A Martingale Approach to Regenerative Simulation*

Peter W. Glynn
and
Donald L. Iglehart
Department of Operations Research
Stanford University, Stanford, CA 94305

Dedicated to Gerald J. Lieberman in recognition of his leadership, encouragement, and friendship.

Abstract

The standard regenerative method for estimating steady-state parameters is extended to permit cycles which begin and end in different states. This result is established using the Dynkin martingale and a related solution to Poisson's equation. We compare the variance constant which appears in the associated central limit theorem with that arising from cycles that begin and end in the same state. The standard regenerative method has a smaller variance constant than does the alternative.

1 Introduction

Simulation is frequently used to estimate performance measures of complex stochastic systems which defy solution by classical methods of analysis. These measures are generally of two kinds: transient (also called terminating) and steady-state. Our focus in this paper is on steady-state measures of system performance.

*This research was supported by Army Research Office contract DAAH04-94-G-0021.

Estimation of steady-state parameters requires the construction of a point estimate from a simulation of a realization of finite length. The nature of stochastic simulation is such that the convergence of the point estimate to the true parameter is quite slow, of order $t^{-1/2}$, where t is the length of the simulation run. This slow convergence often leads the simulator to seek a confidence interval as well as a point estimate for the parameter of interest.

One of the standard methods for constructing these confidence intervals is regenerative simulation. Regenerative simulation is applicable when the underlying stochastic process is regenerative. This method decomposes the simulated sample path into independent, identically distributed (i.i.d.) cycles of random length. Since the cycles are i.i.d., the classical limit laws (strong law and central limit theorem) can be applied to produce strongly consistent point estimates and confidence intervals. The regenerative cycles for standard regenerative simulation begin and end in the same state, which we shall call the homogeneous regenerative method.

This paper studies cycles that may begin and end in different states, which we shall call the heterogeneous method. These cycles may be used to produce a point estimate and confidence interval for the parameter of interest if they are generated independently and if a ratio formula for the parameter can be obtained. The generation of independent cycles is easy, and the ratio formula is established using the Dynkin martingale and solutions to Poisson's equation.

Organization of this paper is as follows. Section 2 is devoted to a quick review of the standard regenerative method; Section 3 concerns the martingale approach; and Section 4 compares the variance constants arising in the central limit theorem for standard regenerative simulation with that from the martingale approach. The main result here is that standard regenerative simulation leads to smaller confidence intervals than does the martingale approach.

2 Regenerative Simulation

For sake of exposition, we shall restrict our attention in this paper to finite state, continuous time Markov chains (CTMC). Much of what follows can be extended to more general regenerative processes. Let $\mathbf{X} = \{X(t) : t \geq 0\}$ be an irreducible, positive recurrent CTMC with state space S and initial probability measure μ; $\mu(A) \equiv P\{X(0) \in A\}$. Denote by P_μ the probability-measure governing the sample paths of $\mathbf{X}$. Under these conditions, it is well-known that

$$X(t) \Rightarrow X \quad \text{as}\, t \to \infty,$$

where $\Rightarrow$ denotes weak convergence. From the general theory of CTMC's, we know that the distribution of X is independent of the initial measure μ. For a given function $f : S \to \Re$, we are interested in the value of $\alpha \equiv E\{f(X)\} = \sum_{x \in S} f(x)\pi_x$, where $\pi_x \equiv P\{X = x\}$, $x \in S$. Our approach here is to simulate a single sample path of $\mathbf{X}$, and use that data to estimate α.

Before discussing the martingale approach to regenerative simulation, we sketch briefly the ideas behind standard regenerative simulation. Since our CTMC is a regenerative process, we are able to decompose the sample path of $\mathbf{X}$ into independent identically distributed (i.i.d.) regenerative cycles. From these i.i.d. cycles we are able to construct a strongly consistent point estimate for α and a central limit theorem (CLT) which forms the basis for constructing confidence intervals for α.

