

Varaince reduction in mean time to failure simulations

Perwez Shahabuddin **1
Victor F. Nicola *
Philip Heidelberger *
Ambuj Goyal *
Peter W. Glynn **2

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

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

ABSTRACT

We describe two variance reduction methods for estimating the mean time to failure (MTTF) in Markovian models of highly reliable systems. The first method is based on a ratio representation of the MTTF and employs importance sampling. The second method is based on a hybrid simulation/analytic technique where the number of simulated transitions are reduced by computing partial results analytically. Experiments with a large example show the effectiveness of both techniques for highly reliable systems.

1. INTRODUCTION

The requirement for highly reliable systems, such as airborne computing systems, is increasing the importance of reliability and mean time to failure (MTTF) prediction during the design phase of these systems. These systems are different than highly available systems, such as banking or airline reservation systems, where continuous operation (availability) is more important. In highly reliable systems any system failure during the mission causes a mission failure. While such systems can typically be modeled as Markov chains (see, e.g., [6]), the size of the corresponding Markov model increases rapidly

with complexity of the system. Thus conventional 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 the measures of interest.

Our goal is to obtain variance reduction methods which are applicable to a broad class of models. Specifically, we are interested in models defined by the reliability and availability modeling language described in [7], so that the techniques can be implemented in a software package and made available to designers in an automatic and transparent fashion. A typical system contains multiple component types with redundant units for each component type. Failure of these systems is usually caused by exhaustion of redundancy or by a combination of component failures leading to a system failure. Failed components may be repaired. If all components are repairable and component failures are Poisson, then a regen-

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

² This work was partially performed while this author was a visiting scientist at IBM Research

erative state for the system (see, e.g., [4]) is the state where all units of all component types are operational. If, in addition, all repair times are exponentially distributed, then, typically, a continuous time Markov chain is obtained. As recommended in [11], we transform the continuous time chain into an appropriate discrete time Markov chain. In [8] the estimation of steady state availability for such systems was discussed. In this paper, we discuss estimation of the MTTF. Specifically, we consider an importance sampling method (see, e.g., [9]) and a hybrid simulation/analytic method.

Importance sampling for rare event simulation has been successfully used in [2], [3], [8], [12], [13] and [14]. Proper selection of the importance sampling distribution makes the rare events more likely to occur which results in a variance reduction. The key, of course, is to choose a good importance sampling distribution. The theory of large deviations was used in [3], [13], and [14] to select an effective distribution for problems arising in Markov chains with “small increments”, random walks, and queueing systems, respectively.

Effective heuristics were used in [2] and [12] to select importance sampling distributions for steady state availability estimation in large machine repairmen-like models. For example, four orders of magnitude reduction in variance was reported in [2] for a system with 70 components (the Markov chain for this model has 2^{70} states). For regenerative systems, steady state performance measures can be expressed as a ratio. In [2], a single importance sampling distribution was used to estimate both the numerator and denominator of this ratio. The distribution used in [2] is dynamic in the sense that it does not correspond directly to a time homogeneous Markov chain. We call this technique Dynamic Importance Sampling (DIS). In [8], different dynamic importance sampling distributions were used to estimate the numerator and denominator inde-

pendently which resulted in additional variance reduction (actually, importance sampling was not used at all to estimate the denominator). We call this technique Measure Specific DIS (MSDIS).

The importance sampling technique discussed in previous papers does not yield significant variance reductions for estimating the MTTF using the standard simple estimator. An intuitive reason for this is as follows. When we make failure events occur more often by choosing an appropriate importance sampling distribution, the value of the estimator ends up smaller than the actual value and, in addition, the likelihood ratio is less than one. This actually ends up producing variance rather than reducing it. However, if we formulate the MTTF problem as a ratio of two expectations (as shown in Section 2), then significant variance reductions can be achieved using importance sampling. In Section 2, both DIS and MSDIS methods are used to estimate the ratio representation of the MTTF. The MSDIS sampling scheme proposed in [8] gives the best results.

