

ON AN INITIAL TRANSIENT DELETION RULE WITH RIGOROUS THEORETICAL SUPPORT

Hernan P. Awad

Department of Management Science
University of Miami
Coral Gables, FL 33124-6544, U.S.A.

Peter W. Glynn

Department of Management Science and Engineering
Stanford University
Stanford, CA 94305-4026, U.S.A.

ABSTRACT

We study an initial transient deletion rule proposed by Glynn and Iglehart. We argue that it has desirable properties both from a theoretical and practical standpoint; we discuss its bias reducing properties, and its use both in the single replication setting and in the multiple replications / parallel processing context.

1 INTRODUCTION

Let $Y = (Y(t) : t \geq 0)$ be a real-valued stochastic process, in which $Y(t)$ represents the output of a simulation at (simulated) time t . Suppose that Y has a steady-state, in the sense that there exists a (deterministic) constant α such that

$$\frac{1}{t} \int_0^t Y(s) ds \implies \alpha \quad (1)$$

as $t \rightarrow \infty$, where \implies denotes weak convergence. The quantity α is known as the steady-state mean of Y , and the problem of computing α via simulation is called the *steady-state simulation problem*.

In view of the law of large numbers (1), the most natural simulation-based estimator for α is the time-average, namely

$$\alpha(t) \triangleq \frac{1}{t} \int_0^t Y(s) ds.$$

One of the challenges in steady-state simulation is dealing with the fact that $\alpha(t)$ is typically a biased estimator of α . The source of this bias lies in the initialization of Y at time $t = 0$ using a distribution that is atypical of steady-state behavior (e.g., initializing a queue in the empty state): this atypical initialization makes the first portion of the simulated path unrepresentative of the steady-state behavior (what is often called the *initial transient problem*), which in turn induces bias in $\alpha(t)$.

The initial transient problem has been the focus of much study in the simulation literature. One of the major approaches to dealing with this problem is to discard the initial portion of the simulation output, since those observations are viewed as being heavily “contaminated” by the initial condition. An *initial transient deletion rule* specifies the mechanism to choose how much of the path should be discarded. (Other approaches to the initial transient problem include *perfect simulation* and the use of *low-bias* estimators.) In this paper we analyze an initial transient deletion rule proposed by Glynn and Iglehart (1987) in the Markov process setting, and argue that it has desirable properties both from the theoretical and practical standpoints.

The initial transient was recognized as a problem early on in the history of the simulation community, and a large number of initial transient deletion rules have been developed since then. Starting with simple heuristics like those of Conway (1963), Fishman (1971), Welch (1983) and Roth and Rutan (1985), the literature has steadily grown with the addition of increasingly sophisticated deletion procedures; see Wilson and Pritsker (1978b) for an early survey and Pawlikowski (1990) for a more recent survey in the queueing context; see also the list in Robinson (2002).

Various works have devised rules to test for the presence of an initial transient: examples are Schruben (1982), Schruben, Singh, and Tierney (1983), and Vassilacopoulos (1989); such tests can then be used to develop (rather sophisticated) initial bias detection rules, as is done for example in Jackway and de Silva (1992), Richet, Jacquet, and Bay (2003), Pawlikowski (1990), and Robinson (2002). Many other statistically-based deletion rules have been developed in the Markov chain Monte Carlo (MCMC) literature, since in that setting the initial transient is an even bigger problem than in steady-state simulation (in MCMC the interest is in sampling from the stationary distribution, rather than computing the steady-state mean). See Cowles and Carlin (1996) and Mengersen, Robert, and Guihenneuc-Jouyaux (1999) for a survey, discussion, and comparison of existing methods.

It is safe to say that the search for better initial transient deletion rules is open both in the steady-state simulation and MCMC communities: empirical studies have shown that the performance of most of the existing rules is typically not robust; see Gafarain, Ancker, and Morisaku (1978), Wilson and Pritsker (1978a), Linton and Harmonosky (2002), Cowles and Carlin (1996).

