Proceedings of the 1995 Winter Simulation Conference
ed. C. Alexopoulos, K. Kang, W. R. Lilegdon, and D. Goldsman

PERFORMANCE ANALYSIS OF FUTURE EVENT SETS

Halim Damerdji

Department of Industrial Engineering
North Carolina State University
Raleigh, NC 27695–7906, U.S.A.

Peter W. Glynn

Department of Operations Research
Stanford University
Stanford, CA 94305–4022, U.S.A.

ABSTRACT

The linked list and indexed list future event sets are investigated here. The interaction hold model and the Jackson network model are the underlying stochastic models considered. For the interaction hold model and for the (doubly) linked list, we find, for example, the mean number of key comparisons performed in order to find a record's insertion point into the list; this is useful when deciding whether to scan from the head or the tail of the list. The distribution of the relative position of the to-be-inserted record is also obtained; for indexed lists this is helpful when deciding the number of sublists and position(s) of the middle pointer(s). The Jackson network model has a realistic event logic, but events are restricted to be exponentially distributed. Because the stationary probabilities can be computed for this model, it is then possible to evaluate and compare the (steady-state) performance of certain future event sets (e.g., linked lists scanned from the head or the tail).

1 INTRODUCTION

In a discrete-event simulation, the future event set (FES), or also simulation calendar, is the data structure that holds the records of the events scheduled to occur in the future. Part of the computational effort in a simulation is devoted to the manipulation of the FES as records get removed, inserted, updated and looked at (e.g., for checking certain conditions). See Devroye (1986). Manipulation of the FES must be efficient, especially for large-scale simulations. Because of the various demands a complex simulation puts on the FES, it is not evident that in such a varied environment the more modern algorithms (e.g., splay trees, binomial queues, etc.) perform uniformly better than the traditional ones (e.g., indexed lists, heaps, etc.). Jones (1986) contains an extensive empirical investigation.

It is clearly of interest to study the performance of the classical data structures (e.g., linked lists) in a simulation environment, but the problem is difficult. There are two general approaches to analyzing the performance of a FES. The first one is to run various real-world simulations and perhaps arrive to conclusions based upon the empirical experiments. The other approach is to assume that the underlying stochastic system is simple enough so that an analysis of the FES performance is tractable. This is the approach taken here. The two stochastic models considered are the interaction hold model and the Jackson network model. The former has a simple event logic but allows for general probability functions, while the latter has a more realistic event logic, but events are restricted to be exponentially distributed. We will be looking at (doubly) linked lists and indexed lists only.

Under the hold model assumption, first considered for FES performance evaluation by Vaucher (1977), the underlying stochastic process consists of a fixed number L of independent renewal processes with identical event lifetime distributions. Each time an event triggers a transition (at the end of its lifetime), its record is removed from the list, the event is then generated anew and its record inserted back into the FES. The hold model where events may have different distribution functions was considered by McCormack and Sargent (1981), and was coined the *interaction* hold model. A multi-server queue with possibly nonequivalent servers and with an infinite supply of customers readily available fits the interaction hold model framework. But, as previously mentioned, the interaction hold model does have a simple event logic. For example, the FES will have constant size L. For the Jackson network model, introduced in Damerdji and Glynn (1995), it is possible to numerically compute the stationary probabilities, and one can then evaluate the performance of certain data structures (e.g., linked list) under this model.

In a simulation, the FES is typically kept sorted, based upon the events' occurrence times. Whenever a record is to be inserted into the FES, a search is performed to find its insertion point (so the FES stays sorted). The search is carried out by comparing the to-be-inserted event's occurrence time with the occurrence times of the successive events in the FES until the correct insertion point is found. The removal of the current event's record and the actual update of the FES are fixed overheads. Algorithm efficiency is measured here in terms of number of comparisons performed to find the insertion point at a transition.

For the interaction hold model, the clock-vector across transitions is viewed as a general state space Markov chain, whose stationary distribution can be derived (see Section 2). The mean of the number of comparisons per transition for a linked list is provided, also in Section 2. It is found, for example, that if the events are all exponentially distributed, it is strictly better to scan from the head of the list (unless the rates are all equal in which case it is equivalent to scan from the tail of the list). Indexed lists, i.e., lists with sublists or also buckets, are investigated under the interaction hold model assumption. Here, we only consider sublists that contain fixed numbers of records. In the case of an indexed list with two sublists and with exponentially distributed events, explicit expressions for the number of comparisons per transition in terms of the event rates are given; see Section 2. Section 3 contains an example.

