Proceedings of the 1987 Winter Simulation Conference
A. Thesen, H. Grant, W. David Kelton (eds.)

MEASURE SPECIFIC DYNAMIC IMPORTANCE SAMPLING FOR AVAILABILITY SIMULATIONS

Ambuj Goyal *
Philip Heidelberger *
Perwez Shahabuddin **[1]

* IBM Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598

** Department of Operations Research
Stanford University
Stanford, California 94305

ABSTRACT

This paper considers the application of importance sampling to simulations of highly available systems. By regenerative process theory, steady state performance measures of a Markov chain take the form of a ratio. Analysis of a simple three state Birth and Death process shows that the optimal (zero variance) importance sampling distributions for the numerator and denominator of this ratio are quite different and are both dynamic in that they do not correspond directly to time homogeneous Markov chains. Analysis of this three state example suggests heuristics for choosing effective importance sampling distributions for more complex models of highly available systems. These heuristics are applied to a large model of computer system availability. The example shows that additional variance reduction over that previously reported can be obtained by simulating the numerator and denominator independently with different dynamic importance sampling distributions.

1. INTRODUCTION

The requirement for highly available systems, such as fault tolerant computing systems, is increasing the importance of reliability and availability prediction during the design phase of these systems. While such systems can typically be modeled as Markov chains (see, e.g., [3]), the size of the corresponding Markov model increases rapidly with complexity of the system. Thus numerical solution techniques are only feasible for relatively small models, i.e., simple systems. Simulation analysis is an alternative approach, however, because system failures are rare, extremely long simulations may be required in order to obtain accurate estimates of availability.

This paper discusses the use of importance sampling (see, e.g., [4]) as a variance reduction technique for simulating highly available systems. Importance sampling for rare event simulation has been successfully used in [1], [7], [9] and [10]. Proper selection of the importance sampling distribution makes the rare events more likely to occur and Walrand [10] gives an additional intuitive explanation as to why importance sampling can be an effective variance reduction technique for rare event simulations. The key, of course, is to choose a good importance sampling distribution. The theory of large deviations was used in [9] and [10] to select an effective distribution for problems of essentially estimating buffer overflow probabilities in queueing networks. Effective heuristics were used in [1] and [7] to select importance sampling distributions for availability estimation in large machine repairmen-like models. For example, four orders of magnitude reduction in

[1] This work was performed while this author was visiting IBM Research.

variance was reported in [1] for a system with 70 components (the Markov chain for this model has 2^{70} states).

In this paper we extend the method of [1] to obtain additional variance reduction. Since the underlying Markov chain models are regenerative, the regenerative method (see, e.g., [2]) can be used to estimate steady state performance measures. Furthermore, highly available systems act essentially like machine repairmen models with very low repairmen utilization. Since the system is typically operational, an operational state is an appropriate choice of regeneration state and the regenerative method is often well suited for this type of simulation (particularly if all components have exponential failure time distributions). For regenerative systems, steady state performance measures can be expressed as a ratio. In [1], a single importance sampling distribution was used to estimate both the numerator and denominator of this ratio. The distribution used in [1] is dynamic in the sense that it does not correspond directly to a time homogeneous Markov chain. In this paper, we use different dynamic importance sampling distributions to estimate the numerator and denominator independently. For a simple three state birth and death process, optimal dynamic importance sampling distributions are derived that lead to zero variance estimates. For estimating system unavailability, the optimal distributions for the numerator and denominator are quite different. Furthermore, the analysis of this three state model leads to a natural and simple heuristic for selecting dynamic importance sampling distributions for more complex models of highly available systems.

In Section 2, we review the regenerative method of simulation for Markov chains and describe the proposed variance reduction technique in general terms. In Section 3, the three state birth and death process is considered and its optimal dynamic importance sampling probability distributions are derived. The heuristic for choosing dynamic importance sampling probability distributions for larger, more complicated models of highly available systems is described in Section 4 along with experiments reporting on variance reductions obtained when this heuristic is applied to a larger model of computer system availability. In some cases, an additional 3 orders of magnitude reduction in variance are obtained over the method of [1]. Section 5 contains concluding remarks.