From a theoretical standpoint, the majority of the existing deletion rules are lacking rigorous support, in the sense that there are no provable guarantees on their performance over a sufficiently rich class of models. A major exception is an avenue of research that has focused on obtaining (computable) bounds on the rate of convergence to stationarity of Markov processes, from which deletion rules with provable performance guarantees can be developed; some noteworthy examples are Diaconis and Stroock (1991), Aldous (1991), Rosenthal (1995a, 1995b), Diaconis and Saloff-Coste (1993), Fill (1991), and Meyn and Tweedie (1994). For some models, sharp thresholds for the length of the warm-up period can be computed—see Aldous and Diaconis (1986), Rosenthal (1996). This approach is particularly suited to study reversible Markov chains, context in which most of its success has taken place. However, simulation models encountered in practice are rarely—essentially never—reversible. In the MCMC context this may not be a critical problem, since the stationary distribution is known up to a multiplicative scale factor, and one has the freedom to build a corresponding Markov chain that will have easier convergence diagnostics (e.g., restricting oneself to reversible dynamics). In contrast, in most steady-state simulation problems the model is given and (in principle) one has no freedom to modify its dynamics. Also, computing such bounds on convergence rates invariably requires detailed knowledge and analysis of the transition kernel of the Markov process (for example, in the discrete-time Markov chain context one needs to bound the second largest eigenvalue of the one-step transition matrix; see, e.g., Brémaud 1999). Such analysis imposes an extra cost to the simulationist, as most discrete-event simulations are implemented without need to calculate an explicit expression for the transition kernel.

Part of what makes the initial transient deletion rule of Glynn and Iglehart (1987) appealing is that it only involves analysis of the simulated output. The deletion period is a measurable function of the simulation output that is used universally across all positive recurrent irreducible discrete state-space Markov chains, and is not permitted to depend on the underlying transition kernel. This feature is of practical interest, since it allows the deletion rule to be fully automated, and easily incorporated into general-purpose simulation software. (One can say that this feature is what defines a truly statistically-based deletion rule.) Additionally, we argue here that there is rigorous support for this

deletion rule: it has provable bias-reducing properties, and one can show it does not delete “too much” of the output.

In Section 2, we describe in detail the above mentioned bias deletion rule, and a variant suited to the multiple replications setting. In Section 3, we describe the natural point estimators associated with applying the rule, and discuss their bias reducing properties. In Section 4, we discuss the implications in terms of completion time in the parallel simulation context. In Section 5, we discuss what are the possible advantages of using such a rule in the single replication context.

2 THE INITIAL BIAS DELETION RULE

The bias deletion rule of Glynn and Iglehart (1987) is designed for the Markov process setting. Specifically, we assume that $Y = (Y(t) : t \geq 0)$ has the representation $Y(t) = f(X(t))$, where $X = (X(t) : t \geq 0)$ is a continuous-time Markov chain living in discrete state space S . Further, we assume X is irreducible and positive recurrent, and denote by $\pi = (\pi(x) : x \in S)$ the (unique) stationary law of X .

The bias deletion rule is inspired by the following observation. Suppose one could generate a S -valued random variable Z independent of X and having distribution π , that is

$$P(Z = x) = \pi(x),$$

$x \in S$. Then, letting $T \triangleq \inf\{s \geq 0 : X(s) = Z\}$, the process $(X(T+s) : s \geq 0)$ is a stationary version of X . Hence, discarding all the observations prior to τ would completely eliminate the initial transient. Of course, in practice one cannot generate Z and T as above, since the distribution π is unknown. One can, however, estimate π from the simulation output; for example, a consistent estimate for π is given by the occupation measure π_t , where

$$\pi_t(x) = \frac{1}{t} \int_0^t I(X(s) = x) ds.$$

This estimate can then be used to generate Z and T as above. Hence, in the single replication case the bias deletion rule takes the following form:

- i.) Simulate $(X(s) : 0 \leq s \leq t)$.
- ii.) Put $Z \triangleq X(tU)$, where U is uniformly distributed on $(0, 1)$ and independent of X . (Note this implies $P(Z = x|X) = \pi_t(x)$.)
- iii.) Put $T(t) \triangleq \inf\{s \geq 0 : X(s) = Z\}$.
- iv.) Discard all observations prior to $T(t)$, keeping $(Y(s) : T(t) \leq s \leq t)$.