In Section 3 we outline a hybrid simulation/analytic method described in [10] and describe its application in estimating the MTTF in simulations of highly reliable systems. In highly reliable systems, a large number of transitions typically occur before the system fails. The hybrid method transforms the original Markov chain into a different, but related, Markov chain. The expected value of an appropriately defined functional of the transformed Markov chain is identical to the MTTF. Furthermore, the expected number of transitions until failure is less in the transformed chain than in the original chain. This reduction usually translates into a variance reduction since, for a given number of total transitions simulated, more failures occur in the transformed chain than in the original chain.

In Section 4, we summarize the results, discuss the relative merits of the two methods and indicate future areas of research.

2. RATIO METHOD

Let $\{Y_s, s \geq 0\}$ be a continuous time Markov chain (CTMC) with state space $E = \{0, 1, 2, \dots, N\}$ (N may be either finite or infinite), and transition rate matrix $Q = (q(i, j))$. Let $q(i) = -q(i, i) = \sum_{j \neq i} q(i, j)$ denote the total rate out of state i . Pick a starting state, say state 0, set $Y_0 = 0$ and for any set B let α_B be the first entrance time of the CTMC to B . In a reliability model, we generally take state 0 to be the state for which all components are operational. Let $F \subset E$ be a set of states (e.g., in a reliability model they might be the set of states where the system is failed) for which $0 \notin F$. If state 0 is a regenerative state and the set F represents a rare set of states, then Walrand [14] applies the approximation $E[\alpha_F] \approx E[\alpha_0] / P\{\alpha_F < \alpha_0\}$ and then uses direct simulation to estimate the numerator $E[\alpha_0]$ and importance sampling to estimate the denominator $P\{\alpha_F < \alpha_0\}$. (Since F is a rare set of states, $P\{\alpha_F < \alpha_0\}$ is very close to zero.) In fact for any CTMC we have the exact relationship

$$E[\alpha_F] = \frac{E[\min(\alpha_0, \alpha_F)]}{P\{\alpha_F < \alpha_0\}} \quad (2.1)$$

which is easily obtained by considering the two cases $\{\alpha_0 < \alpha_F\}$ and $\{\alpha_0 > \alpha_F\}$, and applying the Markov property at time α_0 if $\alpha_0 < \alpha_F$.

Given Equation 2.1 we can construct what appears to be a new estimator for $E[\alpha_F]$. The new estimator is a ratio estimator for the right hand side of Equation 2.1. The observations for this estimator are obtained by sampling regenerative cycles up to $\min(\alpha_0, \alpha_F)$. Perhaps surprisingly, it turns out that this new estimator has the same asymptotic variance as the naive esti-

mator. The reason is that the naive estimator is effectively identical to the new estimator, in the sense that the two estimators are probabilistically identical at the instant at which failures occur. The difference comes from the fact that Equation 2.1 is formulated in terms of the probability of the rare event ($P\{\alpha_F < \alpha_0\}$) which can be estimated effectively using importance sampling.

As recommended in [5] and [11], we simulate only the embedded discrete time Markov chain $\{X_n, n \geq 0\}$ and use deterministic holding times $h_i = 1/q(i)$. The transition matrix of $\{X_n, n \geq 0\}$ is $P = (P_{ij})$ where $P_{ij} = q(i, j)/q(i)$ for $i \neq j$ and $P_{ii} = 0$.

As in [8], we consider two methods of estimating the ratio in Equation 2.1. The first method, called dynamic importance sampling (DIS), uses the same importance sampling distribution (and the same sample paths) to estimate both the numerator and denominator of Equation 2.1. The second method, called measure specific dynamic importance sampling (MSDIS), uses different importance sampling distributions for the numerator and denominator which are then estimated independently. In fact, like Walrand, we use direct simulation (no importance sampling) to estimate the numerator since very stable estimates of $E[\min(\alpha_0, \alpha_F)]$ can be obtained via direct simulation. In addition, for such highly reliable systems, the expected number of transitions in a regenerative cycle is typically small. Expressions for the resulting asymptotic variances are given in [8]; they are closely related to the asymptotic variance obtained when applying the regenerative method.

