

COMPUTATIONAL EFFICIENCY OF BATCHING METHODS

David Goldsman

School of Industrial and Systems Engineering
Georgia Institute of Technology
Atlanta, GA 30332-0205, U.S.A.

Bruce W. Schmeiser

School of Industrial Engineering
Purdue University
West Lafayette, IN 47907-1287, U.S.A.

ABSTRACT

This paper discusses the efficiency of various batching methods for estimating performance parameters from steady-state simulation output, e.g., the steady-state mean. Our primary focus is on issues related to computational and storage requirements of batching methods such as batch means, overlapping batch means, and standardized time series. We also examine the important question of determining the proper batch size for a given estimation problem — how much effort do batch-size algorithms require?

1 INTRODUCTION

Suppose we have steady-state (but not necessarily independent) output, Y_1, Y_2, \dots , arising from a single run of a simulation experiment. For instance, Y_i could represent the waiting time of the i th customer in a certain queueing model, or the transit time of the i th part through a manufacturing system. It is often of interest to estimate a performance measure θ , a property of the steady-state distribution function F_Y . The quantity θ is usually the mean, the variance, or a quantile, but in general can be any performance measure of the model being simulated. We briefly discuss estimation of θ before turning to our main topic of standard-error estimation.

1.1- Natural Estimators

The natural point estimator for θ , denoted by $\hat{\theta}$, is typically the sample mean $\bar{Y}_n = \sum_{k=1}^n Y_k/n$, the sample variance, or a simple function of the relevant order statistics, chosen to mimic the performance measure θ . Computational aspects associated with $\hat{\theta}$ differ little from the corresponding independent-observations case, simply because the natural point estimator is used in both cases. Not only are the natural estimators computationally reasonable, they perform well statistically.

Confusion sometimes arises about the use of the sample variance, S^2 , to estimate the process variance, $\text{Var}(Y_i)$. The sample variance performs well both statistically and computationally. As discussed in Wood and Schmeiser (1994), despite some bias, S^2 is consistent for $\text{Var}(Y_i)$ and has a smaller mean squared error (mse) than other so-called quadratic-form estimators. Hesitation about using S^2 arises because S^2/n , the usual estimator of $\text{Var}(\bar{Y}_n)$ for *independent* observations, is inappropriate for *dependent* observations. But the fact that the natural estimators are appropriate for estimating any θ , regardless of dependency, follows from the empirical distribution function converging to F_Y . Without this property, the ubiquitous histograms created by commercial simulation software would be somewhat misleading.

1.2- Unnatural Estimators

An alternative to the natural estimators are computationally intensive estimators. Bootstrapping and jackknifing are commonly used examples (see Efron 1982 and Efron and Gong 1983). These are sometimes inappropriate in the steady-state context where n is typically large. Because the effort to create n steady-state simulation observations is $O(n)$, the effort involved in calculating all estimators (both $\hat{\theta}$ and the associated standard error) should also be close to $O(n)$. (The notation $f(n) = O(g(n))$ means that $|f(n)/g(n)| \leq C$ as $n \rightarrow \infty$, for some constant C .) Otherwise, time spent computing the estimator could have been used to increase n . This is quite unlike the analysis of real-world data, where the cost of computing is essentially free compared to the cost of collecting data.

1.3- Sampling Error

A complementary concern of simulation output analysis is to estimate the sampling error of $\hat{\theta}$, that is, the error caused by the estimator's randomness. This gives the experimenter an idea of the degree to which

the point estimator $\hat{\theta}$ reflects the true but unknown parameter θ . The sole source of sampling error is the choice of random-number generator and its seed. Estimating the sampling error allows the practitioner to conclude which digits of $\hat{\theta}$ are meaningful. We measure the sampling error by the standard error, the standard deviation of $\hat{\theta}$ (or, almost equivalently, by the variance of $\hat{\theta}$). The standard error can be used to compute a confidence interval for θ or, less commonly, a tolerance interval, or to simply display meaningful digits of $\hat{\theta}$ (Song and Schmeiser 1994).

