

Two Approaches to the Initial Transient Problem

Peter W. Glynn*
Dept. of Operations Research
Stanford University
Stanford, CA 94309-4022
e-mail: glynn@leland.stanford.edu

Abstract

This paper describes two different approaches to dealing with the initial transient problem. In the first approach, the length of the “warm-up period” is determined by obtaining analytical estimates on the rate of convergence to stationarity. Specifically, we obtain an upper bound on the “second eigenvalue” of the transition matrix of a Markov chain, thereby providing one with a theoretical device that potentially can give estimates of the desired form. The second approach is data-driven, and involves using observed data from the simulation to determine an estimate of the “warm-up period”. For the method we study, we are able to use a coupling argument to establish a number of important theoretical properties of the algorithm.

1 Introduction

In many applications settings, it is of interest to compute steady-state performance measures. To be specific, suppose that the system under consideration is described by a Markov process $X = (X(t) : t \geq 0)$ living on state space S . For a given $f : S \rightarrow \mathbb{R}$, the steady-state simulation problem is concerned with the estimation of the time-average limit α defined via the law of large numbers

$$\frac{1}{t} \int_0^t f(X(s)) ds \rightarrow \alpha \quad \text{P}_x \text{ a.s.}$$

as $t \rightarrow \infty$ for all $x \in S$ (assuming such a limit exists), where

$$\text{P}_x(\cdot) \triangleq P(\cdot | X(0) = x).$$

*This research was supported by the Army Research Office under Contract No. DAAL03-91-G-0319

Such laws of large numbers hold, in great generality, for Markov processes exhibiting some type of positive recurrence condition. In addition, for such processes, there typically exists a unique probability distribution π , known as the stationary distribution, such that

$$\alpha = \int_S f(x)\pi(dx).$$

Furthermore, π has the property that if $X(0)$ has distribution π , then X is a (strictly) stationary process. It follows that if X is initiated with distribution π , then the sample mean $\alpha(t)$ given by

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

is unbiased as an estimator of α .

Unfortunately, the distribution π is typically unknown and, consequently, it is generally impractical to generate $X(0)$ from π . As a result, bias is induced in $\alpha(t)$ and the initial segment of the simulation may be unrepresentative of steady-state behavior. This introduces certain complications into the estimation problem that do not arise in Monte Carlo environments in which unbiased estimators may easily be constructed. The “initial transient problem” focuses both on the effect of this initial bias, and on developing effective algorithms for mitigating the impact of this bias.

One means of attacking this problem is to note that the positive recurrent Markov processes (satisfying some sort of aperiodicity condition) typically exhibit “total variation convergence” to stationarity, by which we mean that

$$\|P(X_t \in \cdot) - P_\pi(X \in \cdot)\| \rightarrow 0$$

as $t \rightarrow \infty$, where $X_t \triangleq (X(t+s) : s \geq 0)$ is the “post- t ” process, $P_\pi(\cdot)$ is the distribution under which X has initial distribution π , and $\|\cdot\|$ is the total variation norm defined by

$$\|\eta\| = \sup_A |\eta(A)|$$

for any signed measure η . Thus, if one can compute a time s for which

$$\|P(X_s \in \cdot) - P_\pi(X \in \cdot)\| < \epsilon,$$

we have an ϵ -guarantee that the post- s process is close to stationarity, and hence any data collected subsequent to s should have relatively low bias.

The remainder of this paper describes two different approaches to accomplishing this task. In section 2, we establish a new analytical bound for the total variation distance from stationarity that takes explicit advantage of the known transition structure of the system. By contrast, section 3 is concerned with developing a new method for identifying s that is purely data-driven, and takes no explicit advantage of the transition structure of the system being simulated.

2 Upper Bounds on Rates of Convergence to Stationarity

As indicated in section 1, we are concerned here with developing upper bounds on the rate of convergence to stationarity, as described via the total variation norm. For the remainder of this section, we shall assume that X is an aperiodic, irreducible, discrete-time Markov chain with finite state space (although one would expect appropriate analogs in both continuous time and general state space).

