DISCRETE-TIME CONVERSION FOR SIMULATING SEMI-MARKOV PROCESSES

Bennett L. FOX *

Département d'informatique et de Recherche Opérationelle, Université de Montréal, C.P. 6128, Succursale A, Montréal, Que., Canada H3C 3J7

Peter W. GLYNN **

Mathematics Research Center and Department of Industrial Engineering, University of Wisconsin-Madison, WI 53706, USA

Received June 1985
Revised July 1986

We simulate long-run averages of time integrals of a recurrent semi-Markov process efficiently. Converting to discrete-time by simulating only an imbedded chain and computing the conditional expectations of everything else needed given the sequence of states visited reduces asymptotic variance, eliminates generating holding-time variates, and (when advantageous) gets rid of the future event schedule. In this setting, uniformizing continuous-time Markov chains is not worthwhile. We generalize beyond semi-Markov processes and cut ties to regenerative simulation methodology. Implementation of discrete-time conversion is discussed. It often requires no more work or even less work than the naive method. We give sufficient conditions for work savings. Continuous-time Markov chains, for example, satisfy them.

simulation • variance reduction • semi-Markov processes

1. Introduction

Let X be a positive recurrent semi-Markov process and f be a real-valued function. We want to estimate the time average

$$r = \lim_{T\to\infty}\left(\frac{1}{T}\right)\mathrm{E}\int_0^T f(X(t))\mathrm{d}t, \qquad (1)$$

assuming it exists. This covers most steady-state simulations. Section 3.1 starts with a standard regenerative simulation approach an then converts to discrete time by conditioning on the sequence $Y = (Y_0, Y_1, \ldots)$ of states visited in an imbedded Markov chain. This reduces variance and also, often, the work to simulate. Hordijk, Iglehart and Schassberger [7] adopt the same approach for the special case of continuous-time Markov chains, except that they prove variance reduction by explicit calculation without mentioning the general principle that computing conditional expectations reduces variance. For a discussion of this principle, see, for example, Bratley, Fox and Schrage [2, Sect. 2.6]. Fox and Glynn [5] obtain an analog of these results for certain *finite-horizon* semi-Markov processes. Uniformization pays in [5] but not here as shown in [7] and, with less effort, in Section 3. We refer the reader to Heyman and Sobel [6, Sect. 8.7], for a description of uniformization.

Section 3.2 considers more general processes and cuts ties to the regenerative approach.

Roughly speaking, the total variance in the naive estimator of r is the variance induced by Y plus holding-time variances. Sometimes their relative importance can be assessed *a priori*. Discrete-time conversion eliminates the need to generate holding times explicitly, thereby eliminating their contributions to total variance. The work for discrete-time conversion is sometimes less, sometimes more than for the naive estimator. Section 2 points out that the work for discrete-time conversion can be significantly more than for the naive method

* Research supported by the Natural Sciences and Engineering Research Council of Canada.
** Research supported by the United States Army under contract no. DAAG29-80-C-0041 and by NSF Grant ECS-840-4809.

only if computing conditional expectations given Y is hard. We believe that such cases are atypical. For continuous-time Markov chains, computing these conditional expectations is trivial. Using Erlang's method and its generalizations (Neuts [8]) widens the scope of applicability of continuous-time Markov chains. Section 2 goes on to give widely-satisfied conditions which imply that discrete-time conversion saves work. Coupled with its sure variance reduction, work reduction makes it a certain winner.

2. Implementing discrete-time conversion

In ordinary simulations, without conversion to discrete time, Y is simulated implicitly using a future event schedule (FES) without explicit reference to its transition probabilities. The FES depends also on continuous variates such as holding times. Maintaining it often eats a major chunk of computer time. Even with discrete-time conversion, we could still use a FES to generate Y subsequently replacing all continuous variates by their conditional expectations given Y; however, at least when Y is a Markov chain, we generally prefer to get rid of the FES by using explicit transition probabilities to generate Y.

How do we do the latter efficiently? Let $n(Y_i)$ be the number of possible values of Y_{i+1}. Even when the state space is infinite, $n(Y_i)$ is often finite for all i. When the state space is large, we need a (transition) matrix generator to produce rows as needed along with auxiliary information to expedite variate generation. Suppose $n(Y_i)$ is finite. Given the positive elements of row Y_i and their respective locations, Y_{i+1} can be generated in (small) constant time independent of $n(Y_i)$ using the auxiliary information passed by the matrix generator. For example, this information could be a bucket system (aka: guide table, index table) to localize the search for Y_{i+1} or alias table for row Y_i; e.g., see Bratley, Fox and Schrage [2, Sects. 5.2.1 and 5.2.8, resp.]. Recoding a multidimensional state space as a (one-dimensional) vector, using standard devices, may be desirable.