1.4 Batching

Batching is a classical methodology used in simulation output analysis for estimating the standard error of $\hat{\theta}$. For our purposes, the i th *batch* is composed of the observations $Y_i, Y_{i+1}, \dots, Y_{i+m-1}$ for $i = 1, 2, \dots, n-m+1$, where m is the batch size and n is the run length. The i th batch statistic $\hat{\theta}_i$ — an estimator for θ composed only of observations from the i th batch — is a miniature version of $\hat{\theta}$, again usually a sample mean, sample variance, or simple function of the relevant order statistics. When appropriate, a “grand” estimator can be used in the batch statistic; for example, centering a batch variance on the grand mean is better than centering it on the batch mean (Ceylan 1995).

1.5 Organization

This paper discusses the efficiency of three batching methods for estimating performance parameters from steady-state simulation output. Our primary focus is related to computational and storage requirements of the batching methods. To this end, in §2 we describe some desirable properties that estimators ought to have and summarize the approach for using batching to estimate standard error. §3 discusses some specific batching methods in the context of the steady-state mean estimation problem, where we are interested in estimating the standard error of the sample mean (the true mean’s natural estimator). The batching methods under study may be familiar to the reader — batch means, overlapping batch means, and standardized time series. In §4, we present theoretical considerations in determining the optimal batch size for a given estimation problem. §5 examines the important question of batch-size determination in practical applications — for example, how much effort do various batch-size algorithms require? We comment in §6 on the use of batching for estimating performance other than means. §7 gives some final remarks.

2 STANDARD-ERROR ESTIMATION

Batching methods are often used to estimate the standard error of an estimator. This section deals with the general problem of estimating standard error.

2.1 Properties of Estimators

We first list some properties that a good estimator should possess. These apply whether computing $\hat{\theta}$ to estimate a system performance measure θ or, as is our focus, computing a measure of experiment quality, such as the standard error. For notational convenience, we use $\hat{\theta}$ and θ here.

- **Statistical performance.** An estimator $\hat{\theta}$ should have low bias, variance, and mse. That is, $\text{bias} = \text{E}[\hat{\theta}] - \theta$, $\text{Var}(\hat{\theta})$, and $\text{mse} = \text{bias}^2 + \text{Var}$ should all be small. Ideally the estimator is consistent.
- **Ease of computation.** As discussed above, estimators should involve little more than $O(n)$ calculation operations.
- **Parsimonious storage requirements.** Data storage should be $O(1)$, and certainly not more than $O(n)$.
- **Ease of understanding.** All things being equal, intuitive “simple” estimators are preferred over more complicated estimators.
- **Numerical stability.** Large sample sizes should not cause underflow, overflow, or rounding errors.
- **User-specified parameters.** In the best of all worlds, the user should be required to specify no parameters, such as the number of batches or the batch size.
- **Amenability for use in algorithms.** In our context, we are concerned with how easily the estimators can be incorporated into batch-size determination procedures.

2.2 Batching Estimators

Armed with an idea of what properties an estimator for the standard error should have, we now give a top-level discussion of estimators arising from batching. We combine the generic batch statistics, $\hat{\theta}_i$, $i = 1, 2, \dots, n-m+1$, in order to estimate the variance of $\hat{\theta}$, or its square root, the standard error. In particular, the batch-statistics estimator for the variance of $\hat{\theta}$ is of the form

$$\widehat{\text{Var}}(\hat{\theta}) = \frac{\sum_A (\hat{\theta}_i - \hat{\theta})^2}{d|A|},$$

where A is an appropriate subset of $\{1, 2, \dots, n-m+1\}$, $|A|$ is the cardinality of A , and d is a constant chosen, for example, to yield $E[\widehat{\text{Var}}(\hat{\theta})] = \text{Var}(\hat{\theta})$ for independent data. For any batching method, the value of d is roughly n/m , because the underlying batching idea is that there is an asymptotic constant

$$\sigma^2 \equiv \lim_{n \rightarrow \infty} n \text{Var}(\hat{\theta}),$$

which implies that for large values of m and n the fraction $\text{Var}(\hat{\theta}_i)/\text{Var}(\hat{\theta})$ is approximately n/m .

