative processes with bounded cycle lengths (Algorithm D constructs stationarity detection times);
- wide-sense regenerative processes with simulatable stationary age distribution (Algorithm E constructs stationarity detection times);
- wide-sense regenerative processes with bounded $(p + 1)$ th moment of the cycle length (Algorithm F constructs ϵ -stationarity detection times).

On the other hand, our principal negative result (Theorem 2.2) arises in a setting in which no control whatsoever is placed on the behavior of the cycle length distribution.

The critical role of the tail behavior of the cycle length distribution in developing initial transient detection algorithms comes as no surprise, given the intimate relationship between this tail behavior and rates of convergence for regenerative processes (see, for example, [26]).

One obvious question of interest to the simulation analyst is whether the ability to generate a stationary version of the process can be used to obtain a variance reduction in the context of steady-state simulation. In particular, suppose that $\mathbf{X} = \{X(t)\}_{t \geq 0}$ is a real-valued regenerative process in the classical sense (with i.i.d. cycles) for which $\mathbb{E}[\tilde{Y}_1^2 + \tau^2] < \infty$, where $\tilde{Y}_1 = \int_{T_{(i-1)}^+}^{T_i^+} |X(s)| ds$; to avoid trivial cases, assume that \mathbf{X} is not constant ($X(t) \equiv \alpha$). Letting $\alpha_1(t)$ denote the time average (sample mean) $t^{-1} \int_0^t X(s) ds$, it is well known that

$$\alpha_1(t) \xrightarrow{\text{a.s.}} \alpha = \mathbb{E}X^*(0) = \frac{1}{m} \mathbb{E}Y_1 = \frac{1}{m} \mathbb{E} \int_0^\tau X(s) ds.$$

The conventional approach for estimating the steady-state mean α is to use $\alpha_1(t)$ as point estimator, and under the conditions stated,

$$\sqrt{t}(\alpha_1(t) - \alpha) \xrightarrow{\mathcal{L}} \sigma_1 N(0, 1), \quad t \rightarrow \infty, \quad (6.3)$$

where $\sigma_1^2 = \mathbb{E}Z_1^2/m$, $Z_1 = \int_0^\tau (X(s) - \alpha) ds$.

However, the ability to simulate a stationary version of the regenerative process suggests the following alternative estimator, $\alpha_2(t)$. Assume that the cycle lengths are bounded and that Algorithm D is used to find a stationarity detection time $T_1(0)$, then $\mathbb{E}X(T_1(0)) = \mathbb{E}X^*(0) = \alpha$. Let $A_1 = X(T_1(0))$, proceed to the next cycle and execute Algorithm D a second time to produce a second independent copy A_2 of A_1 . In this way, the simulation of \mathbf{X} over

$[0, t]$ will produce a random number $\chi(t)$ of i.i.d. copies $A_1, A_2, \dots, A_{\chi(t)}$ of $X_{T_1(0)}$, and we can therefore let

$$\alpha_2(t) = \begin{cases} \frac{1}{\chi(t)} \sum_{i=1}^{\chi(t)} A_i & \chi(t) \geq 1 \\ 0 & \chi(t) = 0 \end{cases}$$

Under our moment assumptions, it follows that (see [16]) $\alpha_2(t) \xrightarrow{\text{a.s.}} \alpha$, and

$$\sqrt{t} (\alpha_2(t) - \alpha) \xrightarrow{\mathcal{D}} \sigma_2 N(0, 1), \quad t \rightarrow \infty, \quad (6.4)$$

where $\sigma_2^2 = \mathbf{Var} X_{T_1(0)} \cdot \mathbb{E}\tau^\#$, and $\tau^\#$ is the amount of time for which \mathbf{X} needs to be simulated in order that a copy of $X_{T_1(0)}$ can be calculated ($\tau^\#$ is the sum of the cycle lengths up to the first accepted cycle; the notation indicates that $\tau^\#$ may be thought of as a regeneration point; cf., the proof of Proposition 6.2 below). The following result is shown in the Appendix:

PROPOSITION 6.1. *For any regenerative process satisfying the conditions of Algorithm D, one has $\sigma_2^2 > \sigma_1^2$.*

Thus, one never obtains an efficiency improvement by favoring $\alpha_2(t)$ over $\alpha_1(t)$ (note that σ_1^2 also can be interpreted as the variance per time unit in regenerative simulation; hence, Proposition 6.1 also gives a comparison with that method). This may not appear surprising, as $\alpha_2(t)$ throws away a great deal of information that is incorporated in $\alpha_1(t)$. Nevertheless, a more sophisticated idea will indeed produce variance reduction:

PROPOSITION 6.2. *Under the assumptions of Algorithm D, there exists a constant b (depending on \mathbf{X}) such that the estimator $\alpha_3(t) = (1 - b)\alpha_1(t) + b\alpha_2(t)$ satisfies*

$$\sqrt{t} (\alpha_3(t) - \alpha) \xrightarrow{\mathcal{D}} \sigma_3 N(0, 1), \quad t \rightarrow \infty, \quad (6.5)$$

with $\sigma_3^2 \leq \sigma_1^2$. Furthermore, b can be consistently estimated, i.e., there exists $b^*(t)$ such that $b^*(t)$ can be evaluated from the simulation in $[0, t]$ and $b^*(t) \xrightarrow{\text{a.s.}} b$, $t \rightarrow \infty$, and then (6.5) holds with $\alpha_3(t)$ replaced by $(1 - b^*(t))\alpha_1(t) + b^*(t)\alpha_2(t)$.