2. DYNAMIC IMPORTANCE SAMPLING

In this section, we review the regenerative method of simulation for Markov chains and describe the proposed variance reduction technique in general terms. We assume that the process of interest $\{Y_s, s \geq 0\}$ is a continuous time Markov chain with finite state space $E = \{0, 1, \dots, N\}$ and transition rate matrix $\mathbf{Q} = (q(i,j))$ (see, e.g., [6]). Let $q(i) = -q(i,i) = -\sum_{j \neq i} q(i,j)$ denote the rate out of state i. We assume that the process converges to a steady state distribution, i.e., $Y_s \Rightarrow Y$ where $\Rightarrow$ denotes convergence in distribution. The goal of the simulation is to estimate $r = E[f(Y)]$ for some function f. Let $\{X_n, n \geq 0\}$ denote the embedded, discrete time Markov chain: $\{X_n, n \geq 0\}$ has transition matrix $\mathbf{P}$ where $p(i,j) = q(i,j)/q(i)$ for $i \neq j$ and $p(i,i) = 0$. Pick a regeneration state, say 0, let $X_0 = 0$ and let τ_1 be the first $n > 0$ such that $X_n = 0$. Then the steady state performance measure r of the continuous time Markov chain can be expressed as

$$r = \frac{E[\sum_{n=0}^{\tau_1 - 1} g(X_n)]}{E[\sum_{n=0}^{\tau_1 - 1} h(X_n)]} \qquad (2.1)$$

where $g(i) = f(i)/q(i)$ and $h(i) = 1/q(i)$. Simulation of the discrete time Markov chain is guaranteed to yield a smaller variance than simulation of the continuous time Markov chain (see [5]). Letting s denote a sample path, then $r = E_P[G(s)]/E_P[H(s)]$ where $E_P[G(s)]$ and $E_P[H(s)]$ denote the numerator and denominator, respectively, of Equation 2.1 and E_P denotes expectation using the transition matrix **P**.

If m iid regenerative cycles of the discrete time chain are simulated and if $G_k(s)$ and $H_k(s)$ are the realizations of $G(s)$ and $H(s)$ on cycle k, then r can be estimated by $\hat{r}_m(P) = \sum_{k=1}^{m} G_k(s) / \sum_{k=1}^{m} H_k(s)$. In particular, $\lim_{m\to\infty} \hat{r}_m(P) = r$ with probability one and $\sqrt{m}\,(\hat{r}_m(P) - r) \Rightarrow N(0, \sigma^2(P)/E_P[H(s)]^2)$ where $N(0, \sigma^2)$ denotes a normally distributed random variable with mean zero and variance σ^2 and $\sigma^2(P) = \mathrm{Var}_P[G_k(s) - rH_k(s)]$. Thus the asymptotic variance of $\hat{r}_m(P)$ is $\sigma^2(P)/(E_P[H(s)]^2 m)$.

Our goal is to use importance sampling to derive an estimator with a smaller asymptotic variance than $\hat{r}_m(P)$. First, let P(s) denote the probability of a sample path $s = (X_0, X_1, \ldots, X_{\tau_1})$ using the transition matrix **P**: $P(s) = p(X_0, X_1)p(X_1, X_2) \ldots p(X_{\tau_1 - 1}, X_{\tau_1})$. Now let $P'(s)$ denote the probability of s under another probability distribution. Assuming $P'(s) \neq 0$ whenever $G(s)P(s) \neq 0$, then (under some additional technical conditions)

$$E_P[G(s)] = \sum_s G(s)P(s) = \sum_s G(s)\frac{P(s)}{P'(s)}P'(s) = E_{P'}[L'(s)G(s)] \tag{2.2}$$

where $E_{P'}$ denotes the expectation using the probability distribution P' and $L'(s) = P(s)/P'(s)$ is the likelihood ratio of the sample path s.