Instead of using $\widehat{\text{Var}}(\hat{\theta})$ to estimate $\text{Var}(\hat{\theta})$, one can estimate directly the asymptotic constant σ^2 . The change in the associated batch-statistics estimator is easy — just multiply by n . Of course, this relationship is true only asymptotically, and may require very long runs before the necessary asymptotics apply.

3- SOME BATCHING METHODS

We illustrate batching methods by considering the problem of estimating the steady-state mean μ arising from a simulation process, Y_1, Y_2, \dots . The natural point estimator for μ is the grand sample mean $\hat{\theta} = \bar{Y}_n = \sum_{k=1}^n Y_k/n$. There are a number of popular batching estimators for $\text{Var}(\hat{\theta})$ in the literature. Song and Schmeiser (1993) discuss each of the three estimators described below as well as some others, giving their quadratic-form representations both algebraically and graphically.

3.1- Nonoverlapping Batch Means Estimator

The *nonoverlapping batch means* (NBM) estimator for $\text{Var}(\bar{Y}_n)$ is defined as

$$\hat{V}_N \equiv \frac{m \sum_{j=1}^b (\bar{Y}_{j,m} - \bar{Y}_n)^2}{n(b-1)},$$

where we divide the n observations into b adjacent, *nonoverlapping* batches, each of size m (assume $n = mb$), and where the j th nonoverlapping batch mean is

$$\bar{Y}_{j,m} \equiv \frac{1}{m} \sum_{k=1}^m Y_{(j-1)m+k},$$

$j = 1, 2, \dots, b$. Notice that \hat{V}_N is similar to the sample variance of the batch means. Also notice that, if we define batch estimators by $\hat{\theta}_i = \sum_{k=0}^{m-1} Y_{i+k}/m$, then the nonoverlapping batch means are $\bar{Y}_{j,m} = \hat{\theta}_{(j-1)m+1}$.

If the batch size m is known, then NBM requires $O(n)$ computation and $O(1)$ storage. Additional details on NBM and simple variants thereof are given in Schmeiser (1982), Carlstein (1986), Glynn and Whitt (1991), and Chien, Goldsman, and Melamed (1997).

3.2- Overlapping Batch Means Estimator

The *overlapping batch means* (OBM) estimator for $\text{Var}(\bar{Y}_n)$, proposed by Meketon and Schmeiser (1984), is defined as

$$\hat{V}_O \equiv \frac{m \sum_{i=1}^{n-m+1} (\hat{\theta}_i - \hat{\theta})^2}{n-m(n-m+1)},$$

where we divide the n observations into $n-m+1$ *overlapping* batches, each of size m , and where the i th overlapping batch mean is $\hat{\theta}_i = \sum_{k=0}^{m-1} Y_{i+k}/m$, for $i = 1, 2, \dots, n-m+1$.

OBM is identical to NBM, except that it uses all $n-m+1$ batches and rescales accordingly. Computationally, if m is known, then OBM is $O(n)$, rather than $O(nm)$, if each new batch mean is obtained by adding and subtracting one observation from the previous batch of size m . But to achieve $O(n)$ computation, OBM requires saving the previous m observations, so it needs $O(m)$ storage. OBM and NBM have about the same bias as estimators of $\text{Var}(\bar{Y}_n)$, but for large m and n/m , OBM's variance (and hence its mse) is about 1/3 smaller. See Goldsman and Meketon (1986) and Song and Schmeiser (1995) for details.