We start by noting that for $t \in \mathbb{Z}^+$,

$$\begin{aligned} \|\mathbb{P}(X_t \in \cdot) - \mathbb{P}_\pi(X \in \cdot)\| &= \|\mathbb{P}(X(t) \in \cdot) - \mathbb{P}_\pi(X \in \cdot)\| \\ &\leq \frac{1}{2} \max_x \sum_y |P_{xy}^t - \pi_y|, \end{aligned}$$

where $P = (P_{xy} : x, y \in S)$ is the transition matrix of X . Let π be the unique stationary distribution of X , and let Π be the matrix having all rows identical to π . Since $P\Pi = \Pi P = \Pi^2$, it is evident (via an inductive proof) that for $n \geq 1$,

$$P^n - \Pi = (P - \Pi)^n.$$

Clearly, the rate of convergence of P^n to Π is therefore related to the structure of the eigenvalues of $P - \Pi$. In particular, let $\lambda_1, \lambda_2, \dots, \lambda_d$ be the distinct (complex-valued) eigenvalues of $P - \Pi$, with corresponding (complex-valued) eigenvectors u_1, u_2, \dots, u_d . Note that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log(\|\mathbb{P}(X_t \in \cdot) - \mathbb{P}_\pi(X \in \cdot)\|) \leq \log(\gamma) \quad (1)$$

where $\gamma = \max(|\lambda_i| : 1 \leq i \leq d)$. Thus, a bound on γ yields a bound on the rate of convergence to stationarity.

In view of this, let (λ, u) be an eigenvalue-eigenvector pair corresponding to $P - \Pi$ so that

$$P^n u - \Pi u = \lambda^n u,$$

for $n \geq 1$. In other words,

$$\mathbb{E}_x u(X_n) - \pi u = \lambda^n u(x) \quad (2)$$

for $x \in S$. Assume $\lambda \neq 0$, and set

$$M_n = \lambda^{-n} [u(X_n) - \pi u(1 - \lambda^n)(1 - \lambda)^{-1}].$$

(Clearly, $|\lambda| < 1$ since $(P - \Pi)^n \rightarrow 0$, and thus $(1 - \lambda)^{-1}$ is finite.) We claim that $(M_n : n \geq 0)$ is a martingale with respect to $\mathcal{F}_n = \sigma(X_0, \dots, X_n)$.

To verify this, note that since S is finite, M_n is integrable and adapted to $(\mathcal{F}_n : n \geq 0)$. Furthermore, (2) implies that

$$\begin{aligned} \mathbb{E}_x(M_{n+1}|\mathcal{F}_n) &= \lambda^{-n-1}(\lambda u(X_n) + \pi u - \pi u(1 - \lambda^{n+1})(1 - \lambda)^{-1}) \\ &= \lambda^{-n}(u(X_n) - \pi u(1 - \lambda^n)(1 - \lambda)^{-1}) = M_n, \end{aligned}$$

so $(M_n : n \geq 0)$ is indeed a (complex-valued) martingale.

Let $T(x) = \inf\{n \geq 1 : X_n = x\}$ be the (first) hitting time of x . Set $D_j = M_j - M_{j-1}$ and observe that for $n \geq 1$,

$$\begin{aligned} \mathbb{E}_x M_{T(x) \wedge n} &= \mathbb{E}_x M_0 + \sum_{j=1}^n \mathbb{E}_x D_j I(T(x) \geq j) \\ &= \mathbb{E}_x M_0 + \sum_{j=1}^n \mathbb{E}_x I(T(x) \geq j) \mathbb{E}_x(D_j | \mathcal{F}_{j-1}) \\ &= \mathbb{E}_x M_0. \end{aligned} \tag{3}$$

(Thus, the optional sampling identity continues to hold despite the fact that $(M_n : n \geq 0)$ is not real valued). Note that $M_{T(x) \wedge n} \rightarrow M_{T(x)}$ a.s. (since $T(x)$ is finite valued). Furthermore, if $\mathbb{E}_x |\lambda|^{-T(x)} < \infty$, then the Dominated Convergence Theorem, applied to (3), yields

$$\mathbb{E}_x M_{T(x)} = \mathbb{E}_x M_0,$$

so that

$$\mathbb{E}_x(\lambda^{-T(x)}(u(X_{T(x)}) - \pi u(1 - \lambda^{T(x)})(1 - \lambda)^{-1})) = u(x).$$

But $u(X_{T(x)}) = u(x)$ and $\mathbb{E}_x(1 - \lambda^{T(x)}) \neq 0$ (since $|\lambda|^{-T(x)} \geq |\lambda|^{-1} > 1$), from which it follows that if $\mathbb{E}_x |\lambda|^{-T(x)} < \infty$,

$$u(x) = \pi u(1 - \lambda)^{-1}.$$

Thus, if $\mathbb{E}_y |\lambda|^{-T(y)} < \infty$ for all $y \in S$, evidently we would obtain $u(y) = \pi u(1 - \lambda)^{-1}$ for all $y \in S$. This is a contradiction (as is easily seen by taking π of both sides). Hence, there exists $y \in S$ such that $\mathbb{E}_y |\lambda|^{-T(y)} = +\infty$.

