

INCREASING THE FREQUENCY OF REGENERATION FOR MARKOV PROCESSES

Sigrún Andradóttir

James M. Calvin

Department of Industrial Engineering
University of Wisconsin – Madison
1513 University Avenue
Madison, WI 53706, U.S.A.

Department of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332, U.S.A.

Peter W. Glynn

Department of Operations Research
Stanford University
Stanford, CA 94305, U.S.A.

ABSTRACT

We present and compare two ways of applying the regenerative method of simulation output analysis to simulate the steady-state behavior of irreducible discrete time Markov chains having a finite state space. The “standard” approach involves defining the regenerative cycles based on the times a particular “return” state is visited. The alternative approach involves using “splitting” to define the regenerative cycles. We present a way of selecting the splitting distribution that guarantees that the splitting cycles are shorter than the standard cycles, and that the variance estimates produced by the splitting approach are less variable than the variance estimates produced by the standard approach.

1 INTRODUCTION

The regenerative method of simulation output analysis is designed to obtain strongly consistent point and variance estimates from one long sample path in steady-state simulation of regenerative stochastic processes. In order to apply the regenerative method, it is necessary to identify times when the stochastic process of interest regenerates, and to continue generating a sample path of this stochastic process until it has regenerated a number of times. This can limit the applicability of the regenerative method as regeneration times are sometimes difficult to identify, and as it may take a prohibitive amount of computer time to continue the simulation until several regenerations have been observed.

In this paper, we study the application of the regenerative method to steady-state simulation of irreducible discrete time Markov chains on a finite state space. In this case, it is easy to identify regeneration times as the Markov chain regenerates every time it returns to a given state. Thus, the “standard” way of applying the regenerative method to simulate Markov chains involves selecting any one “return” state in the state space, simulating the chain until this state has been visited several times, and using the times when this state is visited as the regeneration times. Unfortunately, this approach does not always work well as it may be impossible to select a return state that is visited often enough for this approach to be practical. We discuss another way of defining the regeneration times based on using splitting. This approach involves selecting a particular “splitting distribution,” and using the times when the Markov chain has this distribution as regeneration times. It will be shown that when the standard method is used with a return state s , and the splitting method is used with the splitting distribution consisting of the s th row of the transition probability matrix of the Markov chain being simulated, the splitting regenerative method will produce both shorter regenerative cycles and more precise variance estimates than the standard regenerative method.

This paper is organized as follows: in Section 2, we present the standard and splitting approaches to regenerative simulation of discrete time Markov chains. Then in Section 3, we compare the estimators produced by these two methods. Finally, Section 4 contains some concluding remarks.

2 REGENERATIVE SIMULATION OF MARKOV CHAINS

Let $\{X_n\}$ be a discrete time Markov chain with finite state space S and transition probability matrix P . Suppose further that $\{X_n\}$ is irreducible. Then $\{X_n\}$ has a unique stationary distribution π . We want to estimate the steady-state mean $r = E_\pi\{f(X_0)\}$, where $f : S \rightarrow \mathcal{R}$ is a given function (if μ is a probability distribution on S , then E_μ is the expected value operator given that X_0 has distribution μ and P_μ is the corresponding probability measure). We will present and compare two ways of applying regenerative simulation to estimate r . For more details on the regenerative method of steady-state simulation output analysis, the reader is referred to Crane and Lemoine (1977) and to Iglehart and Shedler (1980).