The Jackson network model is discussed in Section 4. An expression for the expected number of comparisons per transition in terms of the stationary probabilities is given for a general FES. An example illustrates our point. The Jackson network model is introduced in the hope that something might be learned (e.g., is it better to scan a linked list from the head rather than the tail?), and then that the exponential-case conclusions might carry over to the general case, at least to some extent. A limited empirical investigation is presented in Section 4. Section 5 is the conclusion.

2 THE INTERACTION HOLD MODEL

The interaction hold model is now described in more detail. We will have clocks keep track of the events' residual lifetimes, i.e., the times remaining for the events to trigger the transition. For each event i, let $G_i(.)$ be its distribution function and $(\tau_{i,j} : j \geq 1)$ its sequence of lifetimes. Let $X_n \equiv (X_{n,1}, \ldots, X_{n,L})$ denote the clock-vector reading at the n'th transition (in fact, at the instant after the transition). The event that will trigger the next transition has the smallest clock reading. If event i, say, triggers the n'th transition, the clocks corresponding to the other events all get updated, while event i gets generated anew from $G_i(.)$. Its clock reading will be $\tau_{i,N_i(n)}$, where $N_i(n)$ is the number of lifetimes of event i that have been observed by the n'th transition. We have that

$$X_n = \sum_{i=1}^{L} I\left[X_{n-1,i} = \min_{j=1,\ldots,L} X_{n-1,j}\right] \Big(X_{n-1,1} - X_{n-1,i}, \ldots, \tau_{i,N_i(n)}, \ldots, X_{n-1,L} - X_{n-1,i}\Big),$$

where $I[.]$ is the indicator event. Since the triggering event is generated anew independently of the past, it is clear that $\{X_n : n \geq 0\}$ is a Markov chain, with state space $\mathcal{R}_+^L = \{(y_1, \ldots, y_L) : y_i > 0, 1 \leq i \leq L\}$. It is shown in Damerdji and Glynn (1995) that this general state space Markov chain is recurrent in the sense of Harris (see Meyn and Tweedie 1993 for a definition), and that its stationary distribution is given by

$$\pi(dx_1, \ldots, dx_L) = \sum_{i=1}^{L} \frac{\lambda_i}{\sum_{\ell=1}^{L} \lambda_\ell} G_i(dx_i) \prod_{j \neq i} \lambda_j \overline{G}_j(x_j) dz_j,$$

where $\overline{G}_j(.) \equiv 1 - G_j(.)$. Let P_π be the probability measure of the Markov chain on its infinite path space with initial distribution π.

As previously mentioned, efficiency of a FES is measured here in terms of number of comparisons (in order to find the correct insertion point) per scan. Doubly linked lists are considered first. Let W_k be the random variable that denotes the relative position (minus one) of the event-record inserted at the k'th transition. Let also $Y_{\mathrm{F},k}$ (respectively, $Y_{\mathrm{B},k}$) represent the number of comparisons (minus one) performed at the k'th transition when scanning from the head (resp., tail) of the list. For example, if at the k'th transition, the generated event has the second smallest time, then $W_k = 1$, $Y_{\mathrm{F},k} = 1$, and $Y_{\mathrm{B},k} = L - 1 - Y_{\mathrm{F},k} = L - 2$. We obtain the following.

Proposition 1.

$$E_\pi Y_{\mathrm{F},1} = \sum_{i=1}^{L} \frac{\lambda_i}{\sum_{\ell=1}^{L} \lambda_\ell} \sum_{j \neq i} \lambda_j \int_0^\infty \overline{G}_i(t) \overline{G}_j(t) dt,$$

where E_π is the expectation under P_π.

This result is a simplified version of an equation that appears in McCormack and Sargent (1981). For identical distributions, see Vaucher (1977). Our goal now is to investigate whether it is more efficient to scan the list from the head or the tail. Certain properties of the distributions may be determining in some cases. This is discussed next.

Definition 1. *A distribution function G, with finite mean $1/\lambda$, is said to be new better (resp., worse) than used in expectation (NBUE (NWUE)) if*

$$\int_x^\infty \overline{G}(t)dt \leq \ (\geq)\ (1/\lambda)\overline{G}(x) \quad \text{for all } x \geq 0.$$

From Barlow and Proschan (1975), it follows that if two distribution functions G_i and G_j, with means $1/\lambda_i$ and $1/\lambda_j$, are NBUE (resp., NWUE), then

$$\int_0^\infty \overline{G}_i(t)\overline{G}_j(t)dt \geq \ (\leq)\ \frac{1}{\lambda_i + \lambda_j}.$$

Because the expression in Proposition 1 contains such an integral, the following corollary ensues.