Let $\beta(y) = \sup\{|\lambda| : \mathbb{E}_y |\lambda|^{-T(y)} < \infty\}$ be the radius of convergence of the probability generating function of $T(y)$. We have just shown above that there exists $y \in S$ such that $\beta(y) \leq |\lambda|^{-1}$. So $|\lambda|^{-1} \geq \min(\beta(y) : y \in S)$ or, equivalently,

$$|\lambda| \leq \max(\beta(y)^{-1} : y \in S),$$

yielding the bound

$$\gamma \leq \max(\beta(y)^{-1} : y \in S).$$

We can summarise the above discussion with the following theorem.

Theorem 1 *Let X be a finite state aperiodic irreducible discrete-time Markov chain. Then,*

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log(\|P(X_t \in \cdot) - P_\pi(X \in \cdot)\|) \leq \log(\max_{y \in S} \beta(y)^{-1}),$$

where $\beta(y) = \sup\{z : E_y z^{T(y)} < \infty\}$.

This result bounds the rate of convergence to stationarity, in terms of the rate at which the chain X returns to the various states of S . In certain settings, probability arguments can then be used to a priori dominate the $\beta(y)$'s.

The theorem above complements the many other results that have been developed in recent years to bound the rate of convergence to stationarity; see, for example, Diaconis and Stroock [3], Fill [4], and Meyn and Tweedie [8]. In certain highly structured models, these analytic tools turn out to be quite powerful, and the bounds obtained are relatively tight. However, in general, it is probably fair to say that for unstructured systems, the bounds are often quite loose and consequently of less practical value. In addition, a glance at (1) makes clear that a bound on γ does not necessarily provide a bound on the total variation distance between $P(X_t \in \cdot)$ and $P_\pi(X \in \cdot)$ (although some analytical tools give a bound also on the total variation distance).

Another criticism of the above approach is that for unbounded functions f (as often arise in engineering applications), a bound on the total variation distance does not translate into a bound on $|E f(X(t)) - E_\pi f(X(0))|$, and hence the information obtained about the bias of $\alpha(t)$ is somewhat limited.

3 A Data-Driven Stationarity Detection Rule

In section 2, our concern was with describing upper bounds on the total variation rate of convergence that takes explicit account of the specific transition structure of the process. However, historically, it is fair to say that the most widely used methods for determining the time to stationarity have made such assessments based purely on the observed data obtained by simulating the system itself; see, for example, Conway [2] and Gafarian, Ancker, and Morisaku [5]. A principal difficulty with this approach is that very few such data-driven methods have come equipped with any theoretical guarantees; see, however, Asmussen, Glynn, and Thorisson [1] for some noteworthy exceptions.

Here, we analyze a data-driven rule first proposed in Glynn and Iglehart [6], and show that it enjoys some theoretically important properties. Our goal is to define a non-negative family of random variables $(T(t) : t \geq 0)$ such that:

$$\mathbf{a} \quad T(t) \leq t \text{ a.s.}, \tag{4}$$

$$\mathbf{b} \quad P(T(t) \in \cdot | X) = P(T(t) \in \cdot | X(s) : 0 \leq s \leq t) \text{ for } t \geq 0,$$

c $\|P(X_{T(t)} \in \cdot) - P_\pi(X \in \cdot)\| \rightarrow 0$ as $t \rightarrow \infty$,

d $(T(t) : t \geq 0)$ is a tight family of r.v.'s (under P).

For a simulation time horizon t , we then view $T(t)$ as the epoch at which the process X is in approximate stationarity. Condition **b** above states that $T(t)$ can be generated once X has been simulated to time t , whereas **a** forces $T(t)$ to be in the interval $[0, t]$. Condition **c** asserts that X is in approximate stationarity at time $T(t)$, whereas **d** rules out detection rules that throw out more and more data as $t \rightarrow \infty$ (e.g. $T(t) = t^{1/2}$).

The rule proposed in Glynn and Iglehart [6] is described by the following algorithm:

1. Simulate X to time t .
2. Generate a uniform r.v. U , independent of X .
3. Set $T(t) = \inf\{s \geq 0 : X(s) = X(Ut)\}$.