If computing conditional expectations given Y takes negligible time, then implementing discrete-time conversion with a FES certainly beats naive estimation – though using a matrix generator may do even better. The generator could pass $f(Y_{i+1})$ times the (typically-precomputed) mean holding time in Y_{i+1} along with Y_{i+1} itself. If computing conditional expectations given Y is not trivial, does discrete-time conversion require less work than straightforward simulation? Roughly speaking, the answer is *yes* if

(i) computing conditional expectations given Y is no harder than generating the corresponding variates,

(ii) generating Y via a matrix generator is easier than generating it via a FES.

Condition (i) holds for continuous-time Markov chains X and all f's in (1). Condition (ii) usually holds when the transition matrix is highly structured, a condition satisfied by reasonable models. In toy examples such as the $M/M/s$ queue these conditions clearly hold, but in complex situations (general X and f) pilot runs are needed to assess relative work and amount of variance reduction. This is worthwhile if straightforward simulation would require thousands of runs to get reasonably accurate estimates. In the regenerative setting, a run corresponds to a cycle.

Discrete-time conversion is not always better and not always worse than straightforward simulation when relative work is considered. No universal conclusions can be drawn from empirical comparisons.

3. Discrete-time conversion reduces variance

In Section 3.1 we sketch how to generalize results of [7] to regenerative semi-Markov processes and at the same time indicate how to simplify proofs. Section 3.2 gives a normal proof and further generalizes the class of processes considered in [7]; more importantly, it allows non-regenerative approaches.

In Section 3, we limit the discussion to a comparison of asymptotic variances. To estimate these variances and construct confidence intervals, see, for example, Bratley, Fox and Schrage [2, Ch. 3]. Standardized time series can also be used, as in Schruben [9], to construct confidence intervals; there, the variance constant acts as a scaling parameter.

3.1. Regenerative framework

Let i be a (convenient) recurrent state and let S_j be the time at which state i is visited for the jth

time. We assume that state i is visited at time 0 and set $S_0 = 0$. Put

$$\tau_j = S_{j+1} - S_j, \tag{2}$$

the length of the jth regeneration cycle. Following, for example, the development in Bratley, Fox and Schrage [2, Sect. 3.7], define

$$V_j = \int_{(S_j,\, S_{j+1})} f(X(t))\mathrm{d}t, \tag{3}$$

$$\bar{V}_n = \frac{1}{n}\sum_{j=0}^{n-1} V \tag{4}$$

$$\bar{\tau}_n = \frac{1}{n}\sum_{j=0}^{n-1} \tau_j, \tag{5}$$

$$D_k = V_k - r\tau_k, \tag{6}$$

$$\hat{r}_n = \bar{V}_n/\bar{\tau}_n, \tag{7}$$

$$\tilde{\tau}_n = \mathrm{E}[\bar{\tau}_n \mid Y], \tag{8}$$

$$\tilde{r}_n = \mathrm{E}[\bar{V}_n \mid Y]/\tilde{\tau}_n. \tag{9}$$

We reach state Y_j at time T_j. Let

$$\alpha_k = T_{k+1} - T_k, \tag{10}$$

$$P(\alpha_k \in \mathrm{d}t \mid Y) = F(Y_k, Y_{k+1}, \mathrm{d}t), \tag{11}$$

$$\mu(x, y) = \int_0^\infty zF(x, y, \mathrm{d}z). \tag{12}$$

We define N_n to be the number of jumps to complete n cycles, as jump N_n ends, we visit state i for the nth time. This gives

$$\tilde{\tau}_n = \frac{1}{n}\sum_{j=0}^{N_n-1} \mu(Y_j, Y_{j+1}), \tag{13}$$

$$\mathrm{E}[\bar{V}_n \mid Y] = \frac{1}{n}\sum_{j=0}^{N_n-1} f(Y_j)\mu(Y_j, Y_{j+1}), \tag{14}$$

so we can compute $\tilde{r}_n$ in (9).

Since X is a semi-Markov process (by assumption), the α_k's are conditionally independent given Y. So Y regenerates implies X regenerates, by (11). Assuming certain mild moment conditions (inequality (24) will do) and mimicking standard arguments, we get

$$\sqrt{n}\,(\hat{r}_n - r) \Rightarrow (\delta/\mathrm{E}\tau_1)\mathrm{N}(0, 1), \tag{15}$$

$$\sqrt{n}\,(\tilde{r}_n - r) \Rightarrow (\sigma/\mathrm{E}\tilde{\tau}_1)\mathrm{N}(0, 1), \tag{16}$$

where

$$\delta^2 = \mathrm{Var}\, D_k, \tag{17}$$

$$\sigma^2 = \mathrm{Var}(\mathrm{E}[D_k \mid Y]). \tag{18}$$

Using the standard variance-reducing property of conditional expectation, we get $\sigma \leqslant \delta$. Clearly, $\mathrm{E}\tau_1 = \mathrm{E}\tilde{\tau}_1$, so $\tilde{r}_n$ has smaller variance than $\hat{r}_n$.