Corollary 1. *If the event lifetimes are all NBUE (resp., NWUE), then*

$$E_\pi Y_{\mathrm{F},1} \geq \ (\leq)\ \frac{1}{\sum_{\ell=1}^L \lambda_\ell} \sum_{i=1}^L \sum_{j \neq i} \frac{\lambda_i \lambda_j}{\lambda_i + \lambda_j}.$$

Equality holds when the events are all exponentially distributed (since the exponential distribution is both NBUE and NWUE). When the distributions are all NBUE, the lower bound of Corollary 1 complements McCormack and Sargent (1981), who provide an upper bound (for general distributions). Consider the following lemma.

Lemma 1. *For any positive and finite numbers λ_1, ..., λ_L, we have that*

$$\frac{1}{\sum_{\ell=1}^L \lambda_\ell} \sum_{i=1}^L \sum_{j \neq i} \frac{\lambda_i \lambda_j}{\lambda_i + \lambda_j} \leq \frac{L-1}{2},$$

with equality if and only if $\lambda_1 = \ldots = \lambda_L$.

An immediate consequence of Proposition 1 and Lemma 1 is the following.

Corollary 2. *If the events are all exponentially distributed, it is strictly better to scan from the head of the list, unless the events have identical rates in which case it is equivalent to scan from the back.*

Because the stationary distribution π of the Markov chain is known, it is possible to compute the distribution $P_\pi[W_1 = .]$ of the relative position of the record to be inserted. See Damerdji and Glynn (1995) for a general expression. If the events are all exponentially distributed, the distribution of W_1 depends only upon the event rates. Let $C(g,h)$ be the number of possible combinations of size h out of a set of size g, $Q_i(j,n)$ be the j'th set of size n out of $\{1,\ldots,L\} - \{i\}$, $R_j(i,n)$ be its complement in $\{1,\ldots,L\} - \{i\}$, and $T_{b,a,q}(i,j,n)$ be the b'th subset of size a out of $Q_j(i,n) - \{q\}$.

Proposition 2. *If the events are all exponentially distributed, we have, for $n = 1,\ldots,L-1$, that*

$$P_\pi[W_1 = n] = \sum_{i=1}^L \frac{\lambda_i^2}{\sum_\ell \lambda_\ell} \Bigg\{ \sum_{j=1}^{C(L-1,n)} \Bigg[\Big(\lambda_i + \sum_{v \in R_j(i,n)} \lambda_v\Big)^{-1} \sum_{q \in Q_j(i,n)} \Big\{ \lambda_q \sum_{a=0}^{n-1} (-1)^a \bullet \Big[\sum_{b=1}^{C(n-1,a)} \Big(\lambda_i + \sum_{v \in R_j(i,n)} \lambda_v + \lambda_q + \sum_{w \in T_{b,a,q}(i,j,n)} \lambda_w\Big)^{-1} \Big] \Big\} \Bigg] \Bigg\}.$$

Also, we have that $P_\pi[W_1 = 0] = \left(\sum_{i=1}^L \lambda_i^2\right) / \left(\sum_{i=1}^L \lambda_i\right)^2$.

In an indexed list, i.e., a list with sublists, the middle pointers point to certain records (possibly dummy records) in the list chosen depending upon the strategy followed. One approach is to keep the sublists with constant sizes. Another strategy is to have the middle records chosen based upon their event times of occurrence. This was the setting of Engelbrecht-Wiggans and Maxwell (1978) and Davey and Vaucher (1980). Here, the former approach is investigated, and so each sublist will have a fixed number of records.

For simplicity of analysis, consider now an indexed list with two sublists only. Let m be the middle record, i.e., the record such that the (unique here) middle pointer is pointing to. Because of the two-way links, four search strategies are possible: the first (resp. second) sublist can be scanned forward from the head (resp. middle) pointer or backward from the middle (resp. tail) pointer. Let $Z_{\mathrm{FF},k}(m)$, $Z_{\mathrm{FB},k}(m)$, $Z_{\mathrm{BF},k}(m)$, and $Z_{\mathrm{BB},k}(m)$ be the number of comparisons (minus one) performed at the k'th transition in order to find the correct insertion point for the different strategies. Explicit expressions can be obtained. For example,

$$E_\pi Z_{\mathrm{FF},1}(m) = 1 + \sum_{n=0}^{m-2} n a(n) + \sum_{n=m-1}^{L-1} (n-m+1) a(n),$$

where $a(n) \equiv P_\pi[W_1 = n]$. See Damerdji and Glynn (1995) for the other three strategies. It is then possible to optimize with respect to m, and also compare the performances for the four strategies. An example is given in the next section. Lists with more than two sublists can be similarly investigated.