Clearly $T(t) \leq t$ always holds. Set $Z(t) = X(Ut)$ and note that

$$P(Z(t) \in \cdot | X) = \pi_t(\cdot)$$

where $\pi_t(\cdot)$ is the empirical distribution of X defined by

$$\pi_t(\cdot) = \frac{1}{t} \int_0^t I(X(s) \in \cdot) ds.$$

Hence, (4)**b** also holds. To establish **c**, we need to restrict the class of processes X under consideration. Specifically, suppose that X is an irreducible, positive recurrent, continuous-time Markov chain living on a finite or countably infinite state space. Then, the law of large numbers for such processes guarantees that for each $x \in S$,

$$\pi_t(x) \triangleq \frac{1}{t} \int_0^t I(X(s) = x) ds \rightarrow \pi(x) \text{ a.s.}$$

as $t \rightarrow \infty$, where π is the (unique) stationary distribution of X . Since S is discrete, it then easily follows that

$$\|\pi_t - \pi\| \rightarrow 0 \text{ a.s.} \tag{5}$$

as $t \rightarrow \infty$.

We now use a coupling argument to complete the proof of **c**. For each $t \geq 0$, let $\tilde{Z}(t)$ be an S -valued r.v. having conditional distribution given by

$$P(\tilde{Z}(t) = x | X, U) = \frac{[\pi(x) - \pi_t(x)]^+}{\sum_y [\pi(y) - \pi_t(y)]^+},$$

where $[y]^+ \triangleq y \vee 0$ for $y \in \mathbb{R}$. Let U' be a uniform r.v. independent of x, U , and $\tilde{Z}(t)$ and set

$$Z^*(t) = Z(t)I(U' \leq \frac{\pi(Z(t))}{\pi_t(Z(t))}) + \tilde{Z}(t)I(U' > \frac{\pi(Z(t))}{\pi_t(Z(t))}).$$

Observe that

$$\begin{aligned} \mathbb{P}(Z^*(t) = x | X) &= \pi_t(x) \left(\frac{\pi(x)}{\pi_t(x)} \wedge 1 \right) \\ &+ \sum_{y \in S} \mathbb{E}(\mathbb{P}(\tilde{Z}(t) = x | X, U) [1 - (\frac{\pi(y)}{\pi_t(y)} \wedge 1)] I(Z(t) = y) | X) \\ &= (\pi(x) \wedge \pi_t(x)) + \sum_{y \in S} \mathbb{P}(\tilde{Z}(t) = x | X, U) [\pi_t(y) - \pi(y)]^+ \\ &= \pi(x). \end{aligned}$$

(We have used the easily established fact that

$$\sum_y [\pi(y) - \pi_t(y)]^+ = \sum_y [\pi_t(y) - \pi(y)]^+.)$$

Hence, $Z^*(t)$ is a r.v. having the stationary distribution that is independent of X . Set $T^*(t) = \inf\{s \geq 0 : X(s) = Z^*(t)\}$. Clearly, the aforementioned properties of $Z^*(t)$ imply that

$$\mathbb{P}(X_{T^*(t)} \in \cdot) = \mathbb{P}_\pi(X \in \cdot)$$

for $t \geq 0$. Since $T^*(t) = T(t)$ on $\{Z(t) = Z^*(t)\}$, it follows that

$$\begin{aligned} &|\mathbb{P}(X_{T(t)} \in B) - \mathbb{P}_\pi(X \in B)| \\ &= |\mathbb{P}(X_{T(t)} \in B) - \mathbb{P}(X_{T^*(t)} \in B)| \\ &= |\mathbb{P}(X_{T(t)} \in B, Z(t) = Z^*(t)) + \mathbb{P}(X_{T(t)} \in B, Z(t) \neq Z^*(t)) \\ &\quad - \mathbb{P}(X_{T^*(t)} \in B, Z(t) = Z^*(t)) - \mathbb{P}(X_{T^*(t)} \in B, Z(t) \neq Z^*(t))| \\ &\leq |\mathbb{P}(X_{T(t)} \in B, Z(t) \neq Z^*(t)) - \mathbb{P}(X_{T^*(t)} \in B, Z(t) \neq Z^*(t))| \\ &\leq \mathbb{P}(Z(t) \neq Z^*(t)). \end{aligned}$$

We have therefore established the following coupling inequality:

$$\|\mathbb{P}(X_{T(t)} \in \cdot) - \mathbb{P}_\pi(X \in \cdot)\| \leq \mathbb{P}(Z(t) \neq Z^*(t)).$$