To be specific, we select a state in E, call it state 0, to be our "return" state. Let $X(0) = 0$ and $T_0 = 0$. The sequence $0 < T_1 < T_2 < \cdots$ are the successive times the process $\mathbf{X}$ returns to 0. We know there will be an infinite sequence of such times. Let $\{\tau_n : n \geq 1\}$ denote the cycle lengths, where $\tau_n = T_n - T_{n-1}$, next we define a sequence of random variables $\{Y_n : n \geq 1\}$ by

$$Y_n \equiv \int_{T_{n-1}}^{T_n} f(X(s))\, ds, \quad n \geq 1.$$

If f is thought of as the cost per unit time for operating the process $\mathbf{X}$, then Y_n is the cumulative cost in the nth cycle. Again the regenerative structure assures that the sequence $\{Y_n : n \geq 1\}$ is i.i.d. The last fact we need is the so-called ratio formula

$$\alpha = \frac{E_0\{Y_1\}}{E_0\{\tau_1\}} \tag{1}$$

which can be proved using the key renewal theorem. This ratio formula holds irrespective of which state was selected for the return state.

There are two approaches to constructing point estimates and confidence intervals for α: use of the simulation data in the interval $[0, t]$, or use of n

regenerative cycles. We shall use only the first approach. The natural point estimate for α is

$$\alpha(t) \equiv \frac{1}{t}\int_0^t f[X(s)]\,ds.$$

From general theory for CTMC's, it is known that

$$\alpha(t) \to \alpha, \quad \text{a.s.} \tag{2}$$

To construct confidence intervals for α we need a CLT in which α appears as a translation constant. The associated CLT (proved using regenerative process theory) is

$$t^{1/2}(\alpha(t) - \alpha)/(\sigma(0)/E^{1/2}\{\tau_1\}) \Rightarrow N(0,1), \quad \text{as } n \to \infty, \tag{3}$$

where $\sigma^2(0) = E_0\{(Y_1 - \alpha\tau_1)^2\}$. Finally, we note that in (3) the constant $\sigma \equiv (\sigma(0)/E^{1/2}\{\tau_1\})$ must be independent of the return state selected since the other terms appearing in (5) do not depend on the return state.

For background information and details on regenerative simulation see Iglehart (1978).

3 A Martingale Approach

Let $Q = \{q_{ij} : i,j \in E\}$ denote the matrix of infinitesimal rates associated with the CTMC $\mathbf{X}$. For all functions u, define the process

$$M(t) \equiv u(X(t)) - u(X(0)) - \int_0^t Qu(X(s))\,ds, \quad t \geq 0. \tag{4}$$

It is well-known (c.f. Karlin and Taylor [1981], p. 310) that $\{Y(t) : t \geq 0\}$ is a P_μ-martingale called the Dynkin-martingale. Next let g be the solution of Poisson's equation

$$Qg = -(f - \alpha e),$$

where e is the vector all of whose components equal 1. This solution vector g is known to exist for our CTMC $\mathbf{X}$. All solutions are of the form $g = (\Pi - Q)^{-1}f + ce$, where c is a constant; see Glynn (1984) for existence of the inverse matrix $(\Pi - Q)^{-1}$. Now replace u in (4) by g and rewrite $M(t)$ as

$$M(t) = g(X(t)) - g(X(0)) + \int_0^t f_c(X(s))\,ds, \quad t \geq 0, \tag{5}$$

where $f_c(x) = f(x) - \alpha$.