The advantage of MSDIS over DIS stems from the fact that, in MSDIS, by using different changes of measure for estimating the numerator and denominator, we can reduce the variance of each estimator, individually, as much as possible. In

DIS a compromise has to be reached because using a change of measure which reduces the variance of one estimator might increase the variance of the other (in addition to the uncertain effect on the covariance).

We use the specific importance sampling method suggested in [8] for simulating models of highly reliable systems of the general type described in the Introduction. Basically, we pick a change of measure which will take the system on its most likely path to a failure state from its current state. In a given state during the simulation where both failure and repair transitions exist, we use the total combined probability of all failure transitions as p' . Individual failure transitions could be selected from a discrete distribution which is in the ratio of their individual failure rates. However, we give more importance to those failure transitions which correspond to the component types which have largest number of components failed, by selecting them with a total conditional probability p'' . Thus, once a failure in a given component type occurs, the redundancy in that type is exhausted quickly. If all component types have the same number of components failed, then the above selection rule is not used. We experimented with the values of p' and p'' for a large example.

To test the method, we considered a model having ten types of components, with two components of each type. Each component type has a failure rate of λ and repair rate of μ . There is one repairman in the system who services failed components with a service in random order queueing discipline. The system is functioning if at least one component of each type is functioning.

Because the numerator and denominator are simulated independently, one must allocate computer time to estimating each

of these quantities when using MSDIS. A simple optimization problem can be formulated to minimize the asymptotic variance. The optimal fraction of time allocated to the numerator is $\beta/(1 + \beta)$ where $\beta = \sqrt{t_N} \sigma_N / (E[\alpha_F] \sqrt{t_D} \sigma_D)$ and t_N (t_D) is the expected time to generate a sample for the numerator (denominator) and σ_N (σ_D) is the standard deviation of a sample for the numerator (denominator). In this application setting, the optimal asymptotic allocation usually suggests devoting so little time to the numerator that, even in moderate sized samples, unstable estimates of its variance may be obtained. Therefore, for practical considerations, we always allocate at least 10% of the run to estimating the numerator.

In the example, a total of 4,000,000 events were simulated for each case. For MSDIS, 400,000 events were used to simulate the numerator and 3,600,000 were used to simulate the denominator (in accordance with our 10% rule of thumb). The results of our experiment are given in Table 1 for different values of λ with $\mu = 1$. As can be seen there is considerable reduction in the confidence interval widths from Direct to DIS to MSDIS. As the values of p' and p'' are increased we get a considerable variance reduction for MSDIS. Eventually for very large values of p' and p'' the variance of the estimator starts increasing. In fact, for these values the variance estimator took longest to stabilize.

The above heuristic works extremely well for balanced systems where each type of component has approximately the same failure and repair rates and about the same amount of redundancy, such as the one described above, and we have found that it can also be very effective for unbalanced systems. Different selection rules for p'' may have to be used to get greater efficiency for unbalanced systems.

Table 1

Mean Time to Failure Estimates for the
Fault-Tolerant System Example with $\mu = 1$.
Simulation Run Lengths = 4,000,000 Events

	$p' = p''$	MTTF with Half-Widths of 99% Confidence Intervals	
		DIS	MSDIS
$\lambda=10^{-3}$	0.5	49304 ± 288	49303 ± 198
Exact=49257	0.9	49215 ± 427	49246 ± 51
Direct=47745 ± 2714	0.99	51383 ± 4158	49170 ± 243
$\lambda=10^{-5}$	0.5	4.9923 × 10 ⁸ ± 3.0 × 10 ⁶	4.9978 × 10 ⁸ ± 2.1 × 10 ⁶
Exact=4.999 × 10 ⁸	0.9	5.0093 × 10 ⁸ ± 4.4 × 10 ⁶	5.0000 × 10 ⁸ ± 6.3 × 10 ⁵
Direct=5.2642 × 10 ⁸ ± 3.1 × 10 ⁸	0.99	5.0309 × 10 ⁸ ± 1.3 × 10 ⁷	4.9999 × 10 ⁸ ± 1.8 × 10 ⁵
	0.999	5.0084 × 10 ⁸ ± 4.1 × 10 ⁷	4.9993 × 10 ⁸ ± 4.0 × 10 ⁴
	0.9999	4.5994 × 10 ⁸ ± 1.2 × 10 ⁸	5.0002 × 10 ⁸ ± 3.4 × 10 ⁴