Letting $L'_k(s)$ denote the likelihood ratio during cycle k, then $\hat{r}_m(P') = \sum_{k=1}^{m} L'_k(s)G_k(s) / \sum_{k=1}^{m} L'_k(s)H_k(s)$ converges to r with probability one and is asymptotically normally distributed with mean r and variance $\sigma^2(P')/(E_P[H(s)]^2 m)$ where

$$\begin{aligned}\sigma^2(P') = {} & \mathrm{Var}_{P'}[L'_k(s)G_k(s)] \\ & - 2r\mathrm{Cov}_{P'}[L'_k(s)G_k(s), L'_k(s)H_k(s)] \\ & + r^2\mathrm{Var}_{P'}[L'_k(s)H_k(s)].\end{aligned} \tag{2.3}$$

We are free to chose P' in an essentially arbitrary manner; in particular P' need not correspond to a time homogeneous Markov chain. A general form for P' is

$$P'(i_0, i_1, \ldots, i_n) = P'(i_1 | i_0)P'(i_2 | i_0, i_1) \ldots P'(i_n | i_0, \ldots, i_{n-1}) \tag{2.4}$$

where

$P'(i_n | i_0, \ldots, i_{n-1}) = P'(X_n = i_n | X_0 = i_0, \ldots, X_{n-1} = i_{n-1})$.

Thus the importance sampling one-step transition distributions may be dynamic in the sense that they can depend on the past history of states and not just the current state as would be the case in a time homogeneous Markov chain. We call this dynamic importance sampling; this was used in [1] to achieve dramatic variance reductions in simulations of highly available systems. In Section 3 it will be shown that a zero variance estimator for $E_P[G(s)]$ of the form $L'(s)G(s)$ cannot be obtained using a static (time homogeneous) importance sampling distribution for a three state birth and death process but that dynamic importance sampling can be used to obtain a zero variance estimator of the numerator $E_P[G(s)]$.

Similarly, a zero variance estimator of the denominator $E_P[H(s)]$ may be constructed using a different dynamic importance sampling distribution. Thus for this simple example, a zero variance estimate of the ratio may be obtained by simulating the numerator and denominator independently using different dynamic importance sampling distributions. This suggests that different dynamic importance sampling distributions should also be used for the numerator and denominator in more complex models. This allows customization of the dynamic importance sampling distributions to account for the particular forms of the two different reward functions g and h. Furthermore, it is clear from Equation 2.3, that choosing P′ to, say, minimize the variance of the numerator may adversely affect the other two terms in the asymptotic variance. If, for example, cycles $1, \dots, m/2$ are simulated using P′ and cycles $m/2+1, \dots, m$ are independently simulated using P″, then $\hat{r}_m(P', P'') = \sum_{k=1}^{m/2} L'_k(s)G_k(s) / \sum_{k=m/2+1}^{m} L''_k(s)H_k(s)$ converges to r with probability one and is asymptotically normally distributed with mean r and variance $\sigma^2(P', P'')/(mE_P[H(s)]^2/2)$ where

$$\sigma^2(P', P'') = \mathrm{Var}_{P'}[L'_k(s)G_k(s)] + r^2\mathrm{Var}_{P''}[L''_k(s)H_k(s)]. \tag{2.5}$$

Because there is no covariance term in Equation 2.5, the use of different distributions allows one to reduce the variance of the numerator without adversely affecting the variance of the denominator and vice versa.

3. A THREE STATE EXAMPLE

In this section, the optimal dynamic importance sampling probability distributions for evaluating $E_P[G(s)]$ and $E_P[H(s)]$ are computed for a three state birth death process. States 0, 1 and 2 have mean holding times T_0, T_1 and T_2, respectively, $p(1,0) = p$ and $p(1,2) = (1-p)$. The steady state measure r of interest is the stationary probability of being in state 2. For availability modeling, we interpret states 0 and 1 to be operational states and state 2 to be the failed state. Thus r is the steady state unavailability. This could be obtained by a regenerative simulation with functions $g(0)$, $g(1)$ and $g(2)$ equal to 0, 0 and T_2, respectively, and functions $h(0)$, $h(1)$ and $h(2)$ equal to T_0, T_1 and T_2, respectively. Assume state 0 is the regenerative state.