The standard way of applying regenerative simulation to estimate r is to select a fixed return state $s \in S$ and to end a regenerative cycle whenever the state s is hit. Equivalently, one can use the regeneration times $T_0^{(1)}, T_1^{(1)}, T_2^{(1)}, \dots$, where $X_0 = s$, $T_0^{(1)} = 1$, and for $m \geq 1$,

$$T_m^{(1)} = \inf\{n \geq T_{m-1}^{(1)} : X_n = s\} + 1.$$

Another way of applying regenerative simulation to estimate r is to select a fixed splitting distribution φ and to end a regenerative cycle whenever the Markov chain $\{X_n\}$ has the distribution φ . For all $i \in S$, let $\lambda_i \in [0, 1]$ be such that

$$P_{ij} \geq \lambda_i \varphi_j, \quad (1)$$

for all $i, j \in S$. Suppose φ is such that at least one of the λ_i , $i \in S$, is strictly positive. Then it is well known that the Markov chain $\{X_n\}$ is regenerative (see Athreya and Ney (1978) and Nummelin (1978)). By equation (1), whenever the Markov chain is in state i , with probability λ_i the next state will have the distribution φ , independently of the past. This means that regeneration has taken place. We will restrict our attention to the case when $\varphi_i = P_{si}$, for all $i \in S$. Then we can take $\lambda_s = 1$, so the Markov chain $\{X_n\}$ is regenerative. For all $n = 0, 1, 2, \dots$, let δ_n be a Bernoulli random variable with parameter λ_{X_n} (so $\delta_n = 1$ with probability λ_{X_n} and $\delta_n = 0$ with probability $1 - \lambda_{X_n}$), and suppose δ_n is independent of both X_0, \dots, X_n and $\delta_0, \dots, \delta_{n-1}$. Let $X_0 = s$, $T_0^{(2)} = 1$, and for $m \geq 1$,

$$T_m^{(2)} = \inf\{n \geq T_{m-1}^{(2)} : \delta_n = 1\} + 1.$$

Then $T_0^{(2)}, T_1^{(2)}, T_2^{(2)}, \dots$ are regeneration times for the Markov chain $\{X_n\}$.

For $k = 1, 2$, and for all $m \geq 1$, let

$$\begin{aligned} \tau_m^{(k)} &= T_m^{(k)} - T_{m-1}^{(k)}, \\ Y_m^{(k)} &= \sum_{n=T_{m-1}^{(k)}}^{T_m^{(k)}-1} f(X_n), \end{aligned}$$

and

$$Z_m^{(k)} = \sum_{n=T_{m-1}^{(k)}}^{T_m^{(k)}-1} [f(X_n) - r].$$

Let $\tau^{(k)} = \tau_1^{(k)}$, $Y^{(k)} = Y_1^{(k)}$, and $Z^{(k)} = Z_1^{(k)}$, for $k = 1, 2$. By the ratio formula for regenerative processes, we have

$$r = \frac{E_\varphi\{Y^{(k)}\}}{E_\varphi\{\tau^{(k)}\}},$$

for $k = 1, 2$. Moreover, we have the following expression for the variance σ^2 :

$$\sigma^2 = \frac{E_\varphi\{[Z^{(k)}]^2\}}{E_\varphi\{\tau^{(k)}\}}, \quad (2)$$

$k = 1, 2$. (It is easy to show that the right-hand side of equation (2) does not depend on k .) Let

$$K_n^{(k)} = \max\{m \geq 0 : T_m^{(k)} \leq n\},$$

for $k = 1, 2$. Then if we simulate the Markov chain until n state transitions have taken place, we get the following strongly consistent estimates for r and σ^2 :

$$r^{(k)}(n) = \frac{\sum_{j=1}^{K_n^{(k)}} Y_j^{(k)}}{\sum_{j=1}^{K_n^{(k)}} \tau_j^{(k)}} = \frac{\sum_{j=1}^{K_n^{(k)}-1} f(X_j)}{T_{K_n^{(k)}}^{(k)} - 1}, \quad (3)$$

$$[s^{(k)}(n)]^2 = \frac{1}{T_{K_n^{(k)}}^{(k)} - 1} \sum_{j=1}^{K_n^{(k)}} \left(Y_j^{(k)} - r^{(k)}(n) \tau_j^{(k)} \right)^2, \quad (4)$$

$k = 1, 2$.