The rule can be easily adapted to the multiple replication setting. In this context m i.i.d. replicates of X , $\{X^i : 1 \leq i \leq m\}$, are simulated, and the rule takes the following form:

- i.) Simulate $(X^i(s) : 0 \leq s \leq t)$, $i = 1, \dots, m$.
- ii.) Generate rv's $\{U_1, \dots, U_m, R_1, \dots, R_m\}$, mutually independent and independent of $\{X^i : 1 \leq i \leq m\}$, with U_i uniformly distributed on $(0, 1)$ and R_i uniformly distributed on $\{1, \dots, m\} \setminus \{i\}$.
- iii.) Put $Z^i \triangleq \sum_{j \neq i} I(R_i = j) X^j(tU_i)$. (Note this implies $P(Z^i = x | X^1, \dots, X^m) = \pi_i^{-i}(x) \triangleq \frac{1}{m-1} \sum_{j \neq i} \frac{1}{t} \int_0^t I(X^j(s) = x) ds$.)
- iv.) Put $T^i(t) \triangleq \inf\{s \geq 0 : X^i(s) = Z^i\}$.
- v.) Discard the observed path of X^i prior to $T^i(t)$, keeping $(Y^i(s) : t \wedge T^i(t) \leq s \leq t)$, $i = 1, \dots, m$.

Observe this construction makes Z^i independent of X^i , which seems intuitively desirable.

One of the appealing features of this bias deletion rule is that the amount of data deleted from each simulation run typically does not grow unboundedly as the simulated time horizon t grows. Specifically, one has the following.

Proposition 1 *Suppose $|S| < \infty$. Then $(T(t) : t \geq 0)$ is bounded in probability. That is, for all $\varepsilon > 0$ there exists $c > 0$ such that $P(T(t) > c) < 1 - \varepsilon$, $t \geq 0$.*

The same result holds for $(T^i(t) : t \geq 0)$ in the multiple replication setting. In a future article, we will include the proofs of this and the remaining results in this paper.

We point out that this bias deletion rule does not involve unspecified parameters to be chosen or tuned by the simulationist. This is a desirable feature from a practical standpoint.

3 ASSOCIATED POINT ESTIMATORS

Once the bias deletion rule has been applied, perhaps the most natural point estimator for the steady-state mean α is the time-average over the remaining portion of the simulation run, given by

$$\alpha_1(t) \triangleq \frac{1}{t - T(t)} \int_{T(t)}^t f(X(s)) ds.$$

A minor concern when using this estimator is that the average is taken over an interval of random length $t - T(t)$, making $\alpha_1(t)$ a ratio of random variables; such ratio estimators are sometimes not very robust for small to moderate values of t . As an alternative, we consider an estimator that simulates the path of X beyond time t , until a path of length t is available for estimation (after applying the deletion rule):

$$\alpha_2(t) \triangleq \frac{1}{t} \int_{T(t)}^{t+T(t)} f(X(s)) ds.$$

Of course, this comes at the price of extra computational time to simulate X over the interval $(t, t + T(t))$; in particular, the total simulation run-length is a random variable, which may

be a disadvantage if one wants completion time guarantees. (Although, by virtue of Proposition 1, the extra simulation time needed is typically small.)

We stress that the caveats just raised on the use of these estimators are very minor; both $\alpha_1(t)$ and $\alpha_2(t)$ seem very reasonable alternatives in practice. Furthermore, they have good (and similar) theoretical properties, as we discuss below.

In the multiple replications setting, the counterparts to the above estimators are

$$\alpha_1^A(t) \triangleq \frac{1}{m} \sum_{i=1}^m \frac{I(T^i(t) \leq t)}{t - T^i(t)} \int_{T^i(t)}^t f(X(s)) ds,$$

and

$$\alpha_2^A(t) \triangleq \frac{1}{m} \sum_{i=1}^m \frac{1}{t} \int_{T^i(t)}^{t+T^i(t)} f(X(s)) ds.$$

Our main result summarizes the bias reducing properties of these estimators.

Theorem 1 *Suppose $|S| < \infty$. There exist constants $\gamma_1, \gamma_2, \gamma_1^A, \gamma_2^A$ such that*

$$E \alpha_i(t) - \alpha = \frac{\gamma_i}{t^2} + o(t^{-2}),$$

$i = 1, 2$ and

$$E \alpha_i^A(t) - \alpha = \frac{\gamma_i^A}{t^2} + o(t^{-2}),$$

$i = 1, 2$.

The proof uses the regenerative structure of CTMCs.

Noting that the corresponding bias expansion for the time average has the form

$$E \alpha(t) - \alpha = \frac{b}{t} + o(t^{-1})$$

(see, e.g., Glynn 1984), we see that use of the bias deletion rule leads to estimators with substantially smaller bias.