The optimal P′(s) and P″(s) are computed from explicit enumeration of all sample paths. Let s_i denote the sample path containing exactly $i+1$ visits to state 1 between visits to state 0. Now, $P(s_i) = p(1-p)^i$, $G(s_i) = iT_2$ and $H(s_i) = T_0 + (i+1)T_1 + iT_2$ for $i \geq 0$. The optimal (i.e., zero variance) P′(s) for estimating $E_P[G(s)]$ can be computed as follows [4],

$$P'^*(s_i) = \frac{P(s_i)G(s_i)}{\sum_s P(s)G(s)} = p^2(1-p)^{i-1}i \tag{3.1}$$

for $i \geq 0$. Similarly, the optimal P″(s) for estimating $E_P[H(s)]$ is given by

$$P''^*(s_i) = \frac{(T_0 + T_1(i+1) + T_2 i)p^2(1-p)^i}{pT_0 + T_1 + (1-p)T_2} \tag{3.2}$$

From Equation 3.1, $P'^*(s_0) = 0$, $P'^*(s_1) = p^2$, $P'^*(s_2) = 2p^2(1-p)$ and so on. That is, in the notation of Equation 2.4 $P'^*(0|0,1) = 0$, $P'^*(0|0,1,2,1) = p^2$, $P'^*(0|0,1,2,1,2,1) = 2p^2/(1+p)$ etc. Therefore, each successive time the simulation enters state 1, the probability of returning to state 0 changes (under both $P'^*(s)$ and $P''^*(s)$). Thus the optimal change of measures for both the numerator and the denominator of Equation 2.1 are dynamic.

The following observations may be made from these optimal importance sampling distributions. In highly available systems, we are specifically interested in the case where p is close to 1, i.e., the probability of entering state 2 in a regenerative cycle is very small. While simulating the numerator using $P'^*(s)$, the probability of returning to state 0 from state 1 is zero if state 2 has not been visited, otherwise, for p close to

1, it is order p for small values of i and goes back to zero for large values of i. On the other hand, while simulating the denominator using $P''^*(s)$, the probability of returning to state 0 from state 1 is order p. Thus, the optimal importance sampling distributions are very different for the numerator and the denominator in rare event simulations.

4. HEURISTICS FOR LARGE MODELS

Based on the observations made on the three state example, heuristics may be developed for larger models. One natural way to do rare event simulations would be that while simulating the numerator, we should force the simulation to enter the desired rare state as quickly as possible, and once this state is entered, we should make $P'(s) = P(s)$ for the remainder of the regenerative cycle. Therefore, typically a small number of transitions take place in the remainder of the regenerative cycle. The heuristic assumes this condition would occur and does not change $P'(s)$ any further. (Instead of a single rare state, there could be a set of rare states with non-zero $g(X_n)$ reward functions.) On the other hand, we should not change the sampling distribution to simulate the denominator (i.e., the length of the regenerative cycle), that is, $P''(s) = P(s)$.

The dynamic importance sampling distribution employed in [1] was the same as above heuristic employed to simulate the numerator. However, both the numerator and the denominator were estimated from the same simulation runs. The $P'(s)$ ($= P''(s)$) was selected such that from any state the probability of getting closer to a rare state was p' and getting closer to the regenerative state was $1 - p'$. Once a rare state was visited in a regenerative cycle, $P(s)$ was used for the remaining part of the regenerative cycle. Experiments done in [1] suggested that p' should be selected as 0.5. Typically, the confidence intervals became wider as p' deviated from 0.5. The heuristic employed in [1] will be referred to as DIS (dynamic importance sampling) while heuristic proposed here will be referred to as MSDIS (measure specific dynamic importance sampling).