The evaluation of b and the construction of $b^*(t)$ is given in the proof (based upon a slightly tricky application of linear control variates) in the Appendix. It will also be seen that unless one has a process with a very special dependence structure, the strict inequality $\sigma_3^2 < \sigma_1^2$ holds. The construction is, however, somewhat complicated; it is also the feeling of the authors that the variance reduction that can be obtained in this way will seldom be very substantial. Further, even though the control variate idea carries over to Algorithm E when simulating a fixed number of cycles of generic length $\tau^\#$, the corresponding estimator can never compete with $\alpha_1(t)$ when we discuss efficiency in terms of the simulation run length t , since $\mathbb{E}\tau^\# = \infty$.

As a consequence, we do not believe that the main simulation contribution of this paper lies in the area of variance reduction. Rather, the focus is on the

initial transient problem. The idea is to produce estimators with lower bias, without adversely affecting the asymptotic variance or the length of the simulation. Our results contribute to this by giving insight into how to construct a random time $T(\epsilon)$ such that the post- $T(\epsilon)$ process is close to stationarity. The estimator

$$\alpha_4(t) = \frac{1}{t} \int_0^t X(T(\epsilon) + s) ds$$

will typically have better bias characteristics than $\alpha_1(t)$, without affecting the asymptotic variance (since $\sqrt{t}(\alpha_4(t) - \alpha) \xrightarrow{d} \sigma_1 N(0, 1)$ as $t \rightarrow \infty$) or the order of magnitude of the length of the simulation for large t .

A second important possibility that the results of this paper create is the ability to use simulation to numerically calculate upper bounds on the rate of convergence of a stochastic system to its steady-state. Such upper bounds are currently of great interest to the probability community; see, for example, [3]. As stated earlier, in the finite state Markov chain setting this is tantamount to using simulation to numerically calculate an upper bound on the second eigenvalue λ of the transition matrix of the chain. In this case, λ may often be available by other means, but for even slightly more complicated processes the difficulties are formidable (e.g., for queues this convergence rate is related to the concept of relaxation time, see [4] Ch. III.10). What we can do is to use the method of coupling (see, e.g., [4], Ch. VI.2, for some basic discussion and [22] for a more comprehensive treatment) to numerically calculate the rate of convergence. The idea is to simultaneously simulate both a stationary version \mathbf{X}^* and the nonstationary version \mathbf{X} of the process (the techniques of this paper would be used to generate \mathbf{X}^*), in such a way that the coupling time is finite—by coupling time we mean a random time κ with the property that $X(t) = X^*(t)$, $t \geq \kappa$. For a positive recurrent Markov chain with a finite or countably infinite state space S , we may start by simulating \mathbf{X}^* and \mathbf{X} independently, take κ to be the first n such that $X_n = X_n^*$, and let the processes be identical after κ —whereas for nonMarkovian processes or processes with a continuous component of the state space, slightly more intricate procedures may be needed (see, e.g., Example 6.1 below). In any case, the tail of the distribution of κ gives an estimate of an upper bound on the total variation distance between the distributions of the stationary and nonstationary versions,

$$\|\mathbb{P}(X(t) \in \cdot) - \mathbb{P}(X^*(t) \in \cdot)\| \leq 2\mathbb{P}(\kappa > t), \quad (6.6)$$

and by simulating i.i.d. replications $\kappa_1, \dots, \kappa_N$ of κ , we can calculate empirical bounds on $\mathbb{P}(\kappa > t)$.

The ability to generate a stationary version of the process can have additional benefits as well. For example, one can estimate quantiles of the stationary distribution by generating i.i.d. samples of the process in stationarity (in much the same way as $\alpha_2(t)$ is constructed). The generation of confidence intervals for quantiles in the i.i.d. setting is less complicated and more straightforward than that in the dependent context, although it is likely

that the asymptotic variance of this new estimator is worse than that of the traditional one. More generally, traditional steady-state estimators like $\alpha_1(t)$ only capture features of the one-dimensional distribution, which does not always tell the whole story; see the discussion in Whitt [31].

7. EXAMPLES

Example 7.1. We give some estimates indicating how the run length of the recursive Algorithm C may depend on the number s of states of the Markov chain $\{X_n\}$. A difficulty is that this run length in general appears to depend on the transition probabilities k_{ij} in a complicated way, and so it may be difficult to make comparisons for different values of s . Here we consider an extremely simple example, a chain that goes to any other state with equal probabilities (thus $k_{ij} = 1/(s-1)$ for $i \neq j$, $k_{ii} = 0$). By the “run length” l_s we understand the expected number of steps of $\{X_n\}$ that need to be generated before $Z = Z^{(s)}$ is observed.