4 COMPLETION TIME GAINS IN THE PARALLEL SIMULATION SETTING

The use of multiple replications estimators is particularly suited to the parallel simulation setting, since each simulation run can be performed in a different processor, speeding up the completion time. In the presence of (massive) parallel processing capability, the reduction in bias achieved by applying the bias deletion rule can be exploited to attain significant savings in completion time, as we show below.

The time-average $\alpha(t)$ typically satisfies a central limit theorem (CLT) of the form

$$\sqrt{t}(\alpha(t) - \alpha) \implies \sigma N(0, 1), \quad (2)$$

as $t \rightarrow \infty$, where $N(0, 1)$ is a standard Gaussian rv and σ^2 is the so-called time-average variance constant (TAVC). The corresponding multiple replications estimator (with no bias deletion whatsoever) is given by

$$\alpha^A(t) \triangleq \frac{1}{m} \sum_{i=1}^m \frac{1}{t} \int_0^t f(X(s)) ds.$$

It satisfies a CLT with the same TAVC, namely

$$\sqrt{mt}(\alpha^A(t) - \alpha) \implies \sigma N(0, 1), \quad (3)$$

as $m, t \rightarrow \infty$, provided that the run length t increases fast enough compared to the number of runs m ; specifically, one needs $t/m \rightarrow \infty$ for (3) to hold (Glynn and Heidelberger 1991).

Relation (3) is often used to compute asymptotic confidence intervals for α , and the length of the confidence interval is proportional to $(mt)^{-1/2}$. Suppose one wants the length of the confidence interval to be of the order 10^{-n} . Then the requirement $t/m \rightarrow \infty$ implies that, as we demand more precision (i.e., as n increases), the simulation run length must increase slightly faster than 10^n (e.g., put $t = 10^{n(1+\varepsilon)}$ and $m = 10^{n(1-\varepsilon)}$, with $\varepsilon > 0$ small). The situation is different when the bias deletion rule is used, as the following result shows.

Theorem 2 Assume $|S| < \infty$. Then, for $i = 1, 2$

$$\sqrt{mt}(\alpha_i^A(t) - \alpha) \implies \sigma N(0, 1), \quad (4)$$

as $m, t \rightarrow \infty$, provided $t^3/m \rightarrow \infty$.

Hence, when using $\alpha_1^A(t)$ or $\alpha_2^A(t)$, the simulation run length need only increase slightly faster than $10^{n/2}$ (e.g., put $t = 10^{n(1/2+\varepsilon)}$ and $m = 10^{n(3/2-\varepsilon)}$, with $\varepsilon > 0$ small), to make the length of the confidence interval decrease as 10^{-n} .

When using $\alpha^A(t)$ or $\alpha_1(t)$ we can identify t with the completion time for the (overall) experiment. When using $\alpha_2^A(t)$ the completion time is $T_{comp} = \max\{t + T^i(t) : 1 \leq i \leq m\}$, but it can be shown that if m and t are chosen as above, then $T_{comp}/t \implies 1$ as $m, t \rightarrow \infty$. Thus, we see that use of the bias deletion rule allows the simulationist to better exploit the availability of (massive) parallel processing capability, achieving desired precision levels with smaller completion time, by using a larger number of shorter runs.

5 USE IN THE SINGLE REPLICATION SETTING

In the single replication setting, the bias reduction properties of $\alpha_1(t)$ and $\alpha_2(t)$ presented in Theorem 1 are not, by themselves, sufficient to claim that these estimators are preferable to the time average $\alpha(t)$. A more appropriate measure of an estimator performance in this setting is given by the mean-square error (MSE). In great generality it can be shown that MSE of the time-average estimator has an expansion of the form

$$E(\alpha(t) - \alpha)^2 = \frac{\sigma^2}{t} + \frac{\eta_0}{t^2} + O(e^{-\lambda t}),$$

for some constants $\lambda > 0$ and η_0 , where $\sigma^2 \triangleq 2 \int_0^\infty E_\pi(f(X(0)) - \alpha)(f(X(s)) - \alpha) ds$; see Awad and Glynn (2006). When the initial transient deletion rule is used, the resulting estimators enjoy similar MSE expansions:

Theorem 3 Suppose $|S| < \infty$. For $i = 1, 2$

$$E(\alpha_i(t) - \alpha)^2 = \frac{\sigma^2}{t} + \frac{\eta_i}{t^2} + O(t^{-3}),$$

as $t \rightarrow \infty$.