Welch (1987) discusses partially overlapping batch means and Fox, Goldsman, and Swain (1991) propose spaced batch means, both of which are other variations on the batches to use. All are $O(n)$ in computation, but OBM is statistically the most efficient.

3.3- Standardized Time Series Estimator

Schruben (1983) defines the *standardized time series* from (nonoverlapping) batch i , $i = 1, 2, \dots, b$, as

$$T_{i,m}(t) \equiv \frac{\lfloor mt \rfloor (\bar{Y}_{i, \lfloor mt \rfloor} - \bar{Y}_{i,m})}{\sigma \sqrt{m}} \quad (0 \leq t \leq 1),$$

where $\lfloor \cdot \rfloor$ is the floor function, and

$$\bar{Y}_{i,j} \equiv \frac{1}{j} \sum_{k=1}^j Y_{(i-1)m+k},$$

for $1 \leq i \leq b$, $1 \leq j \leq m$.

The most well-known standardized time series estimator is based on the area under each $T_{i,m}(t)$ functional,

$$\begin{aligned} A_i &\equiv \frac{\sqrt{12}\sigma}{m} \sum_{k=1}^m T_{i,m}(k/m) \\ &= \frac{\sqrt{12}}{m^{3/2}} \sum_{k=1}^m \left(\frac{m+1}{2} - k \right) Y_{(i-1)m+k}, \end{aligned}$$

for $i = 1, 2, \dots, b$. The *area* estimator for $\text{Var}(\bar{Y}_n)$ is given by

$$\hat{V}_A \equiv \frac{1}{nb} \sum_{i=1}^b A_i^2.$$

If the batch size m is known, then the area estimator requires $O(n)$ computation and $O(1)$ storage.

The area estimator is one of many possible standardized time series estimators. It has some very nice asymptotic bias and variance properties, but is known to converge rather slowly with respect to these properties (Sargent, Kang, and Goldsman 1992). A number of variants of the area estimator try to avoid these slow-convergence problems, e.g., the weighted area estimator of Goldsman, Meketon, and Schruben (1990), the Cramér-von Mises estimator of Goldsman, Kang, and Seila (1997), and the L_p -norm estimator of Tokol et al. (1997).

4- OPTIMAL BATCH SIZE

Batch size selection (or, for nonoverlapping batches, the number of batches) is a problem that software vendors have typically solved by asking the practitioner to choose a value—a potentially problematic approach. Research into automating this choice has been in the context of confidence-interval procedures and mse-optimal standard-error estimation. All of the research has been for NBM or OBM estimators.

In the context of confidence-interval procedures, Fishman (1978), Law and Carson (1979), and Fishman and Yarberrry (1997) provide specific procedures for NBM. Alexopoulos, Fishman, and Seila (1997) consider computational issues for NBM confidence-interval procedures.

An optimal batch size for confidence-interval procedures is difficult to define, however, because the performance of these procedures is inherently multi-dimensional. Schmeiser (1982) considers four relevant performance properties: probability of covering θ , expected half width, variance of the half width, and the probabilities of covering points $\theta_0 \neq \theta$. He argues, without providing a procedure, that if n is fixed and only the first two criteria are deemed important, as is common in the research literature, then there is little statistical reason to use more than thirty NBM batches and substantial incentive to obtain at least ten NBM batches. Similar arguments suggest that the OBM batch size be chosen to be 1/7 to 1/20 of n . More batches are useful to improve the latter two criteria, including obtaining the consistency needed for the asymptotic results underlying sequential methods. In addition, Nelson (1989) argues that more NBM batches are needed while applying control variates, even if only the first two criteria are considered. Unlike algorithms, the informal guidelines are $O(1)$ in both computation and space; the cost is that the choice must be either hard coded or left to the practitioner.