From our analysis it becomes clear why this phenomenon occurs. As p' increases, we move closer to $P'^*(s)$, but further away from $P''^*(s)$, and when p' decreases, the reverse effect occurs. Therefore, for small p' the variance of the numerator dominates, while for large p' the variance of the denominator dominates. This phenomenon was clearly visible when we duplicated the experiments done in [1].

As a simple experiment, we ran the three state example of [1] which has $T_0 = 1/2\lambda$, $T_1 = 1/(\mu + \lambda)$, $T_2 = 1/\mu$ and $p = \mu/(\mu + \lambda)$. In this example, λ is the component failure rate and μ is the repair rate. We used $p' = 0.999$ for the numerator for a total of about 52,500 events and used no change of measure for the denominator for a total of 52,500 events. We completed the regenerative cycle in progress when the desired number of events was exhausted as suggested in [8]. The results obtained from this experiment are compared in Table 1 to those in [1] where the results were obtained from a single simulation run of 105,000 events with $p' = 0.5$. As λ decreases (p increases), we observe a greater improvement factor in the confidence interval widths (the improvement factor is the ratio of the confidence interval widths). The last row in the table shows two orders of magnitude reduction in the confidence interval widths, which means four orders of magnitude reduction in variance over and above the method in [1], which itself is a few orders of magnitude improvement over the direct simulation (i.e., no importance sampling).

Table 1

Half-Widths of 99% Confidence Intervals for Three State Example With $\mu = 1$.

λ	DIS	MSDIS	Improvement Factor
10^{-1}	0.44×10^{-3}	0.13×10^{-3}	3.5
10^{-2}	0.54×10^{-5}	0.64×10^{-6}	8.4
10^{-3}	0.55×10^{-7}	0.17×10^{-8}	32.4
10^{-4}	0.55×10^{-9}	0.69×10^{-11}	79.7
10^{-5}	0.55×10^{-11}	0.56×10^{-13}	98.2

The same heuristic is employed to simulate a larger fault-tolerant system with ten types of components where each type has two components. This yields a Markov chain with 3^{10} states. The exponential failure and repair rates of each component are assumed to be 0.00001 and 1, respectively. There is a single repairman in the system. The failed components are selected for repair according to a Random Order Service Preemptive Resume discipline. The system is considered unavailable if both components in any given type are failed. The measure of interest is the steady state unavailability of the system. Table 2 shows the results obtained using the various importance sampling techniques (the exact value was obtained from the product form queueing network results). In 200,000 events, direct simulation produces a 99% confidence interval that includes negative values and has a relative width of 368%. DIS gives 182 times reduction in the confidence interval widths, and MSDIS gives another 50 times reduction in the confidence interval widths over DIS (for MSDIS we used approximately 180,000 events for the numerator and 20,000 events for the denominator; the contributions of the numerator and denominator to the asymptotic variance are then approximately equal). Thus almost 8.3×10^7 times reduction in variance over direct simulation is obtained.

Table 2

Unavailability Estimates for the Fault-Tolerant System Example.

	Unavailability	Half-Widths of 99% Confidence Intervals
Exact	0.20005×10^{-8}	
Direct	0.40030×10^{-8}	0.73×10^{-8}
DIS	0.20100×10^{-8}	0.40×10^{-10}
MSDIS	0.19997×10^{-8}	0.82×10^{-12}

5. SUMMARY

This paper considered the application of importance sampling to simulations of highly available systems. Analysis of a simple three state Birth and Death process showed that the optimal (zero variance) importance sampling distributions for the numerator and denominator of the ratio in regenerative simulations are quite different and are both dynamic in that they do not correspond directly to time homogeneous Markov chains. Analysis of this three state example suggested heuristics for choosing effective importance sampling distributions for larger models of highly available systems. These heuristics were applied to a large model of computer system availability. The example showed that additional variance reduction over that previously reported can be obtained by simulating the numerator and denominator independently with different dynamic importance sampling distributions.