Explicit expressions can be derived for the constants η_1 and η_2 using the regenerative structure of CTMC's. However, in general it is not apparent how η_1 and η_2 compare to η_0 . In a realistic situation the simulationist will have no guarantee as to whether applying the initial transient deletion rule will lead to a reduction in MSE, compared to the time-average. In any case, the differences in MSE are at most second-order effects: the first-order term in the MSE expansions is the same for all three estimators.

In view of the above, one may be of the opinion that there is no need (or no point) in applying such initial transient deletion techniques in the single replication setting. In a typical simulation experiment, their effect on the MSE is negligible for large values of t , since the initial bias “dies out” faster than the variance of the time-average.

On the other hand, for any given finite t , the simulationist does not a priori know whether t is large enough to make the bias negligible. These bias deletion rules can act as a safeguard against situations in which the initial condition is so drastically atypical of steady-state behavior that the effects of the initial bias still have a significant effect on the time-average at the chosen run-length. Also, use of the bias deletion rule may serve as a way to assess whether the chosen time-horizon is sufficiently long to dampen the initial transient effects.

We illustrate the above discussion by considering an M/M/1 queueing process started with an unusually large queue length:

Proposition 2 Assume X represents the queue length process of an M/M/1 queue with arrival rate λ and service rate $\mu > \lambda$, with the initial condition $X(0) = x$

a.s. Suppose $f(y) = y$, $y \geq 0$ (so that α is the steady state mean queue length). If $t = x^{1+\varepsilon}$, where $0 < \varepsilon < 1/3$, then

$$x^{(1+\varepsilon)/2}|\alpha(t) - \alpha| \implies \infty$$

as $x \rightarrow \infty$, whereas

$$x^{(1+\varepsilon)/2}(\alpha_i(t) - \alpha) \implies \sigma N(0, 1)$$

as $x \rightarrow \infty$, for $i = 1, 2$. Also, if $t = x^{1-\varepsilon}$, then

$$T(t)/t \implies 1$$

as $x \rightarrow \infty$.

In this example, the first busy cycle is completely unrepresentative of steady state behavior, with unusually large queue lengths. That first busy-cycle is only a small portion of the run length when $t = x^{1+\varepsilon}$, but its effect is large enough that the time-average is far from α . The bias deletion rule, however, will typically delete virtually all of the first busy cycle, and the estimators will most likely not be affected by it. On the other hand, when $t = x^{1-\varepsilon}$ the run length is shorter than the time required to empty the queue. In this case, the bias deletion rule deletes most of the run length, which is a strong signal to the effect that the chosen run length is too small to dampen the initial transient effects.