But

$$\begin{aligned} \mathbb{P}(Z(t) \neq Z^*(t) | X) &\leq \mathbb{P}(U' > \frac{\pi(Z(t))}{\pi_t(Z(t))} | X) \\ &= \sum_y [\pi_t(y) - \pi(y)]^+. \end{aligned} \tag{6}$$

On the other hand, the latter sum is just $\|\pi_t - \pi\|$, which (5) asserts goes to zero a.s. The Bounded Convergence Theorem, applied to (6), then yields the conclusion

$$P(Z(t) \neq Z^*(t)) \rightarrow 0$$

as $t \rightarrow \infty$, verifying (4)c. As for condition (4)d, observe that for $x \geq 0$,

$$\begin{aligned} P(T(t) \geq x) &\leq P(T(t) \geq x, Z(t) \neq Z^*(t)) + P(Z(t) \neq Z^*(t)) \\ &\leq P(T^*(t) \geq x) + P(Z(t) \neq Z^*(t)). \end{aligned}$$

But $T^*(t)$ has a distribution independent of t , and is finite-valued. Furthermore, $P(Z(t) \neq Z^*(t)) \rightarrow 0$ as $t \rightarrow \infty$, establishing, for each $\epsilon > 0$, existence of $x = x(\epsilon)$ and $t(\epsilon)$ such that $P(T(t) \geq x) < \epsilon$ for $t \geq t(\epsilon)$. On the other hand, over $[0, t(\epsilon)]$, the non-explosiveness of X guarantees that there exists a finite deterministic set $S_{t(\epsilon)} \subseteq S$ such that

$$P(\pi_{t(\epsilon)}(S_{t(\epsilon)}) = 1) \geq 1 - \epsilon.$$

On the event $\{\pi_{t(\epsilon)}(S_{t(\epsilon)}) = 1\}$ (namely, those outcomes for which X spends the entire interval $[0, t(\epsilon)]$ in $S_{t(\epsilon)}$), $T(u)$ ($0 \leq u \leq t$) is bounded by $\max\{\min\{s \geq 0 : X(s) = y\}, y \in S_{t(\epsilon)}\}$, which is a finite r.v. independent of u . Consequently, we can find $x'(\epsilon)$ for which $P(T(u) > x'(\epsilon)) < 2\epsilon$ for $0 \leq u \leq t(\epsilon)$. This proves the required tightness, and completes the proof of the second major theorem in this paper.

Theorem 2 *If X is an irreducible positive recurrent continuous-time Markov chain taking values in a finite or countably infinite state space, then the algorithm (1)–(3) produces a family of r.v.’s $(T(t) : t \geq 0)$ having properties (4)a–d.*

Note that for a given model, one has no guarantee that a specifically chosen time horizon t will be sufficiently large so that the asymptotics associated with (4)c are in force. While this is clearly a drawback, the same drawback is shared by (for example) most applications of the central limit theorem in a statistical environment (in which one is never certain as to whether the sample size is sufficiently large so as to guarantee a good normal approximation). Note, however, that, as in section 2, bounds on total variation distance do not directly translate into bounds on bias.

Nevertheless, we believe that the algorithm (1)–(3) has sufficient practical merit so as to be worthy of further investigation. Additional properties of this algorithm will be described in a forthcoming paper; see Glynn [7].

References

- [1] Asmussen, S., Glynn, P.W., and Thorisson, H. (1992). Stationarity detection in the initial transient problem. *ACM Transactions on Modeling and Computer Simulation* 2, 130–157.

- [2] Conway, R.W. (1963). Some tactical problems in digital simulation. *Manage. Sci.* 10, 47–61.
- [3] Diaconis, P. and Stroock, D. (1991). Geometric bounds for eigenvalues of Markov chains. *Annals of Applied Probability* 1, 36–61.
- [4] Fill, J.A. (1991). Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *Applied Probability* 1, 62–87.
- [5] Gafarian, A.V., Ancker, C.J., Jr., and Morisaku, T. (1978). Evaluation of commonly used rules for detecting “steady state” in computer simulation. *Nav. Res. Logistics Quart.* 25, 511–529.
- [6] Glynn, P.W. and Iglehart, D.L. (1987). A new initial bias deletion rule. *Proc. of the 1987 Winter Simulation Conference*, 318–319.
- [7] Glynn, P.W. (1995). Asymptotically valid initial transient detection rules. In preparation.
- [8] Meyn, S.P. and Tweedie, R.L. (1994). Computable bounds for geometric convergence rates of Markov chains. *Annals of Applied Probability* 4, 981–1011.