Standard-error estimation is simpler because it allows a single criterion. The usual criterion is the mse of $\widehat{\text{Var}}(\hat{\theta})$ as an estimator of $\hat{\theta}$. If $\hat{\theta}$ is the sample mean, then asymptotic arguments lead to an mse-optimal batch size

$$\vec{m}^* = \left(2n \frac{c_b^2}{c_v} \frac{\gamma_1}{\gamma_0} \right)^{1/3} + 1,$$

where c_b^2 and c_v are the known bias constant and variance constant of the standard-error estimator, $\gamma_j = \sum_{h=-\infty}^{\infty} |h|^j \text{Corr}(Y_i, Y_{i+h})$, and the additive constant is chosen to yield a batch size of one for independent data (Goldsman and Meketon 1986, Song 1988, Song and Schmeiser 1995).

5- DETERMINING BATCH SIZE

Given the asymptotic results of the previous section, the problem of estimating the mse-optimal batch size when θ is the mean can be reduced to estimating γ_1/γ_0 . This ratio can be interpreted as the center of gravity of the non-negative lags of the autocorrelationogram. Two methods have been proposed for estimating the mse-optimal batch size via the center of gravity.

Song (1996) estimates autocorrelations explicitly, using tests of hypotheses to determine which autocorrelations to ignore as negligible. These estimated autocorrelations are then used to estimate the center of gravity. Although this technique works well on the battery of examples presented in the paper, it may require more than $O(n)$ work to estimate the required correlations. Also, the method depends critically on two user-supplied parameters, which must be chosen carefully.

Pedrosa and Schmeiser (1994), based on results in Pedrosa (1994) and Pedrosa and Schmeiser (1993), estimate the center of gravity implicitly. The method, named 121-OBM, is specific to the use of the OBM estimator, which is called four times in three steps. In the first step, OBM is called with a large batch size, with the purpose being to estimate, quite crudely, the sum of the autocorrelations γ_0 . Based on the result from Step 1, Step 2 estimates the center of gravity as a simple ratio of two OBM estimates, using two different batch sizes on the same data. Step 3 then estimates the variance of the sample mean using the batch size implied by the result of Step 2 and the asymptotic batch-size formula.

Step 2 deserves some discussion, since this is where the issue of computational efficiency arises. Avoiding explicit autocorrelation calculations allows $O(n)$ computation in an algorithm where much of the logic is buried in the OBM algorithm, which is called

as a subprogram. The key thought is to notice that the lag window of the OBM estimator, which is the same as the Bartlett spectral estimator, is a triangle. Consider superposing two such lag windows, one using some batch size m and another using $m + 1$. Think about the area within the second but not in the first. Corresponding to lag h is an area proportional to h . So the difference of the two OBM estimators, scaled appropriately, automatically weights each autocorrelation ρ_h by h , as needed in γ_1 , but without calculating ρ_h .

Both Song's method and 121-OBM assume that the data are available as n discrete data points in a vector that can be called repeatedly. Song's method requires the user to provide an algorithm parameter, while 121-OBM does not. On the other hand, Song's method could be adapted to any batching estimator, whereas 121-OBM is based on a property offered only by the OBM estimator.

6 OTHER PERFORMANCE MEASURES

Schmeiser, Avramidis, and Hashem (1990) and Ceylan (1995) consider performance measures other than the mean. Overlapping batch variances (OBV) require $O(n)$ computation and $O(m)$ space. Overlapping batch quantiles require $O(n \ln(m))$ computation if implemented carefully (as in Hashem and Schmeiser 1994). The optimal batch size, even in the asymptotic case, is unknown for these performance measures. Heuristic arguments and empirical evidence, however, suggest that the optimal mse-batch size grows with $n^{1/3}$. Ceylan (1995) shows that the OBV (or NBV) optimal-mse batch size is closely related to the optimal OBM (or NBM) batch size for the process of squared observations, Y_i^2 .