3.2. Generalization and formal proof

Now we neither assume that Y is a Markov chain nor that the α_k's are conditionally independent given Y. The only structural assumptions relating Y and X are (11) and (19). A referee pointed out that (11) can be relaxed to $P[\alpha_k \in \mathrm{d}t \mid \quad F[Y_0, \ldots, Y_{k+1}, \mathrm{d}t]$ with (12) replaced by $\mu(x_0, \ldots, x_{k+1}) = \int_0^\infty zF(x_0, \ldots, x_k, \mathrm{d}z)$ and $\mu(Y_k, Y_{k+1})$ replaced by $\mu(Y_0, \ldots, Y_{k+1})$ in (c) below and elsewhere. With these provisos, define Y_k, α_k and $\mu(x, y)$ as before. Let I be an indicator, $T_0 = 0$, $T_k \uparrow \infty$, and

$$X(t) = \sum_{k=0}^{\infty} Y_k I(T_k \leqslant t < T_{k+1}). \tag{19}$$

Thus, X is more general than a semi-Markov process.

The obvious estimator is

$$\hat{R}_n = \frac{1}{T^n}\int_0^{T_n} f(X(s))\mathrm{d}s. \tag{20}$$

Typically, the expected work to generate $\hat{R}_n$ or $\hat{r}_n$ grows at rate n whether n indexes 'transitions' as in this section or cycles as in Section 3.1. So we compare efficiencies of various estimators with respect to n.

Roughly speaking, to convert to discrete time we compute conditional expectations given Y. The alternative estimator

$$\tilde{R}_n = \sum_{k=0}^{n-1} f(Y_k)\mu(Y_k, Y_{k+1}) \Big/ \sum_{k=0}^{n-1} \mu(Y_k, Y_{k+1}) \tag{21}$$

is certainly plausible. The only difference between $\hat{R}_n$ defined by (20) and $\hat{r}_n$ defined by (7) lies in what n indexes; likewise, for $\tilde{R}_n$ and $\tilde{r}_n$. We show that $\tilde{R}_n$ beats $\hat{R}_n$ in a precise sense, provided that $\tilde{R}_n$ is no harder to compute than $\hat{R}_n$.

We assume that X satisfies

(a) $T_n/n \Rightarrow \beta > 0$ (β constant).

(b) there are finite constants r_1 and σ_1 such that

$$\frac{1}{\sqrt{n}}\int_0^{T_n} f(X(s) - r_1)\,\mathrm{d}s \Rightarrow \sigma_1 \mathrm{N}(0, 1).$$

(c) there are finite constants r_2 and σ_2 such that

$$\frac{1}{\sqrt{n}}\sum_{k=0}^{n-1}(f(Y_k)-r_2)\mu(Y_k, Y_{k+1}) \Rightarrow \sigma_2 N(0,1),$$

(d) the sequences

$$\left\{n^{-1}\sum_{k=0}^{n-1} f(Y_k)\alpha_k : n \geq 1\right\},$$

$$\{T_n/n : n \geq 1\}$$

and

$$\left\{n^{-1}\left[\int_0^{T_n}(f(X(s))-r)\,ds\right]^2 : n \geq 1\right\}$$

are uniformly integrable.

Section 3.3 contains examples.

Not surprisingly, we have:

Proposition. $r_1 = r_2$.

Proof. From (a) and (b), we get $(1/n)\int_0^{T_n} f(X(s))\,ds \Rightarrow r_1\beta$. Equivalently, $(1/n)\sum_{k=0}^{n-1} f(Y_k)\alpha_k \Rightarrow r_1\beta$. From (d), we also get weak convergence to $r_1\alpha$ in the function space L_1. From Chung [4, p. 306] we therefore get

$$\frac{1}{n}\sum_{k=0}^{n-1} E(f(Y_k)\alpha_k \mid Y) \Rightarrow r_1\beta;$$

i.e.,

$$\frac{1}{n}\sum_{k=0}^{n-1} f(Y_k)\mu(Y_k, Y_{k+1}) \Rightarrow r_1\beta.$$

The left side is the numerator of $\tilde{R}_n$. Likewise, for the denominator of $\tilde{R}_n$ we get

$$\frac{1}{n}\sum_{k=0}^{n-1}\mu(Y_k, Y_{k+1}) \Rightarrow \beta.$$

From (c), $\tilde{R}_n \Rightarrow r_2$. Comparing these results proves that $r_1 = r_2$. □