Consider first the version of the algorithm where one state at a time is added. That is, in Algorithm C, F is a one-point set. Let n_t denote the run length needed to create $Z^{(A)}$ when A has t elements. Due to the special structure of the k_{ij} , n_t does not depend on A , and obviously $n_1 = 0$, $n_s = l_s$. Now let G have t elements and F one, say i . To create $T^{(G)}$, Algorithm C goes through the states in G in succession, creates one step of $\{X_n\}$ at a time, and observes whether the next value is i . The expected amount of time required for this to turn out successfully is $s-1$. Then $Z^{(G)}$ is generated, and the algorithm stops strictly later than at the time when the observed value of $Z^{(G)}$ is the state from which a transition to i occurred. The probability of this last event is $1/(t-1)$, and thus $n_{t+1} > (s-1 + n_t)(t-1)$, from which it follows that $n_2 \geq s-1$, $n_3 \geq n_2$, $n_4 \geq 2n_3$, $n_5 \geq 3n_2$, and thus

$$l_s = n_s > (s-1)! \quad (7.7)$$

Now assume instead that the recursive step is carried out by letting F and G be of the same order of magnitude. For convenience, let $s = 2^N$ and let m_k denote the run length needed to create $Z^{(A)}$ when A has 2^k elements. Assume that F and G both have 2^k elements. An upper bound on the run length needed to create $T^{(G)}$ is 2^k (the number of states in G) multiplied by the expected number of steps needed to create an event in the Poisson stream with intensity $Z_j^{(G)}\lambda_j^{(F+G)}$. Arguing as above, this number is

$$\left(\frac{2^N - 1}{2^k} + m_k \right) 2^k.$$

Thus

$$m_{k+1} \leq 2 \cdot 2^k \left(\frac{2^N - 1}{2^k} + m_k \right) 2^k + m_k,$$

where the last m_k comes from step 11 of Algorithm C. Letting $m_k^* = \max\{m_k, (2^N - 1)/2^k\}$, it follows that $m_1 \leq 2^{N+2}$, $m_{k+1}^* \leq 2^{2k+5}m_k^*$ and thus

$$l_{2^N} = m_N \leq m_N^* \leq 2^{N(N+1)+3N}m_1^* \leq 2^{N^2+5N+2}.$$

Substituting $s = 2^N$, this upper bound is

$$4s^2 e^{\log(s^{\log s / (\log 2)^2})} = 4s^5 s^{2.08 \log s}$$

so that “doubling up” of states leads to

$$l_s = O(s^{2.09 \log s}), \quad (7.8)$$

which is clearly much better than the lower bound (7.7) obtained by adding one state at a time.

For processes with a special structure for the transition graph, these estimates can be improved somewhat. Assume, for example, that $S = \{1, \dots, s = 2^N\}$ and that it is known that $\{X_n\}$ has birth-death structure. That is, $k_{i,j}$ is only nonzero for $j = i - 1$ or $j = i + 1$ when $i = 2, \dots, s - 1$, when $i = 0$ only for $j = 0$ or $j = 1$, and when $i = s$ only for $j = s$ or $j = s - 1$. Proceeding again by doubling up the number of states, this leads to Algorithm C being applied for F, G neighboring intervals of length 2^k ($k = 1, \dots, N - 1$). However, when generating T_G , we need not search all states of G to watch for a transition to F , but only the state neighboring to F . If all nonzero entries of K are $1/2$, π_K is uniform on S , and we can argue as above to get

$$m_{k+1} = 2(2 + m_k)2^k + m_k.$$

Asymptotically, this is easily seen to lead to $l_s = O(s^{1.05 \log s})$. The exact values for small N are given in the following table:

N	1	2	3	4	5	6	7	8	9	10
s	2	4	8	16	32	64	128	256	512	1024
m_k	4	28	268	4588	151.468	$9.845.548 = 9.8 \cdot 10^6$	$1.3 \cdot 10^9$	$3.3 \cdot 10^{11}$	$1.7 \cdot 10^{14}$	$1.7 \cdot 10^{17}$

A further possibility for reducing the run length is that $\pi(S_i)$ may be known for some partitioning of S , $S = S_1 + \dots + S_q$, so that we can first select one of the S_i with the known probabilities and next generate only $Z^{(S_i)}$. An example would be random environment models with S_i the event that the environment is in state i , which would typically have a known steady-state probability.

Note that the deterministic algorithms (say Gauss elimination) for computing π typically have complexity $O(s^3)$; however, as discussed in the Introduction, they also require more detailed knowledge of the transition function than the simulation analyst may want to impose.

Example 7.2. This example serves two purposes, the first to provide a nonartificial example of a regenerative process with bounded cycle length and the second to illustrate the use of coupling to estimate the rate of convergence

to the steady-state. We considered an (s, S) inventory system with $s = 0$, $S = 1$ and a demand process which is a superposition of a deterministic unit rate and a compound Poisson process with arrival rate β and jumps which are uniformly distributed on $(0, 1/2)$. Thus, \mathbf{X} drifts downwards at a unit rate between jumps and is instantaneously reset to 1 whenever $(-\infty, 0]$ is hit, which may occur either continuously at zero (no jump) or through a jump of size exceeding the present level $X(t)$ of stock. Note that if no arrivals occur in $[u, u + z)$, we have $X(v -) = X(u) - z \pmod{1}$. The cycle length τ may be taken as the time spent between consecutive visits to state 1. The process is Markovian, but simple explicit formulas for the stationary distribution enabling $X^*(0)$ (and thereby \mathbf{X}^*) to be simulated by standard methods, are not available. However, since $\tau \leq 1$, we are in a position to apply Algorithm D. For example, a coupling time between a stationary version \mathbf{X}^* and a version \mathbf{X} with arbitrary initial conditions can be constructed as follows:

Algorithm G. (Generating a coupling time κ of a (s, S) inventory system)