We have defined two sequences of regeneration times $\{T_m^{(k)}\}$, $k = 1, 2$, for the Markov chain $\{X_n\}$. Since $\lambda_s = 1$, it is clear that a splitting regeneration occurs each time the state s is hit. This shows that $\{T_m^{(1)}\} \subset \{T_m^{(2)}\}$, so we have shown that splitting regenerations take place more frequently than standard regenerations (and thereby that the splitting regenerative cycles are shorter than the standard regenerative cycles). In Section 3, we show that the splitting regenerative method also produces more precise variance estimates. Note that shorter cycles do not always result in more precise variance estimates, as shown by Calvin (1994).

3 A COMPARISON BETWEEN STANDARD AND SPLITTING REGENERATIVE ESTIMATORS

By applying a theorem in Glynn and Iglehart (1987), one can show that the point and variance estimators given in equations (3) and (4) satisfy the following joint central limit theorem:

Theorem 1 *Under the assumptions made on the Markov chain $\{X_n\}$,*

$$n^{1/2} \left(r^{(k)}(n) - r, s^{(k)}(n) - \sigma \right) \Rightarrow \mathcal{N}(0, D^{(k)}),$$

for $k = 1, 2$, where

$$\begin{aligned} D_{11}^{(k)} &= \sigma^2, \\ D_{12}^{(k)} &= \frac{E_\varphi\{[Z^{(k)}]^3\} - 3\sigma^2 E_\varphi\{\tau^{(k)} Z^{(k)}\}}{2\sigma E_\varphi\{\tau^{(k)}\}}, \end{aligned}$$

and

$$\begin{aligned} D_{22}^{(k)} &= \frac{E_\varphi\{[Z^{(k)}]^4\}}{4\sigma^2 E_\varphi\{\tau^{(k)}\}} - \frac{E_\varphi\{\tau^{(k)}[Z^{(k)}]^2\}}{2E_\varphi\{\tau^{(k)}\}} + \\ &\quad \frac{\sigma^2 E_\varphi\{[\tau^{(k)}]^2\}}{4E_\varphi\{\tau^{(k)}\}} + \frac{2[E_\varphi\{\tau^{(k)} Z^{(k)}\}]^2}{[E_\varphi\{\tau^{(k)}\}]^2} - \\ &\quad \frac{E_\varphi\{\tau^{(k)} Z^{(k)}\} E_\varphi\{[Z^{(k)}]^3\}}{\sigma^2 [E_\varphi\{\tau^{(k)}\}]^2}, \end{aligned}$$

$k = 1, 2$.

We now derive an expression for the amount of variance reduction that is obtained by using splitting regenerative cycles rather than standard regenerative cycles. For this purpose, define

$$\eta = \min\{m \geq 1 : \text{state } s \text{ is hit in the splitting regenerative cycle } [T_m^{(2)}, T_{m+1}^{(2)}]\}.$$

Then

$$\tau^{(1)} = \sum_{j=1}^{\eta+1} \tau_j^{(2)}$$

and

$$Z^{(1)} = \sum_{j=1}^{\eta+1} Z_j^{(2)}.$$

Note that η is a geometric random variable with parameter $p = P_\varphi\{\tau^{(1)} \leq \tau^{(2)}\} = P_\varphi\{\tau^{(1)} = \tau^{(2)}\}$ since $\tau^{(1)} \geq \tau^{(2)}$. Therefore $E_\varphi\{\eta\} = (1-p)/p$. By Theorem 1, we have $D_{11}^{(1)} = D_{11}^{(2)} = \sigma^2$. The following result is proved in Andradóttir, Calvin, and Glynn (1994):

Theorem 2 *Under the assumptions made on the Markov chain $\{X_n\}$, $D_{12}^{(1)} = D_{12}^{(2)}$ and*

$$\begin{aligned} D_{22}^{(1)} &= D_{22}^{(2)} + E_\varphi\{\eta\} E_\varphi\{[Z^{(2)}]^2 | \eta \geq 1\} + \\ &\quad (E_\varphi\{\eta\})^2 (E_\varphi\{Z^{(2)} | \eta \geq 1\})^2 \\ &\geq D_{22}^{(2)}. \end{aligned}$$

Theorem 2 quantifies the variance reduction obtained by using splitting rather than standard regenerative cycles. Note that the difference between the quality of splitting versus standard regenerative estimators increases as $p = P_\varphi\{\tau^{(1)} \leq \tau^{(2)}\}$ decreases, and as $E_\varphi\{[Z^{(2)}]^2 | \eta \geq 1\}$ and $E_\varphi\{Z^{(2)} | \eta \geq 1\}$ increase.