7 CONCLUSIONS

Some interesting computational problems for batching remain unsolved. Central to these is that the research literature has assumed that the n observations Y_i are available in a vector, either in memory or stored in a data base. Of course, in practice, n might not be known *a priori*, in which case the size of the vector is unknown. Because n is not known, an appropriate value for m is also unknown. So we need a method that accepts observations sequentially, accumulates some sums in $O(1)$ space and $O(n)$ time, dynamically increases m as is appropriate to the current run length n , and computes standard-error estimates as needed using the current value of m . The natural approach would be something like Fishman's (1978) idea of doubling batch sizes for NBM. Ideally, a computationally efficient algorithm can be developed that

is based on a more statistically efficient estimator.

ACKNOWLEDGMENTS

We thank Gamze Tokol for useful discussions. David Goldsman's work was supported by National Science Foundation Grant DMI-96-22269 to Georgia Tech. Bruce Schmeiser's was supported in part by National Science Foundation Grant DDM-93-00058 to Purdue University.

BIBLIOGRAPHY

- Alexopoulos, C., G. S. Fishman, and A. F. Seila. 1997. Computational experience with the batch means method. In *Proceedings of the Winter Simulation Conference*, ed. S. Andradóttir, K. Healy, D. Withers, and B. L. Nelson. Piscataway, NJ: IEEE.
- Carlstein, E. 1986. The use of subsamples for estimating the variance of a general statistic from a stationary sequence. *Annals of Statistics* 14: 1171–1179.
- Ceylan, D. 1995. *Variance and quantiles in dynamic-system performance: point estimation and standard errors*. Ph.D. Thesis, School of Industrial Engineering, Purdue Univ., West Lafayette, IN.
- Chien, C., D. Goldsman and B. Melamed. 1997. Large-sample results for batch means. *Management Science*, forthcoming.
- Efron, B. 1982. *The Jackknife, the Bootstrap and Other Resampling Plans*. Philadelphia: SIAM.
- Efron, B. and G. Gong. 1983. A leisurely look at the bootstrap, the jackknife, and cross-validation. *American Statistician* 37: 36–48.
- Fishman, G. S. 1978. Grouping observations in digital simulation. *Management Science* 24: 510–521.
- Fishman, G. S. and L. S. Yarberr. 1997. An implementation of the batch means method. *INFORMS Journal on Computing*, forthcoming.
- Fox, B. L., D. Goldsman and J. J. Swain. 1991. Spaced batch means. *Operations Research Letters* 10: 255–266.
- Glynn, P. W. and W. Whitt. 1991. Estimating the asymptotic variance with batch means. *Operations Research Letters* 10: 431–435.
- Goldsman, D., K. Kang, and A. F. Seila. 1997. Cramér-von Mises variance estimators for simulations. *Operations Research*, forthcoming.
- Goldsman, D. and M. Meketon. 1986. A comparison of several variance estimators for stationary increment stochastic processes. Technical Report, School of ISyE, Georgia Tech, Atlanta, GA.
- Goldsman, D., M. Meketon, and L. Schruben. 1990. Properties of standardized time series weighted