1. Generate the length $\beta = T^*(0)$ of the first cycle $\{X^*(t)\}_{0 \leq t < T^*(0)}$ of \mathbf{X}^* by Algorithm A; let $Y^* \leftarrow 1$.
2. Generate \mathbf{X} over the time interval $[0, T^*(0)]$ by independent simulation; let $Y \leftarrow X(T^*(0))$, $t \leftarrow T^*(0)$.
3. Generate Z , an exponential r.v. with rate β ;
let $t \leftarrow t + Z$, $Y \leftarrow Y - Z \pmod{1}$, $Y^* \leftarrow Y^* - Z \pmod{1}$
4. If $Y < Y^*$, let $a \leftarrow Y^* - \frac{1}{2}$, $b \leftarrow Y$;
else let $a \leftarrow Y - \frac{1}{2}$, $b \leftarrow Y^*$
5. Generate V, V^* , independent r.v.'s on $(0, \frac{1}{2})$;
let $Y \leftarrow Y - V$, $Y^* \leftarrow Y^* - V^*$
6. If $a < Y < b$ and $a < Y^* < b$, let $\kappa \leftarrow t$;
else return to 3

We took $\beta = 1$ and $\beta = 2$, and repeated the experiment 500 times for each value of β . Figure 1 gives the empirical values $\sum_1^{500} I(\kappa_i > t)/500$ of the survival probabilities $\mathbb{P}(\kappa > t)$ giving an upper bound on the rate of convergence to stationarity, cf., (6.6), and shows the expected tendency of slower coupling (lower rate of convergence to stationarity) when $\beta = 1$. A simple measure of this tendency is also the observed empirical means of κ , which were 5.34 and 1.49, respectively. Roughly, we may also conclude that the process has become, for all practical purposes, stationary at time $t = 10$ when $\beta = 1$; whereas the corresponding t -value is probably much higher when $\beta = 2$. Note, however, that coupling methods typically produce only upper bounds and that the interpretation of estimates relating to the coupling epoch κ , for this and other reasons, seem to require some care (in our opinion, not least when the mean $\mathbb{E}\kappa$ is studied!)—no matter whether such results are obtained from theory, as say in [2], [3], and [1] or from simulation experiments as in the present work. Our point here, however is not to discuss these issues, but only to point out that in some cases simulation presents a feasible approach.

Example 7.3. For a simple yet nontrivial case where one can actually simulate a r.v. having the stationary age distribution, cf., Algorithm E, consider the M/G/1 queue. If the queue discipline is FIFO (first in first out),




Figure 1

enough is known about the steady-state distribution to allow a stationary version of the system to be simulated. For example, if the service time is phase-type [23], then so is the steady-state waiting time V ([23]), and hence straightforward to generate as initial value for the simulation. In general, we may use the Pollaczek-Khintchine formula $V \stackrel{\cong}{=} R_1 + \dots + R_N$, where N is geometric with rate ρ (the traffic intensity) and R_1, R_2, \dots are i.i.d. with common density $(1 - B(x))/\mu$, where B is the service time distribution and μ its mean; typically, this density will be simulatable given a specific form of B .

Of course, in the FIFO case the need to simulate at all is questionable. However, if instead we are dealing with some other work-conserving discipline, such as some variant of processor-sharing or priority queueing, it will only rarely be the case that full information on the steady-state distribution is available. Then note that, by work conservation, the cycle length τ has the same distribution as for the FIFO case. To generate an r.v. β having the stationary age distribution as required in step 2 of Algorithm B, one simply starts a stationary FIFO system and let β be its first stationary cycle length, noting that the stationary distributions of the age and the excess life of the cycle are the same.

Example 7.4. Our purpose here is to present a further nontrivial case where one can actually simulate an r.v. having the stationary age distribution, cf., Algorithm E. We take $\{X_n\}$ as a discrete time Harris recurrent Markov process, say with state space S and n -step transition probabilities $K^n(x, A) = \mathbb{P}(X_n \in A | X_0 = x)$ on which we impose the minorization (5.1).

The splitting argument for Harris chains (see [4] for the theoretical background and [13] for the simulation implementation) now states that following each $n = 0, m, 2m, \dots$, we may let a regeneration occur at time $n + m$ w.p. λ . In this way we obtain the distribution of the zero-delayed cycle length τ as geometric on a lattice, $\mathbb{P}(\tau = im) = \epsilon(1 - \epsilon)^{i-1}$, $i = 1, 2, \dots$. From this it follows that the stationary age distribution is given by

$$\mathbb{P}(T^*(0) = im + j) = \frac{\epsilon(1 - \epsilon)^{i-1}}{m}, \quad i = 1, 2, \dots, \quad j = 0, 1, \dots, m - 1,$$

which is of a form allowing for straightforward computer generation. When $m > 1$, this is a wide-sense regenerative process that does not (in general) have i.i.d. cycles (they exhibit dependence).

Example 7.5. Consider the M/G/s/N loss system, that is, an s-server system with Poisson arrivals at rate α , service time distribution B , and allowing at most N customers in the system at any time; new customers arriving when N are present are lost. The queue discipline is work-conserving, but otherwise arbitrary. This model will serve in part as a concrete illustration of the idea of Algorithm E, and in part to illustrate the phenomenon of continuous time but lattice-cycle-length distribution.