4 CONCLUSION

We have discussed two ways of applying the regenerative method to steady-state simulation of irreducible discrete time Markov chains on a finite state space. The standard regenerative method uses visits to a fixed return state as regeneration times, whereas the splitting regenerative method uses the times when the Markov chain has a given splitting distribution as regeneration times. We have shown that when the splitting distribution consists of the row of the transition probability matrix of the Markov chain corresponding to the return state of the standard regenerative method, the splitting regenerative method will produce shorter cycles and more precise variance estimates than the standard regenerative method. This improvement obtained through the use of the splitting regenerative method is likely to be most beneficial when the state space of the Markov chain is large, no single state is visited very often, and the transition probability matrix of the Markov chain exhibits a significant amount of row similarity.

ACKNOWLEDGEMENT

The work of the first author was supported by the National Science Foundation under Grant No. DDM-9210679. The work of the third author was supported by the National Science Foundation under Grant No. DDM-9101580 and by the Army Research Office under Contract No. DAAL03-91-G-0319.

REFERENCES

- Andradóttir, S., J. M. Calvin, and P. W. Glynn. 1994. On the choice of regeneration times in Markov chain simulations. Working paper.

- Athreya, K. B., and P. Ney. 1978. A new approach to the limit theory of recurrent Markov chains. *Transactions of the American Mathematical Society* 245:493–501.
- Calvin, J. M. 1994. Return-state independent quantities in regenerative simulation. *Operations Research* 42:531–542.
- Crane, M. A., and A. J. Lemoine. 1977. *An Introduction to the Regenerative Method for Simulation Analysis*. Berlin: Springer-Verlag.
- Glynn, P. W., and D. L. Iglehart. 1987. A joint central limit theorem for the sample mean and regenerative variance estimator. *Annals of Operations Research* 8:41–55.
- Iglehart, D. L., and G. S. Shedler. 1980. *Regenerative Simulation of Response Times in Networks of Queues*. Berlin: Springer-Verlag.
- Nummelin, E. 1978. A splitting technique for Harris recurrent Markov chains. *Zeitschrift für Wahrscheinlichkeitstheorie und verwandte Gebiete* 43:309–318.
- ing, and general theory for stochastic systems.

AUTHOR BIOGRAPHIES

SIGRÚN ANDRADÓTTIR is an Assistant Professor of Industrial Engineering, Mathematics, and Computer Sciences at the University of Wisconsin – Madison. She received a B.S. in Mathematics from the University of Iceland in 1986, an M.S. in Statistics from Stanford University in 1989, and a Ph.D. in Operations Research from Stanford University in 1990. Her research interests include stochastic optimization, simulation and stochastic processes. She is presently a member of ORSA, TIMS, and IIE.

JAMES M. CALVIN is an Assistant Professor of Industrial and Systems Engineering at Georgia Institute of Technology. He received an A.B. in Mathematics from the University of California at Berkeley and a Ph.D. in Operations Research from Stanford University. His interests include average complexity of global optimization and regenerative methods for simulation output analysis.

After receiving his Ph.D. in Operations Research from Stanford University, **PETER W. GLYNN** joined the faculty at the University of Wisconsin – Madison, where he held joint appointments in the Industrial Engineering, Computer Sciences, and Mathematics departments. In 1987, he returned to Stanford University, where he currently holds the Thomas Ford Faculty Scholar Chair in the Department of Operations Research. His research interests include discrete event simulation, computational probability, queue-