3 AN EXAMPLE

We consider here an interaction hold model with ten exponentially distributed events, split into two equal-size groups: events 1 thru 5 have rates $\lambda_1 = \ldots = \lambda_5 = 1.0$, while events 6 thru 10 have rates $\lambda_6 = \ldots = \lambda_{10}$. Three experiments are conducted: λ_6 is chosen such that $1/\lambda_6 = 1.0$ in the first one (the two groups are equal), $1/\lambda_6 = 2.0$ in the second one (the two groups are rather different), and $1/\lambda_6 = 10.0$ in the last experiment (the two groups are very different).

Table 1: Linked List

$1/\lambda_6$	1.0	2.0	10.0
$E_\pi Y_{F,1}$	4.500	4.222	2.826
$E_\pi Y_{B,1}$	4.500	4.778	6.174

For the third experiment, i.e., when the two groups are very different, it is seen from Table (1) that it is far better to scan the list from the head pointer. From Table (2), the probabilities of the relative position of the inserted record are equal in the first experiment because of the memoryless property of the exponential distribution and by symmetry. Note that the probabilities are decreasing (relative to the position) for the second and third experiments. For the third experiment for example, it is almost ten times more likely that the newly inserted record be at the top of the list rather than the bottom (0.166 vs. 0.018). This stresses again that, at least in the exponential case, it is better to scan a linked list from the head rather than the tail.

Once the distribution of W_1 is obtained, it is straightforward to compute $E_\pi Z_{FF,1}(m)$, $E_\pi Z_{FB,1}(m)$, $E_\pi Z_{BF,1}(m)$, and $E_\pi Z_{BB,1}(m)$ for $m = 3, \ldots, 9$, and optimize over m for each strategy. Let m^*_{FF}, m^*_{FB}, m^*_{BF}, and m^*_{BB} be the respective optimum. For the second and third experiments, it is found that the best strategy is forward-forward, closely followed by backward-forward. Backward-backward was the worst strategy. For the third experiment and for the forward-forward strategy, the fourth record is the optimal middle record; in that case the first sublist contains three records, while the second sublist contains seven records. The first sublist has size smaller than 50% that of the second one. Comparing the best performance for indexed list and regular linked list for the third experiment ($E_\pi Z_{FF,1}(m) = 2.299$ and

Table 2: Distribution of W_1

$1/\lambda_6$	1.0	2.0	10.0
$a(0)$	.100	.111	.166
$a(1)$	.100	.109	.164
$a(2)$	.100	.108	.160
$a(3)$	.100	.106	.153
$a(4)$	.100	.104	.140
$a(5)$	.100	.101	.095
$a(6)$	.100	.098	.052
$a(7)$	.100	.093	.028
$a(8)$	.100	.087	.020
$a(9)$	.100	.078	.018

Table 3: Indexed List with Two Sublists

$1/\lambda_6$	1.0	2.0	10.0
m^*_{FF}	6	6	4
$E_\pi Z_{FF,1}(m^*)$	3.000	2.924	2.299
m^*_{FB}	6	6	6
$E_\pi Z_{FB,1}(m^*)$	3.000	3.040	3.120
m^*_{BF}	6	6	4
$E_\pi Z_{BF,1}(m^*)$	3.000	2.957	2.312
m^*_{BB}	6	6	6
$E_\pi Z_{BB,1}(m^*)$	3.000	3.073	3.248

$E_\pi Y_{F,1} = 2.826$), it is not evident that the indexed list far outperforms the regular list given the overhead associated with maintaining the indexed list.

4 THE JACKSON NETWORK MODEL

Open and closed Jackson queueing networks are proposed as stochastic models for FES performance evaluation. As previously mentioned, these queueing network systems have a reasonably complex event logic, although only exponentially distributed event times are considered. When simulating a Markovian system such as a Jackson network, one need not even use a calendar (from the properties of the exponential distribution). This class of stochastic models is proposed in the hope that what might be learned in the exponential case extends to the general case, at least to some extent.