3. HYBRID METHOD

The hybrid method we consider simulates a different Markov chain with the goal of reducing the number of transitions until failure. The method is described more generally in [10] and is based on the equations for the mean accumulated reward until absorption into a set of states (see, e.g., [1] or [11]).

A certain amount of notation is inevitable. For sets A and B and any matrix M , let $M_{AB} = (M_{ij} : i \in A, j \in B)$ and for a vector x let $x_A = (x_i : i \in A)$. Let $\{X_n, n \geq 0\}$ denote a discrete time Markov chain with a finite state space E and transition matrix $P = (P_{ij})$. Let A^c denote the complement of the set A and let τ_B be the first entrance time to the set B . Define the random variable $Y = \sum_{n=0}^{\tau_A-1} f(X_n)$. Let y be the vector whose i 'th component is $y_i = E[Y | X_0 = i]$, and let P' be the matrix with elements $P'_{ij} = P\{X_{\tau_A} = j | X_0 = i\}$ for $i \in A$ and $j \in A^c$. Then y_A satisfies the system of linear equations $y_A = f_A + P_{AA}y_A$ where the vector f has i 'th component $f_i = f(i)$. Assuming that $(I_{AA} - P_{AA})^{-1}$ exists, then

$$y_A = (I_{AA} - P_{AA})^{-1}f_A, \tag{3.1}$$

$$P'_{Aj} = (I_{AA} - P_{AA})^{-1}P_{Aj}, \quad j \in A^c.$$

These equations form the basis for the hybrid method. Let F denote the set of states for which the system is considered failed, S_k be the set of states for which fewer than k components have failed, and $F_k = S_k^c$; we assume $S_k \cap F = \emptyset$. Consider a new Markov chain $\{X'_n, n \geq 0\}$ with transition matrix P' and reward vector f' where $P'_{ij} = P\{X_{\tau_A} = j | X_0 = i\}$ for $i \in S_k, j \in F_k$ (note that $P'_{ij} = 0$ for $i, j \in S_k$), $P'_{ij} = P_{ij}$ for $i \in F_k, f'(i) = E[\sum_{n=0}^{\tau_A-1} f(X_n) | X_0 = i]$ for $i \in S_k$, and $f'(i) = f(i)$ for $i \in F_k$. Let τ'_F be the first entrance time to the set F in this new Markov chain. If $Z = \sum_{n=0}^{\tau_A-1} f(X_n)$ and $Z' = \sum_{n=0}^{\tau'_F-1} f'(X'_n)$, then $E[Z | X_0 = i] = E[Z' | X'_0 = i]$, i.e., one could simulate either $\{X_n\}$ or $\{X'_n\}$ to estimate the expected cumulative reward until failure. Note that if $f(i) = h_i$ where h_i is the mean holding time in state i in the associated continuous time Markov Chain, then $E[Z]$ is the MTTF (for an appropriately defined initial state). Application of this method requires pre-computing P'

and f' which is accomplished by solving Equations 3.1 with $A = S_k$.

The Markov chain $\{X'_n\}$ has the following interpretation: whenever the chain enters a state in S_k with fewer than k failed components it jumps in one transition to a state in F_k with at least k failed components (and the reward vector is modified to be the expected reward earned by the original chain until exit from S_k). Thus many transitions may be saved since, in models of highly reliable systems, the expected number of transitions to go from S_k to F_k may be large. On the set F_k , $\{X'_n\}$ and $\{X_n\}$ are stochastically identical.