Let T be the hitting time of a (possibly randomized) state. Since the state space of $\mathbf{X}$ is finite and $E(T)$ is finite, the optimal stopping theorem applied to the martingale $\{M(t) : t \geq 0\}$ yields

$$\begin{aligned} E_\mu\{M(T)\} &= E_\mu\{g(X(T))\} - E_\mu\{g(X(0))\} + E_\mu\{\int_0^T f_c[X(s)]\,ds\} \\ &= E_\mu\{M(0)\} = 0. \end{aligned}$$

If we now select T so that $X(T) \stackrel{\mathcal{D}}{=} \mu$, then we obtain the ratio formula

$$\alpha = \frac{E_\mu\{\int_0^T f[X(s)]\,ds\}}{E_\mu\{T\}}.$$

This ratio formula is a generalization of (1). Using this ratio formula and i.i.d. cycles based on the stopping time T, the CLT (3) again holds with variance constant given by

$$\sigma_T^2 \equiv E_\mu\{(\int_0^T f_c[X(s)]ds)^2\}/E_\mu\{T\}.$$

There are several ways to select the initial probability measure μ and the stopping time T. To this end, we define the first entrance time to state $x \in E$:

$$T(x) \equiv \inf\{s > 0 : X(s-) \neq x, \quad X(s) = x\}, \quad x \in E.$$

We outline two alternatives. More details will be given in Section 4. Let

(i) μ be a general probability mass function (pmf) on S. Next define the sets

$$A(x) = \{\omega : X(0, \omega) = x\}, \quad x \in S.$$

Since $\bigcup_{x \in S} A(x) = \Omega$, we can define our stopping time T as

$$T(\omega) = T(x, \omega) \quad \text{for} \quad \omega \in A(x). \tag{6}$$

This amounts to a simple randomization over the initial state followed by the standard regeneration method stopping time. Each regenerative cycle begins by generating an independent realization of the initial state using the probability measure μ. For each initial state so generated, the CLT (3) holds with the same variance constant which appears in (3).

(ii) Now let ν be a joint probability mass function on $S \times S$ each of whose marginals has measure μ. (Note that alternative (i) is a special case of this one.) In particular, there may be some advantage in selecting the components to be dependent. Now generate a realization according to ν and use the first component to determine $X(0)$ and the second component to determine the terminal state of the cycle. A sequence of independent cycles of this type would then be generated producing an i.i.d. sequence of pairs $\{(Y_n, \tau_n) : n \geq 1\}$ from which a CLT similar to (3) could be developed. Computation of the variance constant for this CLT is carried out in Section 4.

4 Comparison of Alternatives

Our goal in this section is to compare variance constants appearing in the CLT's associated with the two alternatives sketched at the end of the last section. To this end we define, for $x \in S$,

$$Z(x) = \int_0^{T(x)} f_c(X(s))ds.$$

From the ratio formula (for return state x) $E_x\{Z(x)\} = 0$, and from the remark below (3)

$$E_x\{Z(x)^2\} = \sigma^2 E_x\{T(x)\}, \quad x \in S. \tag{7}$$

(i) In this alternative, μ is an arbitrary pmf on S. The stopping time, T, is defined by (6). By conditioning on the initial state we obtain

$$E_\mu\{(\int_0^T f_c(X(s))ds)^2\} = \sum_{x\in S} \mu(x) E_x\{Z(x)^2\} \tag{8}$$

and

$$E_\mu\{T\} = \sum_{x\in S} \mu(x) E_x\{T(x)\}. \tag{9}$$

Combining (7), (8), and (9), we find that the variance constant in this case is

$$\begin{aligned} & E_\mu\{(\int_0^T f_c(X(s))ds)^2\}/E_\mu\{T\} \\ = & \sum_{x\in S}\mu(x)E_x\{Z^2(x)\}/\sum_{x\in S}\mu(x)E_x\{T(x)\} \\ = & \sigma^2\sum_{x\in S}\mu(x)E_x\{T(x)\}/\sum_{x\in S}\mu(x)E_x\{T(x)\} = \sigma^2. \end{aligned}$$

So for this alternative, we conclude that the variance constant remains σ^2.