For simplicity, we take $N = s$. The states x of $X(t)$ can then be taken to be of the form $(u_1, \dots, u_k, 0, \dots, 0)$, where $k = 0, \dots, s$ is the number of customers present and there are $s - k$ zeros.

We need to impose a mild condition on the tail of B , namely

$$\frac{B(z + y_0)}{B(z)} \geq \delta, \quad z \geq z_0, \quad (7.9)$$

for some $\delta > 0$, $z_0, y_0 < \infty$. For example, (7.9) holds if $\liminf b(z) > 0$ where $b(z)$ is the failure rate of B at z . With $m = z_0 + y_0$, it immediately follows that the probability $B(z + m)/B(z)$ of service termination before m , given z units of service, has been attained and is bounded below by δ for all $z \geq 0$. By a time transformation, we may assume $m = 1$. Let $p(x)$ denote the probability that, starting from $X(0) = x$, no new customers will arrive in $[0, 1]$ and that all customers present at time 0 will have terminated service at time 1. That is,

$$p(x) = e^{-\beta} \prod_{i=1}^k \frac{B(u_i + 1)}{B(u_i)}.$$

It follows immediately that $p(x) \geq e^{-\beta\delta^k} \geq \epsilon$, where $\epsilon = e^{-\beta\delta^s}$ is > 0 and independent of x .

The following algorithm describes how to generate a regenerative cycle, with the cycle length distribution being geometric on the lattice $\{1, 2, \dots\}$, $\mathbb{P}(\tau = i) = \epsilon(1 - \epsilon)^{i-1}$, $i = 1, 2, \dots$:

Algorithm H. (Generating a regenerative geometric cycle τ of a M/G/s/N loss system)

1. Let $t \leftarrow 0$, $x \leftarrow (0, \dots, 0)$.
2. Simulate $\{X(s)\}_{t \leq s \leq t+1}$ starting from $X(t) = x$;
let $P = p(x)$, $X \leftarrow X(t+1)$, $t \leftarrow t+1$.
3. If $X \neq (0, \dots, 0)$, return to 2;
else, go to 4
4. Generate V , a uniform r.v. on $[0, 1]$;
if $V \leq \epsilon/P$, let $\tau \leftarrow t$;
else return to 2.

To simulate the first stationary cycle of the system, construct β in step 1 of Algorithm E by simulating β_0 , W where $\mathbb{P}(\beta_0 = i) = \epsilon(1 - \epsilon)^{i-1}$, $i =$

1, 2, ..., and W is uniform on $(0, 1)$, and putting $\beta \leftarrow \beta_0 - W$. The rest then follows Algorithm E, using Algorithm H to simulate the cycles of the zero-delayed process.

8. APPENDIX: PROOFS

PROOF OF PROPOSITION 2.2. Because π is not a unit point mass distribution, there exists a probability distribution $\nu \neq \pi$ such that π and ν are equivalent (i.e., share the same sets of probability zero). Hence, $\mathbb{P}_{\nu, K}(T = 0) = 1$ so that $\mathbb{P}_{\nu, K}(X_T \in \cdot) = \nu(\cdot)$. Since $\nu \neq \pi$, this implies that T is not a stationarity detection time for K , yielding a contradiction. \square

PROOF OF PROPOSITION 2.3. Assume that Z is a ϵ -stationarity detection rule in the class $\mathcal{N}(K^{(0)})$ associated with the randomized stopping time σ , and choose $0 < \delta < 1/2$ such that $2(1 - \delta) > \epsilon$. Let A be some large integer to be specified later, and define $K^{(\alpha)} \in \mathcal{N}(K^{(0)})$ by

$$k_{ij}^{(\alpha)} = \begin{cases} k_{ij}^{(0)} & i \leq A \\ \alpha + (1 - \alpha)k_{ii}^{(0)} & i > \alpha, i = j, \\ (1 - \alpha)k_{ij}^{(0)} & i > A, i \neq j \end{cases}$$

$0 < \alpha < 1$. For $0 \leq \alpha < 1$, write $\pi^{(\alpha)} = \pi_{K^{(\alpha)}}$ and $\mathbb{P}^{(\alpha)} = \mathbb{P}_{0, K^{(\alpha)}}$ for the $\mathbb{P}_{K^{(\alpha)}}$ -distribution of $\{X_n\}$, starting from $X_0 = 0$, and let $\tau = \inf\{n \geq 1: X_n = 0\}$. Then, letting $M = \max\{X_n: 0 \leq n \leq \sigma\}$, it is easy to see that

$$\mathbb{E}^{(\alpha)}_{\tau} \geq \frac{1}{1 - \alpha} \mathbb{P}^{(0)}(M > A), \quad (8.1)$$

$$\pi_i^{(\alpha)} = \frac{1}{\mathbb{E}^{(\alpha)}_{\tau}} \mathbb{E}^{(\alpha)} \sum_{n=0}^{\tau-1} I(X_n = i) = \frac{1}{\mathbb{E}^{(0)}_{\tau}} \mathbb{E}^{(0)} \sum_{n=0}^{\tau-1} I(X_n = i), \quad i \leq A \quad (8.2)$$