As discussed in [10], if transitions in the two chains require the same amount of computer time, then an approximation to the variance reduction is the ratio $VR = \text{Var}[Z'] E[\tau'_{F}] / \text{Var}[Z] E[\tau_F]$. If, as has been our experience in reliability simulations, $\text{Var}[Z'] \approx \text{Var}[Z]$, then $VR \approx E[\tau'_{F}] / E[\tau_F]$. If, in the original chain, the probability of a failure before a repair is $O(\epsilon_j) \approx 0$ when there are j failed components, then the expected number of transitions until reaching a state with k failed components is $O(1 / \prod_{j=1}^{k-1} \epsilon_j)$ and thus we expect the hybrid method to produce a variance reduction of $VR \approx O(\prod_{j=1}^{k-1} \epsilon_j)$. If $\epsilon_j = \epsilon$ for all j , then $VR \approx O(\epsilon^{k-1})$. This will be verified for several examples.

Consider now the general class of models described in the Introduction with M types of components and redundancy. Figure 1 depicts the transition structure on the set S_2 and its interface to the set F_2 . The set $S_2 = \{0, 1, \dots, M\}$ with state 0 corresponding to no component failures and state i corresponding to one component of component type i . Because of this special structure, Equations 3.1 can be solved analytically for the case $k = 2$: $f'_0 = (f_0 + \sum_{j=1}^M P_{0j} f_j) / (1 - \sum_{j=1}^M P_{0j} P_{j0})$ and $f'_i = f_i + P_{i0} f'_0$ for $1 \leq i \leq M$. Similarly, if E_{ij} is the probability of exiting S_2 via state j given that $X_0 = i$, then $E_{0j} = P_{0j}(1 - P_{j0}) / (1 - \sum_{j=1}^M P_{0j} P_{j0})$ for $1 \leq j \leq M$, $E_{ij} = P_{i0} E_{0j}$,

for $1 \leq i \neq j \leq M$, and $E_{ii} = (1 - P_{i0}) + P_{i0} E_{0i}$, for $1 \leq i \leq M$. Transitions from a state $i \in S_2$ to a state $k \in F_2$ can be simulated as follows: first an exit state $j \in S_2$ is selected with probability E_{ij} , then a state $k \in F_k$ is selected with probability $P_{jk} / (1 - P_{j0})$.

For $k > 2$, systems of equations of size $|S_k|$ need to be solved. This may be feasible for small values of k , but the number of states in S_k depends on the queueing discipline for the repairmen. For example, if each type of component has its own repairman, then $|S_3| = O(M^2)$.

We now consider the application of the hybrid method for estimating the MTTF in two examples. The first example is a model of a fault-tolerant system with two types of components and two units of each type. All units are initially operational. Units of component types 1 and 2 fail at rates λ_1 and λ_2 , respectively. They are repaired according to a FCFS discipline by separate repair facilities at rates μ_1 and μ_2 , respectively. The system is operational if at least one unit of any type is operational. The resulting Markov chain, while having only nine states, is illustrative of the performance of the hybrid method. We applied the hybrid method with $k = 2, 3$ and computed the variance ratio numerically using the techniques described in [11]. For $\mu_1 = \mu_2 = 100$ and $\lambda_1 = \lambda_2 = 1$, we obtained $VR = 0.03$ for $k = 2$ and $VR = 0.0004$ for $k = 3$. If we increase the failure rates by setting $\lambda_1 = \lambda_2 = 10$, then $VR = 0.25$ for $k = 2$ and $VR = 0.028$ for $k = 3$. Most of the variance reductions are due to the reduction in the expected number of transitions until failure and are consistent with the prediction that $VR \approx O(\prod_{j=1}^{k-1} \epsilon_j)$.