Only open networks are discussed here. There are d single-server stations with external arrival rate λ_i to station i. Service at station i is with rate μ_i and is on a first-in-first-out basis. Upon service completion, the customer departs the system with probability p_{i0} or is routed to station j with probability p_{ij}. The state of the system can be described by $\vec{u} = (u_1, \ldots, u_d)$, where u_i is the number of customers at station i. If e_k is the k-unit vector, the state changes from $\vec{u}$ to $\vec{u} + e_i$ if an external arrival to station i occurs, from $\vec{u}$ to $\vec{u} - e_i$ if a departure from station i to the outside system occurs, and from $\vec{u}$ to $\vec{u} - e_i + e_j$ if an end of service at station i occurs and the departing customer is routed to station j.

Let $\chi(.)$ be the stationary distribution, which is given by

$$\chi(\vec{u}) = \prod_{i=1}^{d} (1 - \alpha_i/\mu_i)(\alpha_i/\mu_i)^{u_i},$$

where $\alpha = (\alpha_1, \ldots, \alpha_d)$ is a vector solution to the traffic equations. Let A_i and S_i be a station i generic external interarrival time and service time, respectively. For any state $\vec{w}$, let $b(\vec{w})$ be the set of busy servers in that state. A general expression for the steady-state expected number of comparisons per transition for a general FES can be provided by conditioning on the first two states of the system. We have that

$$\begin{aligned}
&\mathrm{E}[\#\text{ comparisons/transition}] = \\
&\Big(\sum_{\vec{w}}\Big(\sum_{\ell=1}^{d}\lambda_\ell + \sum_{\ell\in b(\vec{w})}\mu_\ell\Big)\chi(\vec{w})\Big)^{-1}\Big\{ \\
&\sum_{\vec{u}:u_i=1;u_j=0}\ \sum_{\vec{v}=\vec{u}-e_i+e_j} \mathrm{E}[\#|S_j]\mu_i p_{ij}\chi(\vec{u}) \\
&+ \sum_{\vec{u}:u_i\geq 2;u_j\geq 1}\ \sum_{\vec{v}=\vec{u}-e_i+e_j} \mathrm{E}[\#|S_i]\mu_i p_{ij}\chi(\vec{u}) \\
&+ \sum_{\vec{u}:u_i\geq 2;u_j=0}\ \sum_{\vec{v}=\vec{u}-e_i+e_j} \mathrm{E}[\#|S_i, \text{then } S_j]\mu_i p_{ij}\chi(\vec{u}) \\
&+ \sum_{\vec{u}:u_i\geq 2}\ \sum_{\vec{v}=\vec{u}-e_i} \mathrm{E}[\#|S_i]\mu_i p_{i0}\chi(\vec{u}) \\
&+ \sum_{\vec{u}:u_i=0}\ \sum_{\vec{v}=\vec{u}+e_i} \mathrm{E}[\#|A_i, \text{then } S_i]\lambda_i\chi(\vec{u}) \\
&+ \sum_{\vec{u}:u_i\geq 1}\ \sum_{\vec{v}=\vec{u}+e_i} \mathrm{E}[\#|A_i]\lambda_i\chi(\vec{u})\Big\}.
\end{aligned}$$

We are assuming here that if an external arrival occurs to an empty station, the next interarrival time is generated and inserted into the FES before the service time. Also, when a departure from a station with a nonempty queue to an empty station occurs, the service time of the first awaiting customer is generated (and inserted) before the service time of the arriving customer.

As previously mentioned, it is possible to compute the stationary probabilities. In order to evaluate the performance of a FES, it suffices then to compute the inner expected values. A similar analysis for closed networks follows (see Damerdji and Glynn 1995). An example for an open network and for the linked list is now given.

Table 4: Linked List, Open Jackson Network

λ	1.0	1.0	1.0	1.0	1.0
μ_1	1.2	2.0	10.0	2.0	10.0
μ_2	1.2	2.0	10.0	3.0	5.0
F	1.883	1.611	1.154	1.522	1.216
B	1.893	1.722	1.645	1.700	1.650
$\hat{F}_{10\%}$	1.982	1.498	1.000	1.000	1.000
$\hat{B}_{10\%}$	1.350	1.500	1.666	1.666	1.666
$\hat{F}_{50\%}$	2.182	1.490	1.000	1.319	1.000
$\hat{B}_{50\%}$	1.399	1.527	1.666	1.574	1.666