We are now ready to compare $\hat{R}_n$ and $\tilde{R}_n$. From (a) and (b),

$$\sqrt{n}(\hat{R}_n - r_1) \Rightarrow (\sigma_1/\beta)N(0,1); \tag{22}$$

from (a) and (c),

$$\sqrt{n}(\tilde{R}_n - r_2) \Rightarrow (\sigma_2/\beta)N(0,1). \tag{23}$$

That confidence intervals wit fixed coverage based on $\tilde{R}_n$ are asymptotically shorter than those based on $\hat{R}_n$ now follows from our main theorem:

Theorem. $\sigma_2^2 \leq \sigma_1^2$.

Proof. By Jensen's inequality for conditional expectations (see Chung [4, p. 302]),

$$n^{-1}E\left(\int_0^{T_n}(f(X(s))-r)\,ds\right)^2$$

$$\geq n^{-1}EE\left[\left(\int_0^{T_n}(f(X(s))-r)\,ds \mid Y\right)^2\right]$$

$$= n^{-1}E\left[\sum_{j=0}^{n-1}(f(Y_j)-r)\mu(Y_j, Y_{j+1})\right]^2.$$

By (b) and (d) the extreme left side converges to σ_1^2. We show that the right side converges to σ_2^2, which requires a uniform integrability argument.

Let

$$A_{nk} = \left\{n^{-1}\left[\sum_{j=0}^{n-1}(f(Y_j)-r)\mu(Y_j, Y_{j+1})\right]^2 \geq k\right\}.$$

Since A_{nk} is Y-measurable, apply Jensen's inequality again to get

$$E\left\{I_{A_{nk}}\left(\frac{1}{n}\sum_{j=0}^{n-1}(f(Y_j)-r)\mu(Y_j, Y_{j+1})\right)^2\right\}$$

$$\leq E\left[I_{A_{nk}}\left(\frac{1}{n}\int_0^{T_n}(f(X(s))-r)\,ds\right)^2\right].$$

Use (c) to see that $\lim_{k\to\infty}\sup_n P(A_{nk}) \to 0$ and then Chung [4, Theorem 4.5.3] (implications of uniform integrability) to see that this implies that the term in brackets goes to 0 uniformly in n as $k \to \infty$. This proves uniform integrability of the term in braces. □

3.3. Examples

(1) Let X be a semi-Markov process for which Y is regenerative: there exists s such that $P[Y_n = s \text{ infinitely often}] = 1$. Set $T = \inf\{n \geq 1: Y_n = s\}$ and assume that $Y_0 = s$ and

$$E\left(\sum_{k=0}^{T-1}(|f(Y_k)|+1)\alpha_k\right)^4 < \infty. \tag{24}$$

Then (a)–(d) hold; for uniform integrability, see Chung [3].

(2) Let X be a semi-Markov process for which Y is a stationary ψ-mixing process. If Y is ψ-mixing, then $\{f(Y_n)\alpha_n\}$ is ψ-mixing with mixing coef-

ficients doublet since

$$
\begin{aligned}
&|P\{f(Y_0)\alpha_0 \in A,\ f(Y_n)\alpha_n \in B\} \\
&\quad - P\{f(Y_0)\alpha_0 \in A\} \cdot P\{f(Y_n)\alpha_n \in B\}| \\
&= |\mathrm{E}\{P\{f(Y_0)\alpha_0 \in A \mid Y\} \\
&\qquad \cdot P\{f(Y_n)\alpha_n \in B \mid Y\}\} \\
&\quad - P\{f(Y_0)\alpha_0 \in A\} \cdot P\{f(Y_n)\alpha_n \in B\}| \\
&= |\mathrm{E}g_1(Y_0, Y_1) \cdot g_2(Y_n, Y_{n+1}) \\
&\quad - \mathrm{E}g_1(Y_0, Y_1) \cdot \mathrm{E}g_2(Y_n, Y_{n+1})| \\
&\quad (\text{where } g_1(x, y) = P\{f(Y_0)\alpha_0 \in A \mid \\
&\qquad\qquad Y_0 = x,\ Y_1 = y\} \\
&\qquad g_2(x, y) = P\{f(Y_n)\alpha_n \in B \mid \\
&\qquad\qquad Y_n = x,\ Y_{n+1} = y\}) \\
&\leqslant 2\psi_n \mathrm{E}g_1(Y_0, Y_1) = 2\psi_n \cdot P\{f(Y_0)\alpha_0 \in A\},
\end{aligned}
$$

the inequality following from Billingsley [1, (20.28)]. Suppose that

$$\sum_{n=0}^{\infty} \psi_n^{\frac{1}{2}} < \infty,$$

and

$$\mathrm{E}[(f(Y_0)+1)\alpha_0]^2 < \infty.$$

Then, (a)–(d) hold; see Billingsley [1, Theorem 20.1], which is a more general functional central limit theorem.