(here (8.1) follows because $K^{(\alpha)}$ adds a number of “self-loops” in states $i > A$). In particular, (8.1) converges to ∞ and (8.2) to 0 as $\alpha \uparrow 1$, and hence, for some α , we have $\pi^{(\alpha)}(\{0, \dots, A\}) < \delta/2$. Choose A such that $\mathbb{P}^{(0)}(Z \leq A, M \leq A) \geq 1 - \delta/2$. Since $\mathbb{P}^{(\alpha)}(Z \leq A, M \leq A) = \mathbb{P}^{(0)}(Z \leq A, M \leq A)$, by construction, we get

$$\begin{aligned} \|\mathbb{P}^{(\alpha)}(Z \in \cdot) - \pi^{(\alpha)}(\cdot)\| &\geq |\mathbb{P}^{(\alpha)}(Z \leq A) - \pi^{(\alpha)}(\{0, \dots, A\})| \\ &= \mathbb{P}^{(\alpha)}(Z \leq A) - \pi^{(\alpha)}(\{0, \dots, A\}) \\ &\geq \mathbb{P}^{(\alpha)}(Z \leq A, M \leq A) - \pi^{(\alpha)}(\{0, \dots, A\}) \\ &\geq 2 \left(1 - \frac{\delta}{2} - \frac{\delta}{2}\right) > \epsilon \end{aligned}$$

(using $\delta < 1/2$ in the second step), a contradiction. \square

Proof that Algorithm F is valid. Suppose that the cycle length τ is not bounded by the constant a , and consider the r.v. $T'_a(0)$ having distribution

$$\mathbb{P}(T'_a(0) \in dx) = \frac{xI(0 \leq x \leq a)\mathbb{P}(\tau \in dx)}{\mathbb{E}[\tau; 0 \leq \tau \leq a]}.$$

We note that the r.v. $T'_a(0)$ can be generated by looking at the subsequence of the τ_i 's for which $\tau_i \leq a$ and then applying to the subsequence the “length-biased” acceptance-rejection technique of Algorithm A (see step 4). If (4.9) is in force, then (for any measurable set B):

$$\begin{aligned} & |\mathbb{P}(\mathbf{X}'(UT'_a(0) + \cdot) \in B) - \mathbb{P}(\mathbf{X}^* \in B)| \\ &= |\mathbb{P}(\mathbf{X}'(UT'_a(0) + \cdot) \in B) - \mathbb{P}(\mathbf{X}'(UT'(0) \in B)| \\ &= \left| \int_0^\infty \mathbb{P}[\mathbf{X}(Ut + \cdot) \in B | \tau_1 = t] (\mathbb{P}(T'_a(0) \in dt) - \mathbb{P}(T'(0) \in dt)) \right| \\ &\leq \int_0^\infty |\mathbb{P}(T'_a(0) \in dt) - \mathbb{P}(T'(0) \in dt)| \\ &= \int_0^a x \left(\frac{1}{\mathbb{E}[\tau; 0 \leq \tau \leq a]} - \frac{1}{m} \right) \mathbb{P}(\tau \in dx) + \int_a^\infty \frac{x}{m} \mathbb{P}(\tau \in dx) \\ &= 1 - \frac{\mathbb{E}[\tau; 0 \leq \tau \leq a]}{m} + \frac{\mathbb{E}[\tau; \tau > a]}{m} \\ &= \frac{2\mathbb{E}[\tau; \tau > a]}{\mathbb{E}\tau}. \end{aligned}$$

We further note that if $\mathbb{E}\tau \geq 1$, then

$$\begin{aligned} & \frac{\mathbb{E}[\tau; \tau > a]}{\mathbb{E}\tau} \\ &= \frac{\mathbb{E}[\tau^{1/2} \cdot \tau^{1/2} I(\tau > a)]}{\mathbb{E}\tau} \\ &\leq \frac{\mathbb{E}^{1/2}\tau \cdot \mathbb{E}^{1/2}[\tau; \tau > a]}{\mathbb{E}\tau} \quad (\text{Cauchy-Schwarz}) \\ &\leq \frac{\mathbb{E}\tau \cdot \mathbb{E}^{1/2}[\tau; \tau > a]}{\mathbb{E}\tau} \quad (\text{since } \mathbb{E}\tau \geq 1) \\ &= \mathbb{E}^{1/2}[\tau; \tau > a], \end{aligned}$$

whereas (8.3) is bounded by $\mathbb{E}[\tau; \tau > a]$ if $\mathbb{E}\tau < 1$. But,

$$\mathbb{E}[\tau; \tau > a] \leq \mathbb{E}[(\tau/a)^p \tau; \tau > a] \leq a^{-p} \mathbb{E}\tau^{p+1}.$$

Suppose that an upper-bound γ on $\mathbb{E}\tau^{p+1}$ can be computed. Then, for any positive $\epsilon < 1$, we can choose $a = a(\epsilon)$ as in step 1 of Algorithm F. By doing

so we guarantee that

$$2 \sup_B |\mathbb{P}(\mathbf{X}'(UT'_a(0) + \cdot) \in B) - \mathbb{P}(\mathbf{X}^* \in B)| < \epsilon. \quad \square$$

PROOF OF PROPOSITION 6.1. Obviously, $\tau^\#$ is the sum of the cycle lengths up to and including the first accepted cycle, and since a/m is the acceptance probability, it follows by Wald's identity that $\mathbb{E}\tau^\# = \mathbb{E}\tau \cdot a/m = a$. Also, it is standard (or seen by a simple conditioning argument based upon Proposition 3.1) that

$$\mathbf{Var} X^*(0) = \mathbf{Var} X_{T_1(0)} = \frac{1}{m} \mathbb{E} \left[\int_0^\tau (X(s) - \alpha)^2 ds \right].$$