Consider a two-station in tandem with external arrivals to the first station only. It follows then that $\lambda_2 = 0$ and $p_{12} = p_{20} = 1$. Let λ be the arrival rate to station 1, fixed to 1. For the stationary probabilities to exist, it must be that $\mu_1 > 1$ and $\mu_2 > 1$. For all the values of μ_1 and μ_2 tried, we find that it is better to scan from the head of the list rather its tail. See Table (4), where F (resp. B) denotes the expected number of comparisons per transition when scanning from the head (resp. tail) of the list. For example, for $\mu_1 = \mu_2 = 10$, it is far better to scan from the head rather than the tail ($F = 1.154$ and $B = 1.645$). When the rates are closer to one another, it is only slightly better to scan from the head.

We now empirically investigate whether results for the exponential case do extend to the nonexponential one. Uniform distributions are considered in the two experiments performed. If r is the rate of an event in the exponential case (and so $1/r$ is the mean), the uniform distributions $U(a, b)$ for that event are such $b - a = 10\%$ of the mean in the first experiment and 50% of the mean in the second. The lifetime distributions have then a small variance in the former case and a large variance in the latter. The tandem-queue was then simulated over one million transitions. Let $\hat{F}_{10\%}$, $\hat{B}_{10\%}$, $\hat{F}_{50\%}$, and $\hat{B}_{50\%}$ be the respective estimates of the expected number of comparisons per scan. See Table (4) for the results. When the rates are all close to one another (e.g., $\lambda = 1$, $\mu_1 = \mu_2 = 1.2$), it is better to scan from the back of the list. We note that when it is far more efficient to scan from the head in the exponential case, it is also better to scan from the head in the nonexponential case as well. When it is not far more efficient to scan from the head in the exponential case, we note that it is better to scan from the back of the list in the nonexponential case.

5 CONCLUSION

Performance evaluation of linked lists and indexed lists has been investigated here. The stochastic models considered, namely the interaction hold model and the Jackson network model are very different. The former has a simple event logic but allows for general distribution functions, while the latter has a more complex event logic, but distribution functions are exponential. Our analytic and empirical results suggest that in a simulation with different types of distribution functions with the short events frequently regenerated upon transition, it is often better to scan a linked list from the head rather than from the tail. However, there will also be situations where it is better to scan from the back of the list. For example, if the distributions are Uniform (100,120), Triangular (90,95,110), and Normal (120,5), it would seem better to scan from the back of the list.

REFERENCES

Barlow, R. E. and Proschan, F. (1975). *Statistical Theory of Reliability and Life Testing: Probability Models*. Holt, Rinehart and Winston, New York.

Damerdji, H. and Glynn, P. W. (1995). Limit theory for performance modeling of future event set algorithms. In preparation.

Davey, D. and Vaucher, J. (1980). Self-optimizing partitioned sequencing sets for discrete event simulations. *INFOR,* **18**, 41–61.

Devroye, L. (1986). *Nonuniform Random Variate Generation.* Springer-Verlag, New York.

Engelbrecht-Wiggans, R. and Maxwell, W. L. (1978). Analysis of the time indexed list procedure for synchronization of discrete event simulations. *Management Science,* **24**, 1417–1427.

Jones, D. W. (1986). An empirical comparison of priority-queue and event-set implementations. *Communications of the ACM,* **29**, 300–311.

McCormack, W. M. and Sargent, R. G. (1981). Analysis of future event set algorithms for discrete event simulation. *Communications of the ACM,* **24**, 801–812.

Meyn, S. P. and Tweedie, R. L. (1993). *Markov Chains and Stochastic Stability.* Springer-Verlag, London.

Vaucher, J. G. (1977). On the distribution of event times for the notices in a simulation event list. *INFOR,* **15**, 171–182.

HALIM DAMERDJI is an Assistant Professor in the Department of Industrial Engineering at North Carolina State University. His research interests are in simulation and applied stochastic processes. He is a member of InfORMS.

PETER W. GLYNN received his Ph.D. from Stanford University, after which he joined the faculty of the Department of Industrial Engineering at the University of Wisconsin-Madison. In 1987, he returned to Stanford, where he currently holds the Thomas Ford Faculty Scholar Chair in the Department of Operations Research. He was a co-winner of the 1993 Outstanding Simulation Publication Award sponsored by the TIMS College on Simulation. His research interests include discrete-event simulation, computational probability, queueing, and general theory for stochastic systems.