4. Uniformization loses

For the special case of continuous-time Markov chains with generator $Q = (q_{ij})$ with $\sup_i \sum_{j \neq i} q_{ij} = \lambda < \infty$, we can uniformize. This corresponds to the equal holding-time method of [7]. The method of Section 3.1, without uniformization, corresponds to the constant holding-time method of [7]. With uniformization, the imbedded discrete-time Markov chain W has *null*-jumps from a state to itself. Thus, to generate any fixed number of regeneration cycles with the equal holding-time method requires more jumps, hence more work. It also gives more variance, as shown in [7].

We now give a simpler proof. Delete all null jumps from W to get Y, giving $\sigma(Y) \subseteq \sigma(W)$ where $\sigma(H)$ is the σ-field generated by H. By [5, Proposition G.1], for example, we get higher variance by conditioning on the latter. Thus,

$$\mathrm{Var}(\mathrm{E}[D_k \mid Y]) \leqslant \mathrm{Var}(\mathrm{E}[D_k \mid W]). \tag{25}$$

Nevertheless, what is the best way to uniformize? To get a legitimate representation, we must choose the (Poisson) clock intensity $\theta \geqslant \lambda$ in the uniformized process. Choosing $\theta = \lambda$ stochastically minimizes the number of jumps in W to simulate a fixed number of regeneration cycles, hence stochastically minimizes work. Hordijk, Iglehart and Schassberger [7] show by explicit calculation that choosing $\theta = \lambda$ also minimizes σ^2 in (18). An outline of an easier proof follows. Subscript W and X to indicate clock intensity. Put X_λ and X_θ on the same probability space using three synchronized streams of common random numbers. Stream 1 generates the common non-null jump subsequences of W_λ and W_θ. Stream 2 generates the null jumps so that each null-jump sequence in W_λ is at most as long as the corresponding null-jump sequence in W_θ. Stream 3 generates clock chimes so that chime j of the θ-clock sounds before or with chime j of the λ-clock. [5, App. B] in a more general setting, spells out the details. Thus,

$$\sigma(W_\lambda) \subseteq \sigma(W_\theta), \tag{26}$$

for all $\theta \geqslant \lambda$. Again using the fact that conditioning on less reduces variance,

$$\mathrm{Var}(\mathrm{E}[D_k \mid \sigma(W_\lambda)]) \leqslant \mathrm{Var}(\mathrm{E}[D_k \mid \sigma(W_\theta)]), \tag{27}$$

illustrating the power of the conditional Monte Carlo approach.

In the setting of Section 3.2 specialized to continuous-time Markov chains, similar arguments give counterparts to (25) and (27). Since uniformization increases both variance and work, don't use it there either.

Acknowledgment

Peter Lewis informed us that he too was aware that using the general principle that computing conditional expectations reduces variance would streamline the proofs in [7].

References

[1] P. Billingsley, *Convergence of Probability Measures*, Wiley, New York, 1968.
[2] P. Bratley, B.L. Fox and L.E. Schrage, *A Guide to Simulation*, Springer, New York, 1983.
[3] K.L. Chung, *Markov Chains with Stationary Transition Probabilities*, Springer, New York, 1967.
[4] K.L. Chung, *A Course in Probability Theory*, Academic Press, New York, 1974.
[5] B.L. Fox and P.W. Glynn, "Discrete-time conversion for finite-horizon Markov processes". Technical Report, 1985.
[6] D.P. Heyman and M.J. Sobel, *Stochastic Models in Operations Research*, Vol. 1, McGraw-Hill, New York, 1982.
[7] A. Hordijk, D.L. Iglehart nd R. Schassberger. "Discrete-time methods of simulating continuous-time Markov chains", *Adv. Appl. Prob.* **8**, 772–788 (1976)
[8] M.F. Neuts, *Matrix-Geometric Solutions in Stochastic Models*, Johns Hopkins University Press, Baltimore, MD, 1981.
[9] L. Schruben, "Confidence interval estimation using standardized time series", *Operations Research* **31**, 1090–1108 (1983).