REFERENCES

- Aldous, D. 1991. Meeting times for independent Markov chains. *Stochastic Processes and their Applications* 38:185–193.
- Aldous, D., and P. Diaconis. 1986. Shuffling cards and stopping times. *American Mathematical Monthly* 93:333–348.
- Awad, H., and P. W. Glynn. 2006. *ACM Transactions on Modelling and Computer Simulation*. To appear.
- Brémaud, P. 1999. *Markov chains, Gibbs fields, Monte Carlo Simulation and queues*. New York: Springer-Verlag.
- Conway, R. W. 1963. Some tactical problems in digital simulation. *Management Science* 10:47–61.
- Cowles, M. K., and B. P. Carlin. 1996. Markov chain Monte Carlo convergence diagnostics: a comparative review. *Journal of the American Statistical Association* 91 (434): 883–904.
- Diaconis, P., and L. Saloff-Coste. 1993. Comparison theorems for reversible markov chains. *The Annals of Applied Probability* 3.
- Diaconis, P., and D. Stroock. 1991. Geometric bounds for eigenvalues of Markov chains. *The Annals of Applied Probability* 1.
- Fill, J. A. 1991. Eigenvalue bounds on convergence to stationarity for nonreversible Markov chains, with an application to the exclusion process. *The Annals of Applied Probability* 1 (1): 62–87.
- Fishman, G. 1971. Estimating sample size in computer simulation experiments. *Management Science* 18:21–37.
- Gafarain, A., C. Ancker, and T. Morisaku. 1978. Evaluation of commonly used rules for detecting steady-state in computer simulation. *Naval Research Logistics* 25:511–529.
- Glynn, P. W. 1984. Some asymptotic formulas for Markov chains with applications to simulation. *Journal of Statistical Computation and Simulation* 19:97–112.
- Glynn, P. W., and P. Heidelberger. 1991. Analysis of initial transient deletion for replicated steady-state simulations. *Operations Research Letters* 10 (8): 437 – 443.
- Glynn, P. W., and D. L. Iglehart. 1987. A new bias deletion rule. In *Proceedings of the 1987 Winter Simulation Conference*, ed. A. Thesen, H. Grant, and D. Kelton.
- Jackway, P., and B. de Silva. 1992. A methodology for initialisation bias reduction in computer simulation output. *Asia Pacific Journal of Operational Research* 9:85–98.
- Linton, J. R., and C. M. Harmonosky. 2002. A comparison of selective initialization bias elimination methods. In *Proceedings of the 2002 Winter Simulation Conference*, ed. E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, 1951–1957.
- Mengersen, K., C. Robert, and C. Guhennec-Jouyau. 1999. MCMC convergence diagnostics: a review. In *Bayesian Statistics 6*, ed. J. Bernardo, J. Berger, A. Dawid, and A. Smith. Oxford University Press.
- Meyn, S., and R. Tweedie. 1994. Computable bounds for geometric convergence rates of Markov chains. *Annals of Applied Probability* 4:981–1011.
- Pawlikowski, K. 1990. Steady-state simulation of queuing processes: a survey of problems and solutions. *ACM Computing Surveys* 22:123–170.
- Richet, Y., O. Jacquet, and X. Bay. 2003. Automated suppression of the initial transient in Monte Carlo calculations based on stationarity detection using the Brownian bridge theory. In *Proceedings of the The 7th International Conference on Nuclear Criticality Safety (ICNC2003)*, 63. Tokai-mura, Japan.
- Robinson, S. 2002. A statistical process control approach for estimating the warm-up period. In *Proceedings of the 2002 Winter Simulation Conference*, ed. E. Yücesan, C.-H. Chen, J. L. Snowdon, and J. M. Charnes, 439–446.
- Rosenthal, J. S. 1995a. Minorization conditions and convergence rates for Markov chain Monte Carlo. *Journal of the American Statistical Association* 90:558–566.
- Rosenthal, J. S. 1995b. Rates of convergence for Gibbs sampling for variance component models. *Annals of Statistics* 23:740–761.

- Rosenthal, J. S. 1996. Analysis of the Gibbs sampler for a model related to James-Stein estimators. *Statistics and Computers* 6:269–275.
- Roth, E., and A. Rutan. 1985. A relaxation time approach for reducing initialization bias in Simulation. In *Proceedings of the 18th Annual Symposium on Simulation*, 189–203.
- Schruben, L. 1982. Detecting initialization bias in simulation output. *Operations Research* 30:569–590.
- Schruben, L., H. Singh, and L. Tierney. 1983. Optimal tests for initialization bias in simulation output. *Operations Research* 31:1167–1178.
- Vassilacopoulos, G. 1989. Testing for initialization bias in simulation output. *Simulation* 52 (4): 151–153.
- Welch, P. 1983. The statistical analysis of simulation results. In *The Computer Performance Modeling Handbook*, ed. S. Lavenberg, 268–328. New York: Academic Press.
- Wilson, J., and A. Pritsker. 1978a. Evaluation of startup policies in simulation experiments. *Simulation* 31:79–89.
- Wilson, J., and A. Pritsker. 1978b. A survey of research on the simulation startup problem. *Simulation* 31:55–59.

AUTHOR BIOGRAPHIES

HERNAN AWAD received B.Eng.Sci. and Industrial Engineer degrees from Universidad de Chile, and a Ph.D. from Stanford University. He is an assistant professor of Management Science at the University of Miami School of Business. He is a member of INFORMS and IEEE. His e-mail address is <h.awad@miami.edu>, and his web page is <moya.bus.miami.edu/~hawad>.

PETER W. GLYNN received his Ph.D. in Operations Research from Stanford University in 1982. He then joined the faculty of the University of Wisconsin at Madison, where he held a joint appointment between the Industrial Engineering Department and Mathematics Research Center, and courtesy appointments in Computer Science and Mathematics. In 1987, he returned to Stanford, where he is now the Thomas Ford Professor of Engineering in the Department of Management Science and Engineering. He also has a courtesy appointment in the Department of Electrical Engineering. He is a member of INFORMS and a fellow of the Institute of Mathematical Statistics. His e-mail address is <glynn@stanford.edu>, and his web page is <www.stanford.edu/dept/MSandE/people/faculty/glynn>.