(ii) Now let ν be a p.m.f. each of whose marginals has p.m.f. μ. The term $\nu(x,y)$ is the probability of the cycle starting in state x and ending in state y. The variance constant in this case is

$$\begin{aligned} & E_\nu\{(\int_0^T f_c(X(s)ds)^2\}/E_\nu\{T\} \\ = & \sum_{x,y\in S}\nu(x,y)E_x\{Z^2(y)\}/\sum_{x,y}\nu(x,y)E_x\{T(y)\}. \end{aligned}$$

We begin by studying

$$E_x\{Z(y)\} = E_x\{\int_0^{T(y)} f_c(X(s))ds.\}.$$

From (5) we know that

$$\int_0^{T(y)} f_c[X(s)]ds = M(T(y)) - g(X(T(y))) + g(X(0)).$$

Taking expectations and applying the optimal sampling theorem we have

$$\begin{aligned} E_x\{Z(y)\} &= E_x\{M(T(y))\} - g(y) + g(x) \\ &= E_x\{M(0)\} - g(y) + g(x) \\ &= g(x) - g(y) \end{aligned}$$

Now averaging over all possible (x,y)-cycles, we have

$$\sum_{x,y}\nu(x,y)E_x\{Z(y)\}$$

$$
\begin{aligned}
&= \sum_{x,y} \nu(x,y)[g(x) - g(y)] \\
&= \sum_{x} \mu(x,y)g(x) - \sum_{y} \mu(x,y)g(y) = 0.
\end{aligned}
$$

So the relevant Z variable for this case has mean 0. To evaluate the variance constant we only need to study the second moment of $Z(y)$:

$$
\begin{aligned}
E_x\{Z(y)^2\} &= E_x\{(M(T(y)) - g[X(T(y))] + g[X(0)])^2\} \\
&= E_x\{[M(T(y)) - g(y) + g(x)]^2\} \\
&= E_x\{M(T(y))^2\} + 2(g(x) - g(y)) \cdot E_x\{M(T(y))\} \\
&+ (g(x) - g(y))^2. \\
&= E_x\{M(T(y))^2\} + (g(x) - g(y))^2.
\end{aligned}
$$

Next we note that

$$
M(T(y)) = \lim_{n\to\infty} \sum_{k=t}^{\lceil nT(y)\rceil} [M(k/n) - M((k-1)/n)];
$$

recall that $M(0) = 0$ and $\{M(t) : t \geq 0\}$ is right-continuous. Now squaring and taking expectations we have

$$
E_x\{M(T(y))^2\} = E_x\{\lim_{n\to\infty} (\sum_{k=1}^{\lceil nT(y)\rceil} [M(k/n) - M((k-1)/n)])^2\}. \tag{10}
$$

Since S is finite, $E_x\{T^2(y)\} < \infty$. This allows us to use dominated convergence to interchange expectation and limit on the right-hand side of (10). Using the fact that $\{M(t) : t \geq 0\}$ is a martingale, $E_x\{T^2(y)\} < \infty$, and S is finite, we can write

$$
E_x\{M(T(y))^2\} = E_x\{\lim_{n\to\infty} \sum_{k=1}^{\lceil nT(y)\rceil} [M(k/n) - M((k-1)/n)]^2\}.
$$

Now using a path-wise argument, we have

$$
E_x\{M(T(y))^2\} = E_x\{\sum_{k=1}^{\delta(y)} (g(Y_k) - g(Y_{k-1}))^2\},
$$

where $\{Y_n : n \geq 0\}$ is the discrete time embedded Markov chain and $\delta(y) = \inf\{n > 0 : Y_n = y\}$.

Now for $x, z \in S$, define

$$u(x,z) \equiv E_x\{\sum_{k=1}^{\delta(z)}[(g(Y_k) - g(Y_{k-1}))^2 - \sigma^2\beta_{k-1}]\}.$$

When $x = z$,

$$u(z,z) = E_z\{Z(z)^2 - \sigma^2 E_z\{T(z)\} = 0, \tag{11}$$

by virtue of (7). In preparation for the comparison of the variance constants associated with the martingale approach and standard regenerative simulation we need two lemmas.