The second example is the model with ten types of components and two components of each type described in Section 2. The repair rate was set to be $\mu = 1$ and the component failure rate was $\lambda = 0.001$. This model was simulated with

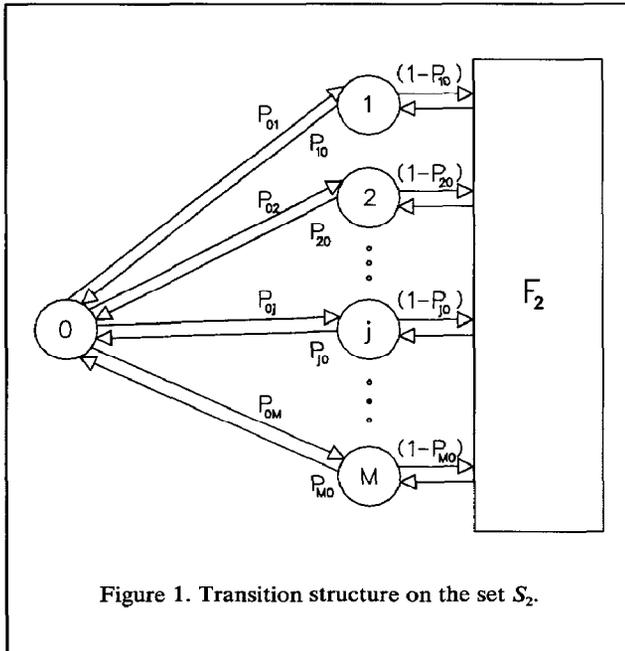


Figure 1. Transition structure on the set S_2 .

$k = 2$ and the estimated variance ratio was $VR = 0.036$. For these parameter settings the predicted variance reduction is $\varepsilon_1 = 0.019$ which is of the same order of magnitude as the observed variance reduction. The apparent factor of two discrepancy between the actual and predicted variance reductions is due to the fact that hybrid method (as we implemented it) requires two transitions, rather than one, to exit S_2 : one to select the exit state in S_2 and another to select the entrance state in F_2 . For a general purpose simulator this represents a realistic implementation of the method. For this example, the best variance reduction using MSDIS yields a variance ratio of 0.00035 (from Table 1, with $p' = p'' = 0.9$).

4. SUMMARY

In this paper we described two methods for estimating the MTTF in a class of Markovian models of highly reliable systems. The first (ratio) method relies on a ratio representation

of the MTTF. The numerator of the ratio is estimated using direct simulation while the denominator is, independently, estimated using importance sampling. The second (hybrid) method combines simulation and numerical techniques. The hybrid method requires some preliminary computations, although in its simplest form these can be done analytically. A different, but related, Markov chain is then simulated. In this Markov chain, the expected number of transitions until failure may be greatly reduced. Both methods produced significant variance reduction in simulations of a large model.

Although the two methods are not necessarily mutually exclusive, the hybrid method, while effective, has a number of disadvantages that reduce its utility when one considers its implementation into a broadly applicable availability simulation package. First, it requires preliminary numerical computations. Second, the method as described here is specifically designed to estimate the MTTF and requires simulation of its own stochastic process. Therefore, the output from this method may not be useful for estimating other performance measures, such as the steady state unavailability. Third, the method is best suited for situations in which the amount of redundancy is low since, in this case, the expected number of transitions to failure in the modified chain remains reasonable. More specifically, the practical implementation of the hybrid method with $k = 2$ uses direct simulation whenever the system is in a state with two or more failed components, whereas the ratio method with importance sampling quickly moves the system through states with more than two failed components to the failed state. Therefore the variance reductions using importance sampling are potentially (and in practice) greater than those using the hybrid method. It might appear that since the ratio method requires simulation of two processes, it too would be of limited value in estimating

other performance measures. However, the technique described in [8] for estimating steady state unavailability also independently simulates the numerator and denominator of a ratio (in that case importance sampling is used for the numerator and direct simulation is used for the denominator). Furthermore, the same importance sampling distribution can be used in estimating both the MTTF and the steady state unavailability so that both of these performance measures can be estimated simultaneously. In addition, it should be possible to apply this technique to simultaneously estimate multiple performance measures at different parameter settings (e.g., different failure rates), since this is an inherent capability of importance sampling.

Therefore, our current research is directed towards further experimentation with and extensions of the importance sampling approach, including improved selection of importance sampling distributions for unbalanced systems, its application in estimating the failure time and interval availability distributions, as well as extensions to gradient estimation and simulations of non-Markovian models.