- area variance estimators. *Management Science* 36: 602–612.
- Hashem S. and B. W. Schmeiser. 1994. Algorithm 727: Quantile estimation using overlapping batch statistics. *ACM Transactions on Mathematical Software* 20: 100–102.
- Law, A. M. and J. S. Carson. 1979. A sequential procedure for determining the length of a steady-state simulation. *Operations Research* 27: 1011–1025.
- Meketon, M. S. and B. W. Schmeiser. 1984. Overlapping batch means: something for nothing? In *Proceedings of the Winter Simulation Conference*, ed. S. Sheppard, U. Pooch and C. D. Pegden, 227–230. Piscataway, NJ: IEEE.
- Nelson, B. L. 1989. Batch size effects on the efficiency of control variates in simulation. *European Journal of Operational Research* 43: 184–196.
- Pedrosa, A. 1994. *Automatic batching in simulation output analysis*. Ph.D. Thesis, School of Industrial Engineering, Purdue Univ., West Lafayette, IN.
- Pedrosa, A. and B. W. Schmeiser. 1993. Asymptotic and finite-sample correlations between obm estimators. In *Proceedings of the Winter Simulation Conference*, ed. G. W. Evans, M. Mollaghasemi, E. C. Russell, and W. E. Biles, 481–488. Piscataway, NJ: IEEE.
- Pedrosa, A. and B. W. Schmeiser. 1994. Estimating the variance of the sample mean: Optimal batch-size estimation and 1-2-1 overlapping batch means. Technical Report SMS94-3, School of Industrial Engineering, Purdue Univ., W. Lafayette, IN.
- Sargent, R. G., K. Kang and D. Goldsman. 1992. An investigation of finite-sample behavior of confidence interval estimates. *Operations Research* 40: 898–913.
- Schmeiser, B. W. 1982. Batch size effects in the analysis of simulation output. *Operations Research* 30: 556–568.
- Schmeiser, B. W., T. Avramidis and S. Hashem. 1990. Overlapping batch statistics. In *Proceedings of the Winter Simulation Conference*, ed. O. Balci, R. P. Sadowski, and R. E. Nance, 395–398. Piscataway, NJ: IEEE.
- Schruben, L. W. 1983. Confidence interval estimation using standardized time series. *Operations Research* 31: 1090–1108.
- Song, W. T. 1988. *Estimators of the variance of the sample mean: quadratic forms, optimal batch sizes, and linear combinations*. Ph.D. Thesis, School of Industrial Engr., Purdue Univ., W. Lafayette, IN.
- Song, W. T. 1996. On the estimation of optimal batch sizes in the analysis of simulation output analysis. *European Journal of Operational Research* 88: 304–309.
- Song, W. T. and B. W. Schmeiser. 1993. Variance of the sample mean: Properties and graphs of quadratic-form estimators. *Operations Research* 41: 501–517.
- Song, W. T. and B. W. Schmeiser. 1994. Reporting the precision of simulation experiments. In *New Directions in Simulation for Manufacturing and Communications*, ed. S. Morito, H. Sakasegawa, K. Yoneda, M. Fushimi, and K. Nakano, 402–407. Tokyo: Operations Research Society of Japan.
- Song, W. T. and B. W. Schmeiser. 1995. Optimal mean-squared-error batch sizes. *Management Science* 41: 110–123.
- Tokol, G., D. Goldsman, D. H. Ockerman, and J. J. Swain. 1997. Standardized time series L_p -norm variance estimators for simulations. *Management Science*, forthcoming.
- Welch, P. D. 1987. On the relationship between batch means, overlapping batch means and spectral estimation. In *Proceedings of the Winter Simulation Conference*, ed. A. Thesen, H. Grant and W. D. Kelton, 309–317. Piscataway, NJ: IEEE.
- Wood, D. Ceylan and B. W. Schmeiser. 1994. Consistency of overlapping batch variances. In *Proceedings of the Winter Simulation Conference*, ed. J. Tew, S. Manivannan, D. Sadowski, and A. F. Seila, 316–319. Piscataway, NJ: IEEE.

AUTHOR BIOGRAPHIES

DAVID GOLDSMAN is an associate professor in the School of Industrial and Systems Engineering at Georgia Tech. He received his Ph.D. from the School of OR&IE at Cornell University in 1984. His research interests include simulation output analysis and ranking and selection. He is an active participant in the Winter Simulation Conference, having been the Program Chair in 1995.

BRUCE W. SCHMEISER is a professor in the School of Industrial Engineering at Purdue University. He received his Ph.D. from the School of ISyE at Georgia Tech in 1975. His primary research interests are the probabilistic and statistical aspects of digital-computer stochastic simulation. He is an active participant in the Winter Simulation Conference, including being Program Chair in 1983 and chairing the Board of Directors during 1988–1990. He is also active in the American Statistical Association, IIE, INFORMS, and Omega Rho, The Operations Research Honor Society.