Lemma 1. For $x, z \in S$, $u(x,z) + u(z,x) = 0$.

Proof. For the embedded Markov chain, $\{Y_k : k \geq 0\}$, let γ denote the first time the chain enters state x after hitting state z, and let N denote the number of x-cycles between time 0 and time γ. Also let V_i denote the sum of the $[(g(Y_{k+1}) - g(Y_n))^2 - \sigma^2\beta_k]$'s over the ith x-cycle. Then using the strong Markov property, Wald's equation, and (11), we have

$$\begin{aligned} u(x,z) + u(z,x) &= E_x\{\sum_{k=1}^{\gamma}[(g(Y_k) - g(Y_{k-1}))^2 - \sigma^2\beta_{k-1}]\} \\ &= E_x\{\sum_{i=1}^{N} V_i\} \\ &= E_x\{N\} \cdot E_x\{V_i\} \\ &= E_x\{N\} \cdot u(x,x) = 0. \end{aligned}$$

Lemma 2. For $x, y, z \in S$,

$$u(x,y) + u(y,z) = u(x,z).$$

Proof. Using Lemma 1, we need to show that

$$u(x,y) + u(y,z) + u(z,x) = 0.$$

Let γ' denote the first time the embedded chain hits x after first hitting y and then hitting z, and let N' denote the number of x-cycles between time 0 and time γ'. From here the proof is the same as Lemma 1 by replacing γ with γ' and N with N'.

With these results in hand, we proceed to the main result.

Theorem 1. For any pmf ν on $S \times S$ having both marginal pmf's equal to μ,

$$\sigma^2 \leq \frac{\sum_{x,y} \nu(x,y) E_x\{Z(y)^2\}}{\sum_{x,y} \nu(x,y) E_x\{T(y)\}}.$$

Proof. From Lemma 2, we obtain

$$\begin{aligned} E_x\{\sum_{k=0}^{\delta(y)-1} (g(Y_{k+1}) - g(Y_k))^2\} &= \sigma^2 E_x\{\sum_{k=0}^{\delta(y)-1} \beta_k\} - u(y,z) + u(x,z) \\ &= \sigma^2 E_x\{T(y)\} - u(y,z) + u(x,z). \end{aligned}$$

Hence

$$\begin{aligned} & \sum_{x,y} \nu(x,y) E_x\{Z(y)^2\} = \\ & \sum_{x,y} \nu(x,y)[\sigma^2 E_x\{T(y)\} - u(y,z) + u(x,z) + (g(x) - g(y))^2] \\ = \quad & \sigma^2 \sum_{x,y} \nu(x,y) E_x\{T(y)\} + \sum_{x,y} \nu(x,y)(g(x) - g(y))^2 \\ > \quad & \sigma^2 \sum_{x,y} \nu(x,y) E_x\{T(y)\}, \end{aligned}$$

the desired inequality.

Thus we conclude from Theorem 1, that the variance constant σ^2 associated with the standard regenerative method is strictly smaller than that obtained by using cycles beginnning in state x and ending in state y with probability $\nu(x,y)$. So while the martingale approach is interesting, it does not lead to shorter confidence intervals.

5 References

Glynn, P.W. (1984). Some asymptotic formulas for Markov chains with applications to simulation. *J. Statistical Comput. Simul.* **19** 97-112.

Iglehart, D.L. (1978). The regenerative method for simulation analysis. Chapter 2 in *Current Trends in Programming Methodology.* Volume III, *Software Modeling.* K.M. Chandy and R.T. Yeh, eds. Prentice-Hall, Englewood Cliffs, NJ.

Karlin, S. and Taylor, H.M. (1981). *A Second Course in Stochastic Processes.* Academic Press, New York.