Thus,

$$\begin{aligned} \sigma_1^2 &= \frac{\mathbb{E}Z_1^2}{m} = \frac{1}{m} \mathbb{E} \left[\int_0^\tau (X(s) - \alpha) ds \right]^2, \\ \sigma_2^2 &= \mathbf{Var} X^*(0) \cdot \mathbb{E}\tau^\# = \frac{a}{m} \mathbb{E} \int_0^\tau (X(s) - \alpha)^2 ds. \end{aligned}$$

However, let t be fixed and let U be uniform on $[0, t]$. Then

$$\mathbb{E}[(X(U) - \alpha)^2 | X] \geq \mathbb{E}^2[(X(U) - \alpha) | X] \quad \text{conditionally upon } X, \quad (8.5)$$

$$\begin{aligned} \frac{1}{t} \int_0^t (X(s) - \alpha)^2 ds &\geq \left[\frac{1}{t} \int_0^t (X(s) - \alpha) ds \right]^2, \\ \alpha \cdot \mathbb{E} \int_0^\tau (X(s) - \alpha)^2 ds &\geq \mathbb{E} \left[\tau \int_0^\tau (X(s) - \alpha)^2 ds \right] > \mathbb{E} \left[\int_0^\tau (X(s) - \alpha) ds \right]^2, \end{aligned} \quad (8.6)$$

where (8.5) follows by Cauchy-Schwarz (equality in the last step of (8.6) would imply $X(t) = \alpha$ for all $t \in [0, \tau)$, which is impossible because we have excluded the case $X(t) \equiv \alpha$). Inserting (8.6) in (8.4) completes the proof. \square

PROOF OF PROPOSITION 6.2. Let $T^\#(1) = \tau_1^\# = \tau^\#$, let $T^\#(2)$ be the sum of the cycle lengths up to and including the second accepted cycle, $\tau_2^\# = T^\#(2) - T^\#(1)$, and so on. Then the $T^\#(n)$ are regeneration points for \mathbf{X} (obtained by randomized stopping of the initial regeneration points) and there are $\chi(t)$ of them before t . Let $\bar{\cdot}$ denote averaging from 1 to $\chi(t)$, and put

$$Z_i^\# = \int_{T^\#(i-1)}^{T^\#(i)} X(s) ds, \quad \alpha_1^\#(t) = \frac{\bar{Z}^\#}{\bar{\tau}^\#}.$$

Then the $(Z_i^\#, \tau_i^\#, A_i)$, $i = 1, 2, \dots$, are i.i.d., $\alpha_1^\#(t)$ is a regenerative estimator, and it is standard that $\alpha_1^\#(t)$ and $\alpha_1(t)$ are asymptotically equivalent, in

the sense that the difference is $o(\sqrt{t})$. Thus, in the proof, we may replace $\alpha_1(t)$ by $\alpha_1^\#(t)$.

Let $f(Z, \tau) = Z/\tau$, $g(Z, \tau, A) = (1 - b)f(Z, \tau) + bA$ and let f_Z , f_τ and g_Z , g_τ , g_A denote the partial derivatives evaluated at $(\mathbb{E}Z^\#, \mathbb{E}\tau^\#)$ and $(\mathbb{E}Z^\#, \mathbb{E}\tau^\#, \mathbb{E}A)$, respectively. Then

$$\begin{aligned}\alpha_1(t) &= g(\bar{Z}^\#, \bar{\tau}^\#), \quad \alpha_1^\#(t) = f(\bar{Z}^\#, \bar{\tau}^\#, \bar{A}), \\ f(\mathbb{E}Z^\#, \mathbb{E}\tau^\#, \mathbb{E}A) &= g(\mathbb{E}Z^\#, \mathbb{E}\tau^\#, \mathbb{E}A) = \alpha,\end{aligned}$$

and applying the Delta method shows that asymptotically,

$$\begin{aligned}\alpha_1(t) &\approx \alpha + f_Z(\bar{Z}^\# - \mathbb{E}Z^\#) + f_\tau(\bar{\tau}^\# - \mathbb{E}\tau^\#), \\ \alpha_4(t) &\approx \alpha + g_Z(\bar{Z}^\# - \mathbb{E}Z^\#) + g_\tau(\bar{\tau}^\# - \mathbb{E}\tau^\#) + g_A(\bar{A} - \mathbb{E}A) \\ &= \alpha + (1 - b)(f_Z(\bar{Z}^\# - \mathbb{E}Z^\#) + f_\tau(\bar{\tau}^\# - \mathbb{E}\tau^\#)) + b(\bar{A} - \mathbb{E}A).\end{aligned}$$

Interpreting $A - f_Z Z^\# - f_\tau \tau^\#$ as a control on $f_Z Z^\# + f_\tau \tau^\#$, standard results on linear control variates show that taking

$$b = -\frac{\mathbf{Cov}(f_Z Z^\# + f_\tau \tau^\#, A - f_Z Z^\# - f_\tau \tau^\#)}{\mathbf{Var}(A - f_Z Z^\# - f_\tau \tau^\#)}$$

will ensure $\sigma_3^2 \leq \sigma_1^2$, and that we will have $\sigma_3^2 < \sigma_1^2$ unless the covariance in the definition of b vanishes; if for no other reason than that because of the randomization in step 5 of Algorithm D, we find it hard to think of examples where this can happen.