REFERENCES

- [1] Chung, K.L. (1967). *Markov Chains With Stationary Transition Probabilities*, Second Edition. Springer-Verlag, New York.
- [2] 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.
- [3] Cottrell, M., Fort, J.C. and Malgouyres, G. (1983). Large Deviations and Rare Events in the Study of Stochastic Algorithms. *IEEE Transactions on Automatic Control* **AC-28**, 907-920.
- [4] Crane, M.A. and Iglehart, D.L. (1975). Simulating Stable Stochastic Systems, III: Regenerative Processes and Discrete Event Simulations. *Operations Research* **23**, 33-45.
- [5] Fox, B.L. and Glynn, P.W. (1986). Discrete-Time Conversion for Simulating Semi-Markov Processes. *Operations Research Letters* **5**, 191-196.
- [6] 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.
- [7] Goyal, A. and Lavenberg, S.S. (1987). Modeling and Analysis of Computer System Availability. *IBM Journal of Research and Development*, **31** 6, 651-664.
- [8] Goyal, A., Heidelberger, P. and Shahabuddin, P. (1987). Measure Specific Dynamic Importance Sampling for Availability Simulations. *1987 Winter Simulation Conference Proceedings*. A. Thesen, H. Grant and W.D. Kelton (eds.). IEEE Press, 351-357.
- [9] Hammersley, J.M. and Handscomb, D.C. (1964). *Monte Carlo Methods*. Methuen, London.
- [10] Heidelberger, P. (1979). A Variance Reduction Technique That Increases Regeneration Frequency. *Current Issues in Computer Simulation*. N.R. Adam and A. Dogramaci (eds.). Academic Press, Inc., 257-269.
- [11] Hordijk, A., Iglehart, D.L. and Schassberger, R. (1976). Discrete Time Methods for Simulating Continuous Time Markov Chains. *Adv. Appl. Prob.* **8**, 772-788.
- [12] Lewis, E.E. and Bohm, F. (1984). Monte Carlo Simulation of Markov Unreliability Models. *Nuclear Engineering and Design* **77**, 49-62.
- [13] Siegmund, D. (1976). Importance Sampling in the Monte Carlo Study of Sequential Tests. *The Annals of Statistics* **4**, 673-684.
- [14] 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

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

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. Heidelberg is the Program Chairman of the 1989 Winter Simulation Conference and an Associate Program Chairman of the 1988 Winter Simulation Conference. He 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 Heidelberg
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598
(914) 789-7156

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 and manages Fault-Tolerant Systems Theory Group at the IBM Thomas J. Watson Research Center in Yorktown Heights, New York. His research interest include fault-tolerant computing, computer architecture, communication networks and performance evaluation. Dr. Goyal is a member of the IEEE Computer Society.

Ambuj Goyal
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704

Yorktown Heights, New York 10598
(914) 789-7504

VICTOR F. NICOLA received the B.S. and M.S. degrees in electrical engineering from Cairo University, Egypt, and Eindhoven University of Technology, The Netherlands, respectively. In 1979 he joined the scientific staff of the Measurement and Control Group at Eindhoven University. From 1983 to 1986 he was at Duke University, Durham, North Carolina, where he received the Ph.D. degree in Computer Science. Currently he is a Research Staff Member at the IBM Thomas J. Watson Research Center in Yorktown Heights, New York. His main research interests are in the areas of performance and reliability modeling of computer systems and fault-tolerance.

Victor Nicola
IBM Research Division
Thomas J. Watson Research Center
P.O. Box 704
Yorktown Heights, New York 10598
(914) 789-7502

PETER GLYNN is an associate professor in the Department of Operations Research at Stanford University. He received a B.S. from Carleton University in 1978 and a Ph.D. in operations research from Stanford University in 1982. From 1982 to 1987, he was an associate professor in the Industrial Engineering Department of the University of Wisconsin at Madison. His current research interests include simulation, queueing theory and applied probability. He is a member of ORSA, TIMS, IMS and the Statistical Society of Canada.

Peter Glynn
Department of Operations Research
Stanford University
Stanford, California 94305
(415) 725-0554