We have observed similar improvements for many models of computer system availability. Whenever DIS works well, MSDIS works even better. Furthermore, since for unavailability estimation $r \approx 0$, by Equation 2.5 the variance of the numerator is the dominant term in the asymptotic variance of the ratio. Therefore, our current research is concentrating on selecting good importance sampling distributions for the numerator.

ACKNOWLEDGEMENTS

We would like to thank A. Conway and S.S. Lavenberg for helpful suggestions and discussions relating to this work.

REFERENCES

[1] Conway, A.E. and Goyal, A. (1987). Monte Carlo Simulation of Computer System Availability/Reliability Models. *Proceedings of the Seventeenth Symposium on Fault-Tolerant Computing*. Pittsburgh, Pennsylvania, 230-235.

[2] Crane, M.A. and Iglehart, D.L. (1975). Simulating Stable Stochastic Systems, III: Regenerative Processes and Discrete Event Simulations. *Operations Research* **23**, 33-45.

[3] Geist, R.M. and Trivedi, K.S. (1983). Ultra-High Reliability Prediction for Fault-Tolerant Computer Systems. *IEEE Transactions on Computers* **C-32**, 1118-1127.

[4] Hammersley, J.M. and Handscomb, D.C. (1964). *Monte Carlo Methods*. Methuen, London.

[5] Hordijk, A., Iglehart, D.L. and Schassberger, R. (1976). Discrete Time Methods for Simulating Continuous Time Markov Chains. *Adv. Appl. Prob.* **8**, 772-788.

[6] Karlin, S. and Taylor, H.M. (1975). *A First Course in Stochastic Processes*, Second Edition. Academic Press, New York.

[7] Lewis, E.E. and Bohm, F. (1984). Monte Carlo Simulation of Markov Unreliability Models. *Nuclear Engineering and Design* **77**, 49-62.

[8] Meketon, M.S and Heidelberger, P. (1982). A Renewal Theoretic Approach to Bias Reduction in Regenerative Simulations. *Management Science* **24**, 173-181.

[9] Siegmund, D. (1976). Importance Sampling in the Monte Carlo Study of Sequential Tests. *The Annals of Statistics* **4**, 673-684.

[10] Walrand, J. (1987). Quick Simulation of Rare Events in Queueing Networks. *Proceedings of the Second International Workshop on Applied Mathematics and Performance/Reliability Models of Computer/Communication Systems*. G. Iazeolla, P.J. Courtois and O.J. Boxma (eds). North Holland Publishing Company, Amsterdam, 275-286.

AUTHOR'S BIOGRAPHIES

AMBUJ GOYAL received the B.Tech. degree from the Indian Institute of Technology, Kanpur, and the M.S. and Ph.D. degrees from The University of Texas at Austin. At present, he is a Research Staff Member at the IBM Thomas J. Watson Research Center in Yorktown Heights, New York. His research interest include computer architecture, communication networks, performance evaluation and fault-tolerant computing. Dr. Goyal is a member of the IEEE Computer Society.

Ambuj Goyal
IBM Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598
(914) 789-7504

PHILIP HEIDELBERGER received a B.A. in mathematics from Oberlin College, Oberlin, Ohio, in 1974 and a Ph.D. in Operations Research from Stanford University, Stanford, California, in 1978. He has been a Research Staff Member at the IBM Thomas J. Watson Research Center in Yorktown Heights, New York since 1978. His research interests include modeling and analysis of computer performance and statistical analysis of simulation output. Dr. Heidelberger is an Associate Editor of *Operations Research*. He is a member of the Operations Research Society of America, the Association for Computing Machinery and the IEEE.

Philip Heidelberger
IBM Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598
(914) 789-7156

PERWEZ SHAHABUDDIN received his B.Tech. degree in Mechanical Engineering from the Indian Institute of Technology, Delhi, in 1984. After working for a year in Engineers India Limited, India, he joined the Ph.D. program in Operations Research at Stanford University, California. His research interests include simulation techniques, queueing theory and diffusion processes. He is a student member of the Operations Research Society of America.

Perwez Shahabuddin
Department of Operations Research
Stanford University
Stanford, California 94305
(415) 723-9383