To obtain an estimator $b^\#(t)$ of b , just estimate the 3×3 covariance matrix of $(Z^\#, \tau^\#, A)$ by the empirical covariance matrix (as when doing regression adjustment for linear control variates). Note that

$$f_Z = \frac{1}{\mathbb{E}\tau^\#}, \quad f_\tau = -\frac{\mathbb{E}Z^\#}{\mathbb{E}^2\tau^\#}$$

can be estimated by inserting the empirical means. Insert these estimates in the definition of b . \square

REFERENCES

1. ALDOUS, D. Meeting times for independent Markov chains. *Stoch. Process. Appl.* 38 (1991), 185–193.
2. ALDOUS D., AND DIACONIS, P. Shuffling cards and stopping times *Am. Math.* 93 (1986), 333–348.
3. ALDOUS D., AND DIACONIS, P. Strong uniform times and finite random walks. *Adv. Appl. Math.* 8 (1987), 69–97.
4. ASMUSSEN, S. *Applied Probability and Queues*. New York, 1987.
5. ASMUSSEN, S., AND THORISSON, H. A Markov chain approach to periodic queues. *J. Appl. Probab.* 24 (1987), 215–225.
6. ATHREYA, K. B., AND NEY, P. A new approach to the limit theory of recurrent Markov chains. *Trans. Am. Math. Soc.* 245 (1978), 493–501.
7. BRATLEY, P., FOX, B. L., AND SCHRAGE, L. *A Guide to Simulation*. 2nd ed. Springer, New York, 1987.

8. CHUNG, K. L. *A Course in Probability Theory*. 2nd ed. Academic Press, New York, 1974.
9. CONWAY, R. W. Some tactical problems in digital simulation. *Manage. Sci.* 10 (1963), 47–61.
10. DOOB, J. L. *Stochastic Process*. Wiley, New York, 1953.
11. FISHMAN, G. S. *Principles of Discrete Events Simulation*. Wiley, New York, 1978.
12. GAFARIAN, A. V., ANCKER, JR., C. J., AND MORISAKU, T. Evaluation of commonly used rules for detecting steady-state in computer simulation. *Nav. Res. Logistics Q.* 25 (1978), 511–529.
13. GLYNN, P. W. Regenerative simulation of Harris recurrent Markov chains. Tech. Rep. 18, Dept. of Operations Research, Stanford Univ., 1982.
14. GLYNN, P. W. A GSMP formalism for discrete event systems. *Proc. IEEE* 77, 1989, 14–23.
15. GLYNN, P. W., AND IGLEHART, D. L. Conditions under which a Markov chain converges to its steady-state in finite time. *Probab. Eng. Inf. Sci.* 2 (1988), 377–382.
16. GLYNN, P. W., AND WHITT, W. The asymptotic efficiency of simulation estimators. *Oper. Res.* 40 (1992), 505–520.
17. IGLEHART, D. L. The regenerative method for simulation analysis. In *Current Trends in Programming Methodology*, Chandy and Yeh, Eds., Prentice Hall, Englewood Cliffs, N.J., 1978.
18. KEANE, M., AND O'BRIEN, G. L. A Bernoulli factory. Manuscript, 1992.
19. KELLY, F. P. *Reversibility and Stochastic Networks*. Wiley, New York, 1979.
20. KELTON, W. D., AND LAW, A. M. A new approach for dealing with the startup problem in discrete event simulation. *Oper. Res.* 30 (1982), 641–658.
21. LAW, A. M., AND KELTON, W. D. *Simulation Modeling and Analysis*. 2nd ed., McGraw Hill, New York, 1991.
22. LINDVALL, T. *Lectures on the Coupling Method*. Wiley, New York, to appear, 1992.
23. NEUTS, M. F. *Matrix-Geometric Solutions in Stochastic Models*. Johns Hopkins University Press, Baltimore, Md., 1981.
24. ROLSKI, T. *Stationary Random Processes Associated with Point Processes. Lecture Notes on Statistics 5*. Springer-Verlag, New York, 1981.
25. SCHRUBEN, L. W. Detecting initialization bias in simulation output. *Oper. Res.* 30, 1982, 569–590.
26. THORISSON, H. The coupling of regenerative processes. *Adv. Appl. Probab.* 15 (1983), 531–561.
27. THORISSON, H. Construction of a stationary regenerative process. *Stoch. Proc. Appl.* 1991.
28. THORISSON, H. Time- and cycle stationarity, 1991. In preparation.
29. WELCH, P. D. A graphical approach to the initial transient problem in steady state simulations. In *Proceedings of the 10th IMACS World Congress on Systems, Simulation and Scientific Computation* (Montreal, 1982), pp. 219–221.
30. WELCH, P. D. The statistical analysis of simulation results. In *The Computer Performance Modeling Handbook*, S. S. Lavenberg Ed. Academic Press, New York, 1983.
31. WHITT, W. Untold horrors of the waiting room: What the steady-state distribution will never tell. *Manage. Sci.* 29 (1983), 395–408.
32. WILSON, J. R., AND PRITSKER, A. A. B. A survey of research on the simulation startup problem. *Simulation* 31 (1978), 55–58.
33. WILSON, J. R., AND PRITSKER, A. A. B. Evaluation of start-up policies in simulation experiments. *Simulation* 31 (1978), 79–89.