Guest Editorial

This special issue of *ACM TOMACS* on “Simulation Optimization” could perhaps be subtitled, “What was old is new (and better) again.” The four articles in this issue illustrate how some of the very first ideas in simulation optimization have been refined, extended and completely reworked in an ongoing attempt to bring rigor, and practicality, to an important topic.

Simulation optimization means searching for the settings of controllable decision variables that yield the maximum or minimum expected performance of a stochastic system that is represented by a simulation model. Perhaps the most natural approach is to try to import ideas from deterministic optimization into the stochastic setting by taking into account the uncertainty from estimation. For example, one of the earliest proposals for doing this, statistical response surface methodology (RSM), combines regression and gradient search (see, for instance, Myers and Montgomery [2002]).

A second early approach was to apply hill-climbing heuristics to simulation problems; Farrell [1977] surveyed these ideas. Hill-climbing heuristics had shown robust performance in tough deterministic optimization problems, and the simulation problems were at least as difficult because little or nothing might be known about the underlying response surface. Techniques for statistical ranking and selection (R&S) (see, for instance, Bechhofer et al. 1995), invented in the 1950’s, proved useful in helping these heuristics make local improvement decisions in the presence of sampling variability.

Stochastic approximation (SA) has been around for over half a century now [Robbins and Munro 1951; Kiefer and Wolfowitz 1952], but it is still an active area of research [Kushner and Yin 1997]. In terms of simulation optimization, a primary barrier has been the availability of gradients, especially for high-dimensional problems (lots of decision variables) motivated by complex stochastic systems. Thus, gradient estimation dominated research from the 1980s onward.

In this volume, we revisit many of these concepts in four articles that emphasize the opportunities, and difficulties, inherent in simulation experiments.

In the spirit of RSM, Homem-de-Mello’s “Variable-Sample Methods for Stochastic Optimization” explicitly combines deterministic optimization

Authors’ addresses: M. C. Fu, The Robert H. Smith School of Business and Institute for Systems Research, University of Maryland, College Park, MS 20742; email: mfu@isr.umd.edu; B. L. Nelson, Department of Industrial Engineering and Management Sciences, Northwestern University, 2145 Sheridan Road, Evanston, IL 60208-3119; email: nelsonb@northwestern.edu.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2003 ACM 1049-3301/03/0400-0105 \$5.00

techniques within the setting of performance estimation from sampling. This approach, called “sample path optimization” or “sample average approximation” in the literature, has been well developed over the past decade or so, but it has been based on the assumption that a very large one-time sample is taken, on which the deterministic optimization is also carried out a single time. In contrast, this article considers a sequential scheme in which the deterministic optimization problem is solved repeatedly using an estimate of the objective function based on *different* (and thus relatively smaller) samples at each iteration of the algorithm. A general framework is established for providing convergence guarantees and estimation error bounds for such procedures, thus allowing deterministic algorithms to be employed in a rigorous, adaptive manner. A variant of random search is used to illustrate the framework.

In “Discrete-Event Simulation Optimization Using Ranking, Selection and Multiple Comparison Procedures: A Survey,” Swisher et al. provide an up-to-date reference on recent advances in R&S and multiple-comparison procedures specifically developed for use in computer simulation experiments; they address issues such as variance reduction, steady-state simulation and large numbers of alternatives. R&S procedures are multistage experiment designs that provide statistical guarantees of finding the best among a finite number of alternatives, whereas multiple-comparison procedures provide statistical inference about the sizes of the differences among the alternatives. The authors provide a useful table of recommended procedures by problem type.

R&S procedures can be considered optimization techniques in their own right, when the number of alternatives is relatively small, but, more recently, they have once again been used to support general-purpose optimization algorithms, such as the ones in “A Combined Procedure for Optimization via Simulation,” by Pichitlamken and Nelson. In this article, a framework and several specific algorithms integrate ideas from adaptive random search, R&S, and hill-climbing heuristics into a provably convergent algorithm for discrete-decision-variable problems subject to linear inequality constraints. As part of this effort, a R&S procedure is designed with optimization in mind to make local improvement decisions efficiently, and correctly.

As alluded to earlier, when applying SA algorithms to simulation optimization, gradient estimation is key. “Two-Timescale Simultaneous Perturbation Stochastic Approximation Using Deterministic Perturbation Sequences,” by Bhatnagar et al. considers the setting where only estimates of the objective function itself are available from a stochastic simulation model; that is, direct gradient estimates using techniques such as perturbation analysis are unavailable. If the problem has a large number of parameters to optimize, traditional finite difference schemes can be computationally impractical, so gradients based on perturbing the objective function with respect to the parameters simultaneously, but *randomly*, have proven to be a viable alternative. The article investigates the idea of replacing random independent and identically distributed perturbation sequences with *cyclic deterministic* sequences that preserve the averaging properties of the gradient estimates required for establishing theoretical convergence. Similar in spirit to quasi-Monte Carlo, where low discrepancy is the objective, a type of decrease in variability is being sought.

We conclude with a brief explanation of the editorial process, especially in light of the outcome that each of the co-editors ended up with an accepted article in the issue. Workload for the submitted papers was divided between the two of us according to topic area, but an article in which one of us was a co-author was handled entirely by the other co-editor. Every article was sent to at least two anonymous referees, with the respective co-editor acting as the associate editor handling the article.

ACKNOWLEDGMENT

We wish to acknowledge the efforts of the referees for their conscientious and timely reports, and the Managing Editor John Konkle for his support throughout the process. Lastly, we thank the Editor David Nicol for inviting us to organize and put together this special issue.

Michael Fu
University of Maryland, College Park, MD
Barry Nelson
Northwestern University, Evanston, IL

REFERENCES

- BECHHOFFER, R. E., SANTNER, T. J., AND GOLDSMAN, D. M. 1995. *Design and Analysis of Experiments for Statistical Selection, Screening, and Multiple Comparisons*. Wiley, New York.
- FARRELL, W. 1977. Literature review and bibliography of simulation optimization. In *Proceedings of the 1977 Winter Simulation Conference*. 117–124.
- KIEFER, J. AND WOLFOWITZ, J. 1952. Stochastic estimation of the maximum of a regression function, *Ann. Math. Stat.* 23, 462–466.
- KUSHNER, H. J. AND YIN, G. G. 1997. *Stochastic Approximation Algorithms and Applications*, Springer-Verlag, New York.
- MYERS, R. H. AND MONTGOMERY, D. C. 2002. *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*, 2d ed. Wiley, New York.
- ROBBINS, H. AND MONRO, S. 1951. A stochastic approximation method. *Ann. Math. Stat.* 22, 400–407.